Title
Timing, uncertainty, and opportunity cost: Lessons for ecosystem modification on the Colorado river

Journal
Journal of Environmental Economics and Management, 138

ISSN
0095-0696

Authors
Donovan, Pierce
Bair, Lucas S
Reimer, Matthew N
et al.

Publication Date
2026-06-01

DOI
10.1016/j.jeem.2026.103358

Timing, uncertainty, and opportunity cost: Lessons for ecosystem modification on the Colorado River

Pierce Donovan[†], Lucas S. Bair[††], Matthew N. Reimer[†††], Michael R. Springborn[‡], and Charles B. Yackulic[††]

April 2026

Abstract

While conservation goals have long been pursued through traditional species-augmenting actions, a broader set of episodic ecosystem modification (EEM) actions, such as hydropower dam releases, prescribed fire, and beach nourishment, is garnering attention. EEM actions face several implementation challenges stemming from high opportunity costs, delayed effect mechanisms, reliance on monitoring for deployment timing, and outcome uncertainty due to infrequent use. In this paper, we study the use of EEM actions in the form of designer flows—ecologically-motivated releases of water into regulated river segments—to maintain a viable population of a threatened native fish species in the Colorado River. We demonstrate how the cost-effectiveness of EEM actions can be hampered by the complex and delayed effects on species viability, but enhanced through targeted monitoring for timing deployment and experimentation for reducing uncertainty about effectiveness.

Keywords: ecosystem modification, designer flows, hydropower, Glen Canyon Dam, Colorado River, integrated assessment, adaptive management, value of information, population viability, shadow value viability (JEL: C61, Q25, Q47, Q57)

[†]Department of Economics, University of Nevada, Reno
Corresponding Author, `pierce.donovan@unr.edu`; `piercedonovan.github.io`
[††]U. S. Geological Survey, Southwest Biological Science Center
[†††]Department of Agricultural and Resource Economics, University of California, Davis
[‡]Department of Environmental Science and Policy, University of California, Davis
We are grateful for discussion with the scientists of the USGS Southwest Biological Science Center and comments from audiences at the annual conference of the Western Economics Association, the Royal University of Agriculture, and the Royal University of Phnom Penh. We also thank David Finnoff and three anonymous reviewers for their valuable suggestions for this paper. The Glen Canyon Dam Adaptive Management Program provided financial support for this project. Any use of trade, firm, or product names is for descriptive purposes only and does not imply endorsement by the U.S. Government.

1 Introduction

Conservation managers have access to an ever-growing toolkit of environmental management actions. In addition to a longstanding set of traditional actions—such as species harvest or stocking—a broader set of ecosystem modification actions is emerging that can achieve conservation goals. Many of these newer actions are one-time measures and are already relatively well-studied, such as dam removal (Lewis et al., 2008; Jaeger, 2025) and introducing a non-native species for biocontrol of an invader (Higgins et al., 1997; Fenichel et al., 2010; Paul et al., 2020). In contrast, another set of ecosystem modification actions is dynamic, involving episodic use over time—for example, prescribed fire (Hashida et al., 2025), selective flooding of land areas (Garnache, 2015), beach nourishment (Smith et al., 2009), and dam releases that modify temperature, facilitate fish migration, or disadvantage invasive species (Bryan et al., 2013; Crespo et al., 2019). These may provide advantages relative to traditional actions in certain circumstances, yet their efficient use in conservation has received less attention from economists.

These episodic ecosystem modification (EEM) actions share several common implementation challenges. They often require implementation at scale, introducing large opportunity costs that state-contingent benefits may not always justify. They are often used in combination with traditional actions, creating benefit in understand their efficient joint deployment. Their effects are typically delayed, increasing both sensitivity to discounting assumptions and potential constraints to their use as a rapid response measure. Their episodic use at irregular intervals makes questions of timing central; optimal timing likely depends on environmental conditions, raising the importance of monitoring. Their novelty heightens parametric uncertainty over efficacy, underscoring the importance of learning by doing. Finally, their implementation must often weigh future threats and stressors that are expected to intensify rather than diminish.

In this paper, we introduce a general model that illustrates the common implementation challenges of EEM actions in conservation decision-making. We illustrate its use with an application to the conservation of humpback chub (*Gila cypha*), a threatened fish species, which is under competitive and predatory pressure from non-native trout species in a large dam-regulated ecosystem downstream of the Glen Canyon Dam on the Colorado River. Here, traditional actions in the form of trout removals have been used to control the non-native population (Coggins et al., 2011). We consider the additional use of EEM actions via designer flows—ecologically-motivated dam releases that, in this case, are designed to disadvantage non-native trout. This involves temporary up- and down-ramping of the flow rate from Glen Canyon Dam—at the cost of optimal hydropower generation—

to selectively attract juvenile trout to higher elevations during a flooding period and strand them as river elevation recedes (Giardina et al., 2024).

The conservation of humpback chub provides a compelling context for illustrating and evaluating the implementation challenges of EEM actions. First, the effective scale of designer flows requires significant foregone hydropower generation. Second, managers deploy trout removal in combination with designer flows over time. Third, the delayed conservation effect of ecosystem modification interacts with discounting to influence incentives for its use. Fourth, without monitoring of environmental conditions (stochastic trout recruitment) the population vulnerable to designer flows is unknown. Fifth, experimentation with ecosystem modification—i.e., learning-by-doing—can reduce parametric uncertainty about designer flow efficacy. Finally, the balance between conservation actions changes as threats to the humpback chub worsen due to the emergence of an even more voracious non-native predator.

We solve for an optimal feedback policy for ongoing humpback chub viability at minimal cost, mapping any potential state of the system into an ideal deployment of designer flows and trout removals. Our results inform a detailed and actionable policy for the Glen Canyon Dam Adaptive Management Program, which can be adapted to other regulated river systems based on the presence or absence of certain critical features. We demonstrate how the cost-effectiveness of EEM actions can be hampered by the complex and delayed effects on species viability, but enhanced through targeted monitoring for timing deployment and experimentation for reducing uncertainty about effectiveness.

In the next section, we provide a general conceptual model of EEM and illustrate three salient examples in natural resource management: designer flows, beach nourishment, and prescribed fire. In Section 3, we describe the details of our integrated assessment of humpback chub conservation in the Colorado River. Sections 3.1, 3.2, and 3.3 present our setting, bioeconomic model and optimization problem, and computational solution method, while Sections 3.4 and 3.5 specify opportunity costs of altered hydropower generation and an experimentation program. Section 4 explores the conditions under which ecosystem modification is valuable with several numerical simulations. These include a limited baseline optimization (4.1) and extensions including the effect of enhanced ecosystem monitoring (4.2), ecological changes (4.3), and experimentation (4.4). Section 5 concludes with further discussion regarding ecosystem modification in our setting.

2 A conceptual model of episodic ecosystem modification

In this section, we provide a simple conceptual model that highlights the unique features and consequences of EEM actions, using a stylized version of the non-native species management problem in Section 3 as a running example. We then discuss how this model applies to other EEM actions, namely prescribed fire and beach nourishment.

We begin by defining a focal state variable of interest, S_t, that represents a component of the environment that supports amenities or disamenities (e.g., a biological population). To illustrate, let S_t denote a population of a damaging, non-native fish species in year t located within a dam-regulated river. The evolution of this state is given by

$$S_{t+1} = g_S\left(S_t, X_t, A_t\right). \tag{1}$$

The evolution of S_t depends on two endogenous factors. First, an intermediate state X_t reflects a relevant component of the ecosystem in which S_t is managed, such as the river stage. Second, a traditional management action A_t directly affects S_t, such as culling the non-native species.

The dynamics of the intermediate state variable, X_t, depend on an EEM action, M_t, and stochastic shock, ε_t,

$$X_{t+1} = g_X\left(X_t, M_t, \varepsilon_t\right). \tag{2}$$

M_t, for example, could take the form of altered dam releases—designer flows—that strand the non-native species. In our setting, this process selectively targets juvenile fish, thereby delaying the EEM action's effect on future damage from adult populations and creating a difference in the relative timing of the impacts of the two management actions. Although the absolute timing of these impacts is application-specific, the traditional action generally has an immediate impact, whereas the EEM action's effect is relatively delayed because of the multi-step pathway through which M_t ultimately impacts S_t.

We highlight a specific class of exogenous shocks specific to the evolution of X_t that is relevant for EEM actions. The ε_t term represents a potentially *observable* shock that, when known, provides salient information for choosing whether to implement M_t. For example, ε_t could represent the juvenile recruitment of the non-native species, i.e., the subpopulation vulnerable to stranding. Thus, observing ε_t through targeted monitoring effort would determine if environmental conditions were especially suitable for an EEM action. Notably, this monitoring is emphasized only for the EEM action due to its sensitivity to environmental conditions beyond the focal state S_t. Besides ε_t, additional unobservable noise terms in the dynamics are left implicit as these are not central nor unique to EEM.

A manager's objective is to maximize the net benefits of managing S_t through joint deployment of traditional and EEM actions, subject to Equations 1 and 2:

$$\max_{A_t, M_t} \sum_{t=0}^{T} \beta^t \cdot \mathbb{E}\left[B\left(S_t\right) - C\left(A_t, M_t\right) \mid P_t^M, P_t^\varepsilon\right] . \tag{3}$$

The parameter β denotes the discounting of future values over the time horizon, T, and $B(\cdot)$ and $C(\cdot)$ denote benefits and costs. Contemporary benefits—for example, derived from reducing the non-native species—are rewarded as the state S_t responds to prior management actions. While benefits from both actions are delayed, those stemming from M_t take longer to accrue relative to those of A_t due to the dynamics in Equations 1 and 2. Since the scale of EEM actions is often large and suitability depends on an observable shock, optimal implementation will be episodic, depending on the state of the system. For example, altering hydropower generation over long, high-demand periods is expensive and the impact of these designer flows depends on the stochastic number of non-native recruits.

Both observational and parametric uncertainty are particularly salient for the choice of whether to deploy the EEM action, M_t. First, in the observational uncertainty case, a shock represented by ε_t drawn from a distribution P_t^ε determines whether current conditions are more or less suitable for an EEM action (e.g., observing recruitment of the non-native species as discussed above). Second, in the parametric uncertainty case, conditional on state and shock variable levels, the efficacy of M_t is not initially well-known, but can be learned with use. Beliefs over the efficacy of M_t are represented by the distribution P_t^M and are updated with each implementation of M_t. This learning-by-doing is costly but modifies the calculation of the expectation in Equation 3. The benefit of M_t is therefore twofold: not only does the action influence S_t (with delay), but it also increases knowledge on the size of this effect and enables future decision making to be more consistent with the true efficacy of the action. For example, each use of a designer flow could provide information about the stranding-driven mortality rate for the non-native species.

To illustrate how this conceptual model translates to other settings involving EEM, consider two additional examples of EEM actions: beach nourishment and prescribed fire. In our first example, a coastal community may be concerned with maintaining substantial beach depth, an intermediate state (X_t), to enhance coastline stability, a focal state of interest (S_t), in the face of storm risk. A manager could take a traditional action (A_t) to immediately protect coastline stability by protecting existing vegetation (e.g., grasses, mangroves, etc.) or by limiting damaging recreational activity. Depending on beach and habitat conditions and on the observation of relevant stochastic shocks (ε_t), such as seasonal storm forecasts or natural sand availability, the manager also decides to episodically employ beach nourishment (M_t) to augment beach depth, and in turn, coastline stability.

In our second example, forest managers may be concerned with reducing fuel loads, an intermediate state (X_t), to protect neighboring communities from wildfire hazard, a focal state of interest (S_t). Once a fire ignites, a manager can take a traditional action (A_t) to directly address the wildfire hazard by suppressing fires with retardant drops or fire lines. Depending on fuel loads (X_t) and burning conditions (ε_t)—such as wind, relative humidity, and air temperature—managers can employ prescribed burns (M_t) to reduce future fuel loads and, in turn, wildfire hazard. As with beach nourishment, prescribed fires are carried out episodically, given the high opportunity costs of permitting and liability rules and stochastic suitability of burning conditions.

Prior economic studies incorporating ecosystem modification—like those concerning beach nourishment (e.g., Smith et al., 2009; Gopalakrishnan et al., 2017; Cutler et al., 2020) and prescribed fires (e.g., Yoder, 2004; Amacher et al., 2005, 2006; Crowley et al., 2009)—establish several important fundamentals regarding their optimal use; however, they are less suited to address several key features of *episodic* ecosystem modification actions that we emphasize here. This literature generally models the implementation of ecosystem modification actions as a modified Faustmann rotation problem to determine the optimal time interval between their periodic use. Beach erosion and fuel accumulation are often modeled as deterministic, while wildfire arrival is assumed to follow a known stochastic process. Under this structure, the optimal schedule for ecosystem modification is fixed or stationary, and managers do not respond episodically to realized environmental conditions (see Hashida et al., 2025, for an exception). Moreover, these models typically do not incorporate delays in the realization of benefits from beach nourishment or prescribed burns. As a result, prior models often abstract from elements central to many ecosystem modification applications: the importance of monitoring to condition deployment on real-time environmental signals, the value of learning about effectiveness through implementation, and the limitations of ecosystem modification as a rapid-response tool. Our conceptual framework is designed to capture this alternative class of problems, where ecosystem modification is episodic, information-dependent, and characterized by uncertain and delayed impacts.

While the optimal use of EEM actions will differ across applications, the structural features highlighted in the conceptual model allow us to draw several general conclusions. First, traditional actions are more likely to be used as controls than EEM actions. Traditional actions directly combine with the focal state, S_t, that determines the flow of amenities (or damages), whereas EEM actions operate indirectly through the intermediate state, X_t, yielding delayed and more uncertain benefits. As a result, EEM actions may be relatively unattractive when effective traditional actions are available. EEM actions will become more attractive when their opportunity costs are low, conditions for the focal state are dire, observed shocks indicate high suitability, or when traditional actions are weak substitutes for EEM actions. Second, learning-by-doing through EEM experimentation is only likely to be worthwhile when traditional actions are less compelling (ineffective or highly costly) or when dire conditions motivate simultaneous use of both actions.

3 Integrated assessment of designer flows

In this section, we present an application of our general framework for the optimal joint deployment of traditional and EEM actions to achieve a conservation goal. We begin with a description of our setting in Section 3.1 and a biological model that depicts the dynamics and interspecific interactions of trout and humpback chub populations in the Colorado River downriver of the Glen Canyon Dam in Section 3.2. We then introduce our environmental manager's problem, which is to minimize the expected cost of trout-reducing actions (trout removals and designer flows), subject to maintaining (probabilistically) the humpback chub population above a viability threshold. Our solution method is laid out in Section 3.3. We explain our approach to modeling the opportunity cost of the EEM action in Section 3.4, which we characterize as the forgone economic value of hydropower generation when dam releases deviate from those that meet energy generation objectives. Finally, we extend the environmental manager's problem to include adaptive management to reduce uncertainty in the efficacy of designer flows in Section 3.5.

3.1 Setting

The Colorado River provides habitat for the threatened native humpback chub (hereafter chub) and a valuable non-native rainbow trout (*Oncorhynchus mykiss*) and brown trout (*Salmo trutta*) (both hereafter trout) sport fishery (Yackulic et al., 2014; Bair et al., 2016). Contemporary operation of the dam facilitates a self-reproducing population of trout in Glen Canyon which periodically migrates downriver into the Marble and Grand Canyons, especially following large recruitment events (i.e., after the spawning and maturation of a large cohort of juvenile trout) (Korman et al., 2015). Trout that migrate downriver to areas in which chub are found can decrease the survival of juvenile chub and ultimately lower adult populations. In particular, management focuses on the population of chub located around the confluence of the Colorado River and Little Colorado River in

Grand Canyon National Park, 125 km below Glen Canyon Dam. Chub spawn in the Little Colorado River and a portion of juveniles regularly disperses into the mainstem of Colorado River where they are vulnerable to competitive and predatory pressure from adult trout (Yard et al., 2011; Yackulic et al., 2014, 2018b).

Figure 1 shows a map of the study area. Our model tracks subpopulations of trout in three regions: Glen Canyon, which contains trout spawning grounds just downriver of Glen Canyon Dam, Marble Canyon, a 100 km reach of the Colorado River through which some trout disperse downriver, and the Trout Removal Reach, which contains the confluence of the Little Colorado River and the Colorado River mainstem and marks the beginning of the Grand Canyon. The boundaries of the Trout Removal Reach are based on the distribution of chub that spawn in the Little Colorado River. A smaller Juvenile Chub Monitoring Reach is regularly surveyed to make inferences regarding abundances of trout and chub in the broader Trout Removal Reach. Our model additionally tracks chub recruitment in the Little Colorado River.

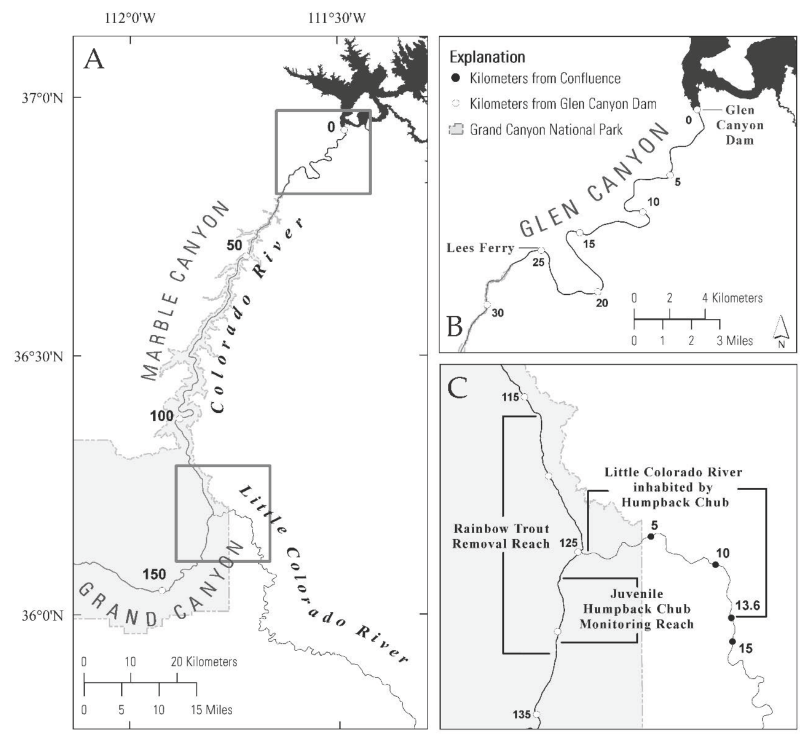

Figure 1: Map of (A) the study area in the Colorado River ecosystem, with detail of (B) the tailwater below Glen Canyon Dam, and (C) chub monitoring and trout removal reaches near the confluence of the Colorado River and Little Colorado River.

Our manager has two opportunities to diminish the trout population to reduce pressure on the chub population. First, for our representative EEM action, operations at Glen

Canyon Dam can be temporarily altered to create periods of sustained high flows followed by a sharp reduction in flow rate (hereafter summarized as the designer flow). This results in a dramatic drop in river elevation that selectively strands juvenile trout—which prefer shallow waters—on higher elevation gravel bars and reduces population growth (Giardina et al., 2024). Efficacy can be improved if data on juvenile populations are available, but this requires complementary monitoring effort. The designer flow can be implemented once per year and involves altering flows at the beginning of May, June, and July. The regional scope of a designer flow's stranding effects is constrained to the Glen Canyon region of our model, i.e., immediately downriver of Glen Canyon Dam.

The second control option is traditional management via removal of adult trout (hereafter the removal action) that occupy the Trout Removal Reach. This approach involves sending teams into the Grand Canyon to conduct electrofishing trips that last approximately two weeks. During past implementation of removals, approximately six trips were scheduled in a year and we assume that this level of effort would be continued (Coggins et al., 2011). These trips provides a second opportunity to reduce a cohort of trout, regardless of whether the designer flow was previously implemented during their juvenile stage. Like the designer flow option, this action can be implemented in each year. However, one key difference is that the cost of the removal action is less sensitive to external factors (e.g., electricity prices) than the designer flow.

Stakeholders of the Glen Canyon Dam Adaptive Management Program are interested in rigorous decision rules for these two trout-reducing actions. Optimal management entails yearly deployment of neither, either, or both of these actions to maintain chub viability at minimal cost. The rest of this section lays out an integrated assessment model designed to meet this objective.

3.2 Viability model

Our biological model adds significant spatial detail to previous analyses of this system that concerned trout removal, namely Bair et al. (2018) and Donovan et al. (2019). Our parameterization (numerical values in Section 4, Table 1) for the model below is calibrated using data from related fieldwork.

We model adult population dynamics at an annual timestep t, and observations are taken prior to management actions. Trout and chub populations evolve according to

$$X_{t+1} = (X_t - y_t + x_t) \cdot \gamma_x \tag{4}$$

$$Y_{t+1} = (Y_t + y_t) \cdot f_y (A_t) \cdot \gamma_y \tag{5}$$

$$S_{t+1} = \max \left\{ (S_t \cdot \gamma_s + s_t \cdot f_s (Y_t)) \cdot \mathbb{1} \left[S_t > \overline{S} \right], \overline{S} \right\} , \tag{6}$$

where X_t, $Y_t \geq 0$, and $S_t \geq \overline{S}$ denote stocks of adult trout upstream in Marble Canyon, adult trout downriver in the Trout Removal Reach, and adult chub within the same reach. Their respective lowercase versions denote new additions to the adult stocks through trout recruitment in Glen Canyon (x_t), trout migration from Marble Canyon (y_t), and chub recruitment in the Little Colorado River (s_t). The natural annual survivorship rate for each stock is provided by the γ terms. The two survival functions $f_y(A_t)$ and $f_s(Y_t)$ (defined below) reflect additional pressures on the two species, with the first capturing the impact of removal action (A_t) and the second reflecting the competition and predation effects of trout on juvenile chub. Trout are generalist feeders and variation in chub abundance has not been shown to alter their growth and survival (Korman et al., 2021; Yard et al., 2023).

A chub population that falls below the viability threshold $\overline{S}$ is considered to be at high risk of extinction. Thus, our manager wishes to avoid falling to $\overline{S}$. We program $\overline{S}$ as an absorbing state, with our decision problem below terminating upon reaching this threshold. In practice, entering this "quasi-extinction" state triggers a new decision-making regime that is beyond the scope of our model (U.S. Fish and Wildlife Service, 2002).

New additions to each stock via recruitment (x_t, s_t) or migrants (y_t) are given by

$$x_t = \psi_x \cdot f_x (M_t) \cdot \exp \left(\varepsilon_{x,t} \right) \tag{7}$$

$$y_t = \psi_y \cdot X_t \tag{8}$$

$$s_t = \psi_s \cdot \varepsilon_{s,t} , \tag{9}$$

where the ψ terms govern the share of a given sub-population that out-migrate. Trout recruitment is subject to additional pressure from designer flows (M_t), with post-designer flow survival captured by $f_x (M_t)$. Studies of trout dynamics in this system have shown that recruitment varies substantially over time with variation in environmental conditions, not density dependence, playing a key role (Korman et al., 2021; Yard et al., 2023). This variation is approximated with a log-uniform distribution. Drivers of chub recruitment are less understood, but recruitment variation is less extreme and can be represented by a uniform distribution (Dzul, 2021). Thus, trout and chub recruitment are represented as stochastic processes with variance driven by the distributions of the ε terms, given by

$$\varepsilon_{x,t} \overset{\text{iid}}{\sim} \text{uniform}\left(\alpha_x, \kappa_x\right) \tag{10}$$

$$\varepsilon_{s,t} \overset{\text{iid}}{\sim} \text{uniform}\left(\alpha_s, \kappa_s\right) \ . \tag{11}$$

The post-designer flow survival of trout recruitment in Glen Canyon, post-removal action survival of trout in the Trout Removal Reach, and survival of juvenile chub as a function of adult trout abundance in the Trout Removal Reach are

$$f_x\left(M_t\right) = 1 - \theta_M \cdot M_t \tag{12}$$

$$f_y\left(A_t\right) = 1 - \theta_A \cdot A_t \tag{13}$$

$$f_s\left(Y_t\right) = \left(\text{logit}^{-1}\left(\mu + \lambda \cdot Y_t\right)\right)^{12} \ , \tag{14}$$

where the θ, μ, and λ parameters specify the strength of these mortality impacts, which were estimated by Yackulic et al. (2018b). Since the timing window for both the removal and designer flow actions is ecologically restricted and both actions are implemented at scale and are costly, they are both limited to one implementation per year: $A_t, M_t \in \{0, 1\}$. The intensity of these actions can, in theory, be scaled down; however, both have a minimum scale below which implementation is assumed to have little impact.

From these dynamic equations, it takes an average of two years for Glen Canyon trout recruits to mature and arrive at the Trout Removal Reach (via Marble Canyon), and are thus potentially exposed to both control actions (A_t, M_t) at different life stages. Figure 2 provides a conceptual diagram of state interactions and action timing.

The use of EEM is novel here; thus, we only have a partial understanding of the efficacy of designer flows. We incorporate this uncertainty in our model by allowing θ_M to be an unknown parameter with a distribution that reflects our current state of knowledge,

$$\theta_M \sim \text{uniform}\left(\underline{\theta}, \overline{\theta}\right), \tag{15}$$

where these [potentially reducible] bounds were obtained through expert elicitation and recent hypsometric analysis (Runge et al., 2015; Giardina et al., 2024).

Our system exhibits each of the features that define the general EEM problem. First, episodic use of designer flows is ideal, as costs are large and fixed, but benefits are state-dependent. Second, when chub viability is threatened, our environmental manager may deploy designer flows in place of, or alongside, trout removals. Third, the effect of designer flows is delayed relative to trout removals, making them less appealing if costs can

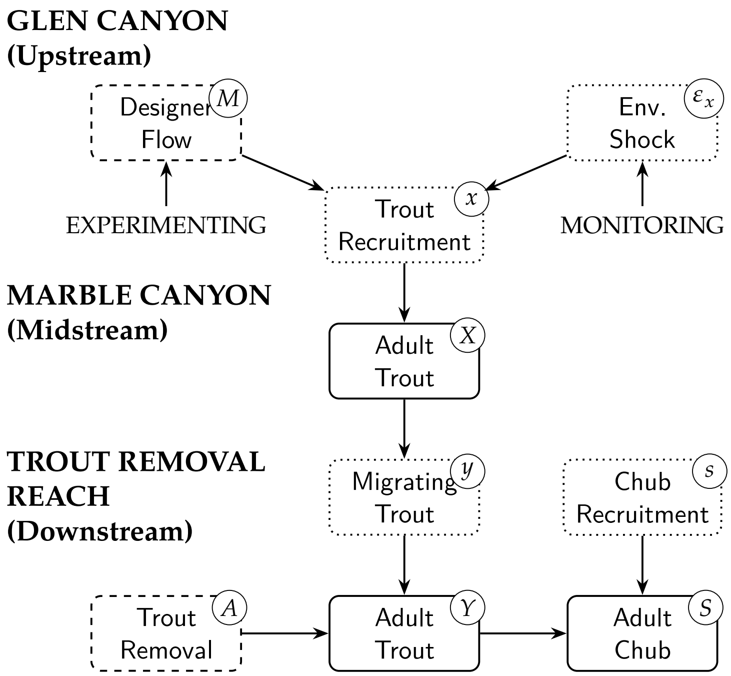

Figure 2: Conceptual diagram of the integrated assessment model of designer flows. Dashed nodes represent control variables and solid nodes represent state variables. We assess model sensitivity to (1) the cost of designer flows (M); (2) the predation intensity of adult trout (Y) on chub (S); and whether or not the manager implements (3) learning-by-doing via experimenting with designer flows, and (4) monitoring to observe the environmental shock.

be postponed until the future. Fourth, an accurate prediction of designer flow impact depends on the monitoring of trout recruitment (while no additional information is needed to effectively deploy trout removal). Fifth, designer flows face additional parametric uncertainty in efficacy due to their novelty, but this may be reducible. Finally, if threats to the humpback chub worsen, the earlier impact of the designer flow may be preferred, as removals permit a trout cohort to pressure the chub population for an additional year.

To match language used in downlisting goals for threatened and endangered species (e.g., that of the humpback chub; U.S. Fish and Wildlife Service, 2002), we are concerned with the cost-effective reduction of trout competition and predation to maintain the adult chub population (S_t) above a viability threshold ($\overline{S}$) over a rolling time horizon (T) with some degree of confidence (Δ). An optimal policy minimizes the expected present-valued cost (with future costs discounted via β) of trout-reducing actions such that this viability goal is continually satisfied for the next T years, an objective which we summarize as

$$\min_{A_t, M_t} \ \mathbb{E}_{\varepsilon}\left[\sum_{t=0}^{\infty} \beta^t \cdot C\left(A_t, M_t\right)\right] \tag{16}$$

$$\text{s.t. } P\left(\bigcap_{\tau=t}^{t+T} S_\tau > \overline{S}\right) \geq \Delta,\ t = 0, 1, \ldots \infty\ , \tag{17}$$

and is subject to the dynamics in Equations 4-15. The expectation is taken over the two stochastic recruitment processes. The total cost within a decision period is given by

$$C(A_t, M_t) = c_A \cdot A_t + c_M \cdot M_t\ , \tag{18}$$

where c_A is the cost of the removal action and c_M is the cost of a designer flow, the latter of which is determined by a model of future energy costs described in the next section.

Our environmental manager's viability horizon of T years is nested in an infinite horizon optimization problem since, in each year t, we are concerned with survivorship over the next T years. This rolling window constraint cannot be recast as something like $S_t > \overline{S},\ \forall t$ because no policy can support this inequality indefinitely due to the stochasticity of the underlying dynamics. Additionally, the time horizon of the cost minimization problem cannot be made finite because this structure incentivizes a reduction in care when there is no immediate risk of reaching $\overline{S}$ within the remaining years before T. A managed species would therefore be left in a diminished state with significant risk of reaching $\overline{S}$ post-T that would cause ex post loss once the manager extends their program beyond the first T years. Thus, the rolling horizon representation of the viability goal ensures both a feasible objective and a time-consistent environmental manager.

3.3 Solution method

Our manager's problem can be equivalently cast as solving for the fixed point of a Bellman equation, which exists given the time consistency of our modeled manager. Let $Z_t = \{X_t, Y_t, S_t\}$ represent the observed set of state variables, $\varepsilon_t \subseteq \left\{\varepsilon_{x,t}, \varepsilon_{s,t}\right\}$ denote unobserved shocks, ζ_t specify shocks that become observed via additional monitoring, and $G\left(Z_t, \zeta_t, \varepsilon_t\right)$ represent the population dynamics specified in Equations 4-15. This notation allows for the possibility of observing a subset of the stochastic elements in the dynamics of $G(Z_t, \zeta_t, \varepsilon_t)$. Then the optimal policy $\{A(Z_t, \zeta_t), M(Z_t, \zeta_t)\}$ satisfies

$$V(Z_t) = \min_{A_t, M_t} C(A_t, M_t) + \beta \cdot \mathbb{E}_\varepsilon [V(Z_{t+1}) \mid A_t, M_t, Z_t, \zeta_t] \tag{19}$$

$$\text{s.t. } P\left(\bigcap_{\tau=t}^{t+T} S_\tau > \overline{S}\right) \geq \Delta \tag{20}$$

$$\text{and } Z_{t+1} = G(Z_t, \zeta_t, \varepsilon_t). \tag{21}$$

While state Z_t embodies all of the information needed to determine a probability distribution for the next decision period's state, information gathered on otherwise unobserved shocks ε_t will reduce the spread of this distribution and potentially improve cost-effectiveness. For example, obtaining the information signal $\zeta_t = \varepsilon_{x,t}$ through trout recruitment monitoring prior to making the designer flow decision reveals the exact impact of that potential action and subsequently increases the precision of knowledge about next year's adult trout population in Marble Canyon. We explore the value of this reduction of observational uncertainty in Section 4.2.

The problem specified in Equations 19-21 is computationally intractable. The viability requirement in Equation 20 does not maintain the Markov property required by standard dynamic programming techniques because evaluating the joint chance constraint requires information from the next T periods in the future. To address this, we use a recently developed shadow value viability approach (Donovan et al., 2019; Donovan and Springborn, 2022). A dynamic programming algorithm specified by Donovan et al. (2019) imposes an estimable penalty Ω on our manager if the chub population falls to $\overline{S}$. This provides a virtual incentive for the manager to avoid states that approach the threshold by taking on costly management action. If crossing the threshold marks the end of a decision problem, we can encode this as irreversible. As shown by Donovan and Springborn (2022), an irreversible threshold produces a state-dependent conservation incentive,

$$\omega(Z_t) = \Omega \cdot \mathbb{E}_{T_f}\left[\beta^{T_f} \mid Z_t\right], \tag{22}$$

where the expectation is taken over the distribution of the time to failure T_f, which itself is conditional on the optimal policy set by Ω. The shape of this penalty function is concave and decreasing in S_t and increasing in X_t and Y_t, which reflects an increasing intensity as the risk of crossing $\overline{S}$ increases.

Given any arbitrary value of Ω, a solution to the problem above—without the viability constraint—is found via value function iteration (Judd, 1998), and the joint chance constraint is checked after the optimization step. We therefore iterate over candidates for Ω

to find the lowest penalty that induces the manager to meet (or exceed) the viability goal from every feasible starting state (some states nearest $\overline{S}$ will fail more often than Δ). This feasible region is referred to as the viability kernel, $\{Z\}_k$ (Donovan et al., 2019; De Lara and Doyen, 2008; Oubraham and Zaccour, 2018). In summary, we solve Program 23:

$$\min \Omega \;\text{ s.t. } P\left\{ \bigcap_{\tau=t}^{t+T} S_\tau > \overline{S} \;\middle|\; A^\Omega, M^\Omega, Z_t \in \{Z\}_k \right\} \geq \Delta \, , \tag{23}$$

where $\{A^\Omega, M^\Omega\}$ solve the following cost-minimization problem conditional on Ω:

$$V^\Omega(Z_t) = \min_{A_t, M_t} \left\{ C(A_t, M_t) + \beta \cdot \mathbb{E}_\varepsilon \left[V^\Omega(Z_{t+1}) \mid A_t, M_t, Z_t, \zeta_t \right] \right\} \tag{24}$$

$$\text{s.t. } V^\Omega\left(S_t = \overline{S}\right) = \Omega \tag{25}$$

$$\text{and } Z_{t+1} = G(Z_t, \zeta_t, \varepsilon_t) \; . \tag{26}$$

After solving for $V^\Omega(Z_t)$ and the optimal policy $\{A^\Omega(Z_t, \zeta_t), M^\Omega(Z_t, \zeta_t)\}$, we can recover $V(Z_t)$ directly by calculating the expected present costs of the optimal policy.

3.4 Opportunity cost of episodic ecosystem modification

The opportunity costs of EEM actions are often more complex than those of traditional management actions and require estimation over a range of possible scenarios. This section quantifies the uncertainty in future values of designer flow costs (c_M). Between years, the energy prices a dam operator faces vary significantly, and this subsequently changes the opportunity cost of hydropower. Here, we compute a range of likely values for c_M that feed into the optimization problem in Section 3.2. This range represents the near-term uncertainty in the opportunity cost of hydropower due to exogenous shocks to Colorado River Basin hydrology or demand in the electricity sector.

The reach of the Colorado River running through Glen, Marble and Grand Canyons is regulated by Glen Canyon Dam (GCD). Annual releases from GCD are allocated to meet Upper Colorado River Basin water deliveries to the Lower Colorado River Basin, while monthly and daily regulation is driven by hydropower and other downriver resource objectives (U.S. Department of the Interior, 2016). Hydropower generation at GCD is utilized in the Western Interconnection, a power grid encompassing the western United States and parts of Canada and Mexico. Given the low marginal cost of hydropower, GCD is used as a load-following facility that adjusts output throughout the day to maximize the value of energy generated as demand changes, allowing avoidance of alternative high-cost electricity generation from other sources during peak demand periods.

Designer flows result in dam releases that deviate from those that meet energy generation objectives. At GCD, dam operators have proposed implementing an annual designer flow cycle, which would consist of three separate designer flows in the months of May, June, and July (U.S. Department of the Interior, 2016). The opportunity cost of a designer flow cycle is the unrealized economic value of energy generation when deviating from the river flow that meets energy generation objectives. To estimate this foregone economic value, we model energy generation under a dam operator selecting hourly flow through GCD to maximize the economic value of energy subject to operational constraints, with and without designer flows. The difference in the optimal economic value of energy generation with and without designer flows informs c_M in Equation 18.

Because peak energy demand occurs in August—outside of the proposed designer flow schedule—additional generation capacity is available on the grid during designer flow implementation. Thus, we characterize the opportunity cost of a designer flow as short-run, i.e., from paying another available generator to supply the forgone energy generation rather than from building new energy generation capacity.

To begin, we model hydropower production (W, in MWh) generated at GCD as a function of hourly flow Q_h (in cfs) through power-generating turbines and reservoir elevation L (feet above mean sea level) (a standard engineering setup following Yackulic et al. (2024)),

$$W\left(Q_h, L\right) = \sigma \cdot Q_h \cdot \left(L - \overline{L}\right) . \tag{27}$$

Q_h is assumed to be constant over an hourly time step h. Because hourly flow through the GCD is assumed to have negligible impact on reservoir elevation, L is exogenous to the contemporaneous hourly flow rate Q_h. L is further assumed to be constant within a month, as elevation varies within tenths of a percent over this timeframe (Bureau of Reclamation, 2024). Energy generation is proportional to reservoir elevation above the tailwater elevation $\overline{L}$. Consequently, σ converts reservoir height to hydraulic pressure, which ensures the full calculation produces a measure of energy (hydraulic pressure × total flow volume).

To identify optimal hourly energy generation over the planning horizon, we utilize the following linear hydropower optimization model which follows Harpman (1999). Within each month, an operator's objective is to maximize the economic value of hydropower,

$$\max_{Q_1,\ldots,Q_H} \sum_{h \in H_m} p_h \cdot W\left(Q_h, L\right) , \tag{28}$$

subject to several operational constraints,

$$\begin{array}{rll} \sum_{h \in H_m} Q_h \leq & \text{max monthly volume (cfs),} & \text{for } m \in \{\text{May, June, July}\} \\ Q_{h=j} \geq & \text{min off-peak flow,} & \text{for } j \in \text{off-peak hours} \\ Q_{h=i} \geq & \text{min on-peak flow,} & \text{for } i \in \text{on-peak hours} \\ Q_h \leq & \text{max flow} & \\ Q_{h-1} - Q_h \leq & \text{max down ramp} & \\ Q_h - Q_{h-1} \leq & \text{max up ramp} & \\ |Q_h - Q_{h-k}| \leq & \text{max flow change in 24 hours,} & \text{for } k \in 24 \text{ hour period.} \end{array} \tag{29}$$

where p_h is the predicted hourly price (\$/MWh) during each hour h within a given monthly window (H_m). Since hydropower generated from the GCD constitutes less than 2% of the generating capacity in the regions of the Western Electricity Coordinating Council where it's marketed, we assume the operator is a price-taker and considers the predicted prices p_h to be exogenous (EIA, 2023). We use historical Palo Verde hub prices to construct a representative weekly price vector for May, June, and July for several years, resulting in a range of hourly prices for each month in which designer flows could take place.

Each designer flow in each month results in three consecutive days of high water releases, followed by an abrupt decline to very low levels for six hours; motivation, design, and reporting of experiments involving this flow pattern are detailed in Korman et al. (2005). When designer flows are implemented, Equation 28 is thus subject to

$$Q_h = \left\{ \begin{array}{rl} 20000 \text{ cfs,} & \text{if } h \in H_{hf} \\ 5000 \text{ cfs,} & \text{if } h \in H_{lf}, \end{array} \right\}. \tag{30}$$

where H_{hf} is the collection of "high flow" hours spanning 8:00 AM - 8:00 AM on the first three weekdays of each month, and H_{lf} is the six hour "low flow" period following each three day high flow event in each month. This "go high then low" strategy completes a trout-stranding designer flow cycle (Giardina et al., 2024). This schedule contrasts a flow schedule prioritizing maximum energy value, which would instead follow demand through the day, subject to operation constraints, maximizing flow during early morning and late evening and minimizing flow during high renewable generation mid-day and lower night time demand.

We solve for the optimal release schedule with and without designer flow constraints. The opportunity cost of implementing a designer flow (c_M) is then the difference between the economic value of hydropower generation under the optimal schedule with no designer flow and the optimal schedule subject to designer flow constraints. We repeat this calculation for each constructed p_h vector, yielding a range of values that we employ in sensitivity analysis. The central estimate is used in our bioeconomic model.

3.5 Adaptive management and EEM experimentation

EEM in the form of designer flows is not well-studied in the Colorado River ecosystem. Their use will not only manage the system under study, but provide useful information on their true efficacy. This will inform changes to designer flow utilization in a post-learning optimal policy and lead to reduced future program costs if they are found to be more effective than previously thought. This section defines an experimental program that employs costly and frequent use of designer flows in the short run to reduce parametric uncertainty in the model outlined in Sections 3.2 and 3.3.

We extend our base model, where beliefs about the designer flow efficacy θ_M follow a uniform distribution in Equation 15, to allow the true value of θ_M to be learnable at a date in the near future. Specifically, we consider a simple adaptive management strategy that chooses whether to implement a learning phase to determine this value. If the learning phase is implemented, upon determination of θ_M, the program switches to a new optimal policy derived via Program 23 based on this new information. The cost of this adaptive management program can then be compared to the costs of the optimal policy in the non-adaptive program that has no learning.

During the short-run learning phase, designer flows are used according to a predetermined learning regime (described below), and the removal action is not used due to its cost and redundancy; any policy with frequent designer flow implementation is more than sufficient to meet the viability objective, though at a higher cost than could be achieved with a mix of the two control options. Upon determination of the true designer flow efficacy—stylized as a random draw from the θ_M distribution, described in Section 3.2, arriving after some amount of time T_L—the learning-savvy environmental manager calculates the optimal policy for the next management phase by solving Program 23 conditional on this θ_M value. The total learning program cost, given new information arriving in year T_L, is

$$C\left(T_L \mid \{M_t \mid 0 \leq \ t \leq T_L\}\right) = \sum_{t=0}^{T_L} \beta^t \cdot c_M \cdot M_t + \beta^{T_L+1} \cdot \mathbb{E}_{\theta_M}\left[\mathbb{E}_{Z_{T_L}}\left[V\left(Z_{T_L} \mid \theta_M\right)\right]\right], \quad (31)$$

where the summation captures total costs during the learning stage and the second term provides the present discounted value of future costs, i.e., the expected costs of an updated optimal program from year T_L that accounts for new information. The continuation value is calculated using iterated expectations; we first calculate the expected long-run costs of the post-learning optimal policy, $V\left(Z_{T_L}\right)$, under each potential true value of θ_M that the manager could observe using our approach in Section 3.3, then calculate the expectation of these conditional expected costs over the θ_M belief distribution. We reduce each $V\left(Z_{T_L} \mid \theta_M\right)$ to a single value using a third expectation over the distribution of potential starting values given an expected number of designer flows in the learning phase.

Next, we determine rules for designer flow implementation during the learning phase. These rules will determine the frequency of designer flows and thus a distribution for the learning time T_L. We consider two types of learning regimes, indexed by R: (1) a designer flow is implemented only in select high trout recruitment years (defined as an event above the $1-\phi$ percentile; $R = \textit{select}$), or (2) a designer flow is implemented in all years ($R = \textit{annual}$). The probability of a high recruitment event, ϕ, defines a threshold recruitment event size, above which the signal from a single designer flow experiment is notably stronger (Yackulic et al., 2018a). The expected cost of the learning program, conditional on ϕ, T_L and R, is

$$\begin{aligned} C\left(T_L \mid R\right) &= \sum_{t=0}^{T_L} \beta^t \cdot \phi^{\mathbb{1}[R=select]} \cdot c_M + \beta^{T_L+1} \cdot \mathbb{E}_{\theta_M}\left[\mathbb{E}_{Z_{T_L}}\left[V\left(Z_{T_L} \mid \theta_M\right)\right]\right] \\ &= \phi^{\mathbb{1}[R=select]} \cdot \frac{1-\beta^{T_L+1}}{1-\beta} \cdot c_M + \beta^{T_L+1} \cdot \mathbb{E}_{\theta_M}\left[\mathbb{E}_{Z_{T_L}}\left[V\left(Z_{T_L} \mid \theta_M\right)\right]\right], \end{aligned} \tag{32}$$

which assumes that under the *select* regime the chance of a designer flow in each period is independent of previous decisions (i.e., that recruitment is not an autocorrelated process).

The learning stage continues until the true designer flow efficacy is determined. The probability of learning the true value of θ_M each year a designer flow is implemented is denoted by π. More literally, we consider this as the probability that stakeholders find that a given experiment establishes sufficient confidence in an updated estimate of θ_M, and set this policy parameter in accordance with expert elicitation within Runge et al. (2015). This interpretation reduces our learning model from a fully-Bayesian process while retaining our desired cost comparison. Thus, the likelihood of learning (i.e., gaining sufficient stakeholder confidence) in a single period is given by

$$p(R) = \phi^{\mathbb{1}[R=select]} \cdot \pi \,. \tag{33}$$

The probability that the manager learns for the first time in year T_L is determined by calculating the likelihood of failing to learn in every year before T_L, multiplied by the likelihood in learning in period T_L,

$$P\left(T_L \mid R\right) = \left(1 - p(R)\right)^{T_L} \cdot p(R) \,. \tag{34}$$

With either learning strategy, the time to learning is distributed geometrically, and the expected learning time will be longer for the *select* regime. We have calculated that the population viability objective is still met under either regime frequency, even in the case where the true θ_M is at the bottom of its possible range of values. Thus, the only meaningful difference between the non-learning and either of the two adaptive management programs is their cost.

A simplified cost comparison provides more intuition with respect to predicting the cost-effectiveness of the learning regime. If the distribution of Z_{T_L} doesn't change with T_L (trout populations would trend downward in expectation over time as a longer learning interval results in more designer flows, thus we relax this idea for our results), the expected cost of the learning program is

$$\begin{aligned}\mathbb{E}_{T_L}\left[C\left(T_L \mid R\right) \mid R\right] &= \sum_{T_L=0}^{\infty} P\left(T_L \mid R\right) \cdot C\left(T_L \mid R\right) = \frac{\phi^{\mathbb{1}[R=select]} \cdot c_M}{1-\beta} \\ &- \frac{p(R)\cdot\beta}{1-\beta+p(R)\cdot\beta} \cdot \left(\frac{\phi^{\mathbb{1}[R=select]} \cdot c_M}{1-\beta} - \mathbb{E}_{\theta_M}\left[\mathbb{E}_{Z_{T_L}}\left[V\left(Z_{T_L} \mid \theta_M\right)\right]\right]\right) \,. \end{aligned} \tag{35}$$

Equation 35 illustrates that, regardless of the learning regime, the expected cost above equals the perpetuity cost of staying in the learning phase forever, minus the expected discounted long-run savings from leaving it.

The viability objective is met under both of these learning regimes, but the more aggressive *annual* flow regime results in higher chub abundances and lower trout abundances over a similar period of implementation. The less aggressive *select* regime utilizes fewer experiments and is less costly in the learning phase on average, but the lower number of observed designer flows results in slower learning and later adoption of an informed policy post-learning.

4 The value of ecosystem modification

In this section, we evaluate several implementation challenges of EEM. In our baseline results (Section 4.1), we solve Program 23 for the optimal joint deployment of traditional (trout removal) and EEM (designer flow) actions. In Section 4.2, we introduce the ability to monitor trout recruitment, which can reduce uncertainty in the impact of a designer flow implementation. The results from these two simulations allow us to characterize the general consequences of the timing of implementation and the delayed conservation effect of the designer flow. In Sections 4.3 and 4.4, we consider a larger non-native species threat that represents increasing stress to the chub and experimentation to reduce uncertainty in the efficacy of designer flows. In sum, our results illustrate common features of EEM actions, explain the mechanisms that make designer flow implementation favorable or unfavorable, and address several questions deemed important by stakeholders of the Colorado River ecosystem. We provide a code package that replicates this analysis and allows customization for specification testing and model adaptation in an online appendix. All parameter values and their sources are presented in Table 1.

Table 1: Parameter definitions with values and sources.

Parameter	Description	Value	Source
α_x, κ_x	trout recruitment bounds	11,14	(Korman et al., 2012)
α_s, κ_s	chub recruitment bounds	4000, 35000	(U.S. Department of the Interior, 2016)
ψ_y	share of trout adults to JCMR	0.009	(Korman et al., 2015)
ψ_s	share of chub recruits to JCMR	0.1	(Yackulic et al., 2014)
ψ_x	share of trout recruits to MC	0.1735	(Korman et al., 2015)
γ_y, γ_x	adult trout natural survival	0.61 (0.69)	(Korman et al., 2015)
γ_s	adult chub natural survival	0.83	(Yackulic et al., 2014)
c_A	cost of [five] trout removal trips	$450K	(Bair et al., 2018)
c_M	cost of designer flow	$250K-$650K	Present Paper
θ_A	removal action efficacy	0.28	(Korman et al., 2012)
$\underline{\theta}, \overline{\theta}$	designer flow efficacy bounds	0.1, 0.5	(Runge et al., 2015; Giardina et al., 2024)
μ, λ	trout viability effects	5, -0.0009 (-0.0063)	(Yackulic et al., 2018b)
T	viability time horizon	20 years	(U.S. Department of the Interior, 2016)
σ	energy conversion $(MWh/cfs \cdot ft)$	0.0000744	(Waldo et al., 2021)
$\underline{L}$	reservoir elevation	3575 ft	(Waldo et al., 2021)
$\overline{L}$	tailwater elevation	3177 ft	(Waldo et al., 2021)
ϕ	high trout recruitment probability	0.25	(Yackulic et al., 2018a)
π	probability of learning event	0.33	(Yackulic et al., 2024)
β	discount factor	0.97	(U.S. EPA, 2024)
$\overline{S}$	lower abundance limit of chub	4000	(Donovan et al., 2019)
Δ	viability confidence	0.90	(Donovan et al., 2019)

Notes: Brown trout parameters are in parentheses. In non-learning simulations, designer flow costs are $450,000. π, β, $\overline{S}$, and Δ are policy parameters chosen to reflect stakeholder concerns and knowledge.

4.1 Several mechanisms lead to limited use of ecosystem modification

In our baseline scenario, our manager observes and responds to adult populations of both chub and rainbow trout, and recruitment of neither species is observed. This section displays an actionable policy for our setting and lays out some baseline stylistic characteristics and limitations of designer flow policy. Further, this section allows us to build upon more predictable results with additional layers of complexity in subsequent sections.

The optimal policy function in Figure 3 presents the adult chub abundance threshold below which it is optimal to use designer flows (3A) and/or the removal action (3B), conditional on adult trout abundances upstream (vertical axes) and downstream (horizontal axes). The color dimension illustrates the *highest* chub population for which action is prescribed for every combination of the trout population variables. Any lower chub population—ceteris paribus—also necessitates action, and any healthier state in the chub dimension does not require action.

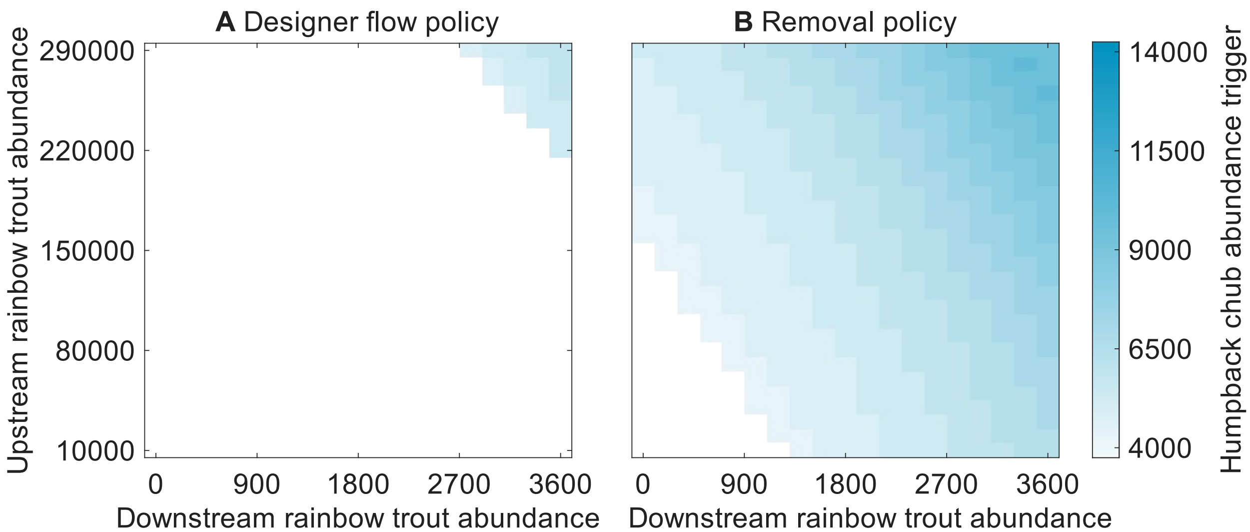

Figure 3: Optimal control policy representing the threshold chub abundance (shading) below which (A) designer flows and (B) removal actions are used, as a function of the adult trout abundances upstream (vertical axes) and downstream (horizontal axes). The shading indicates the highest chub population (and all populations below this) where each binary action is undertaken. In the white (non-shaded) region, it is never optimal to employ the management action.

In the white (non-shaded) region, it is never optimal to employ the management action; that is, there is no chub population above the chub viability threshold ($\overline{S} = 4,000$) for which it is optimal to take the action. Lighter shades imply a stringent threshold for action (i.e., the chub population must fall significantly to trigger a designer flow). Darker shades indicate that even for high chub abundances, action is still optimal. For example, in the top right corner—at maximum abundances of both upstream and downstream adult trout—it is optimal to employ the removal action whenever chub abundance falls below 9,500 and, in addition, designer flows when chub abundance dips under 6,000. As might be expected, chub thresholds are lower when trout abundances are lower and higher when trout abundances are higher.

While designer flows alone are sufficient for achieving viability (Appendix A), they are rarely employed in the optimal policy compared to the removal action. Given that both management actions coincidentally have similar use costs and impacts on trout, there is a preference for the removal action (Figure 3): designer flows must be implemented at an earlier life stage of trout and have a delayed conservation effect on the chub population, while the removal action can be delayed (lowering present-valued costs) until the trout population in the Trout Removal Reach becomes untenable. Our optimal policy illustrates that, given monitoring constraints, designer flows only contribute to cost-effective chub viability in dire situations, i.e., when chub abundance is very low, trout abundance is very high, and the removal action is already implemented. However, as we show next, this situation is unlikely to occur if a manager utilizes this policy.

Figure 4 shows the probability mass function for adult chub, upstream adult trout, and downstream adult trout populations after $T = 20$ years under optimal management. The chub population persists in numbers well above the population viability threshold, and trout abundance is relegated to relatively low levels.

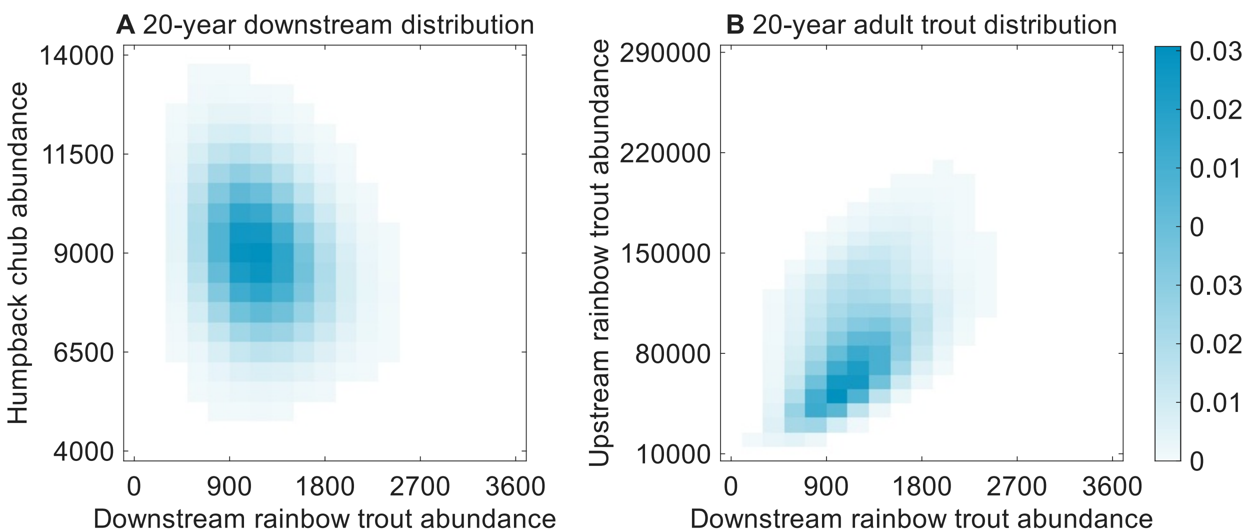

Figure 4: Probability mass function resulting from the optimal policy in the baseline scenario for (A) chub and downstream trout and (B) upstream and downstream trout. The omitted state variable in either panel is held at its 20-year modal value. This figure is conditional on starting state approximating current conditions ($[X, Y, S]$ = [30,000, 1,600, 7,000]), however, the result is not sensitive to alternative starting states (unless extreme).

The primary implication of Figure 4 is that under optimal management, the availability of designer flows does not generate a meaningful decrease in operating costs, as the rarely-visited states that use designer flows do not contribute significantly to the expected present cost calculation. This is because both actions are infrequent—particularly designer flows, as mentioned above. For example, even for high trout populations, designer flows aren't considered until the chub population falls below an abundance of 6,000 (Figure 3A), which rarely occurs (less than seven percent of the time) under optimal management (Figure 4A). This result is ultimately driven by the large and lasting impact that either management action has on trout populations: once action is undertaken, it takes several years on average before the next action is warranted because the stochastic component of trout recruitment spans three orders of magnitude, and events at the top of this range are much less likely.

4.2 Addressing observational uncertainty increases designer flow use

Designer flows have state-dependent benefits. In this section, we demonstrate how a manager with access to monitoring information can make more effective and frequent use of designer flows and the extent to which they shift their preference toward them. While the effects of the removal action may be more predictable, designer flows can prevent one additional year of a given trout cohort's impact on chub by being deployed earlier in the trout life-cycle. In the baseline model, this benefit is not worth the earlier expense. However, with the monitoring of trout recruitment—the subpopulation impacted by designer flows—the timing of designer flows can be selectively implemented when they would have maximum impact (in high recruitment years) and avoided when they would have minimum impact (in low recruitment years). This selectivity raises the marginal benefits of designer flows and limits the tradeoff between the two management actions to a matter of discounting versus delivery of smaller trout cohorts to the Trout Removal Reach.

When informed by trout recruitment monitoring, our policy becomes a function of four state variables, including trout recruitment, upstream adult trout, and downstream adult trout and chub. The less-informed manager in our baseline model faced a distribution of recruitment outcomes, which only resolved once adult trout populations were observed in subsequent years. In contrast, the recruitment-informed manager is able to leverage information on trout recruits one full year earlier—while designer flows can still have an impact on the recruits.

We show the recruitment-informed designer flow policy in Figure 5, which has been modified from the policy function in Figure 3 to account for the extra state variable. Panels A and B of Figure 5 show the designer flow policy for two levels of chub while the color shading responds to the new state variable, trout recruitment. This figure is read similarly to Figure 3, but now designer flows are optimal when *trout recruitment* is at or *above* the level indicated by the color bar. For smaller recruitment events, designer flows are not optimal. The designer flow policy is highly sensitive to trout recruitment: when designer flows are optimal, recruitment is typically quite high, and this recruitment trigger decreases as the adult trout population increases or the chub population decreases, all else equal. The gray region of the state space shows where recruitment does not trigger a designer flow, given the state of the other three subpopulations.

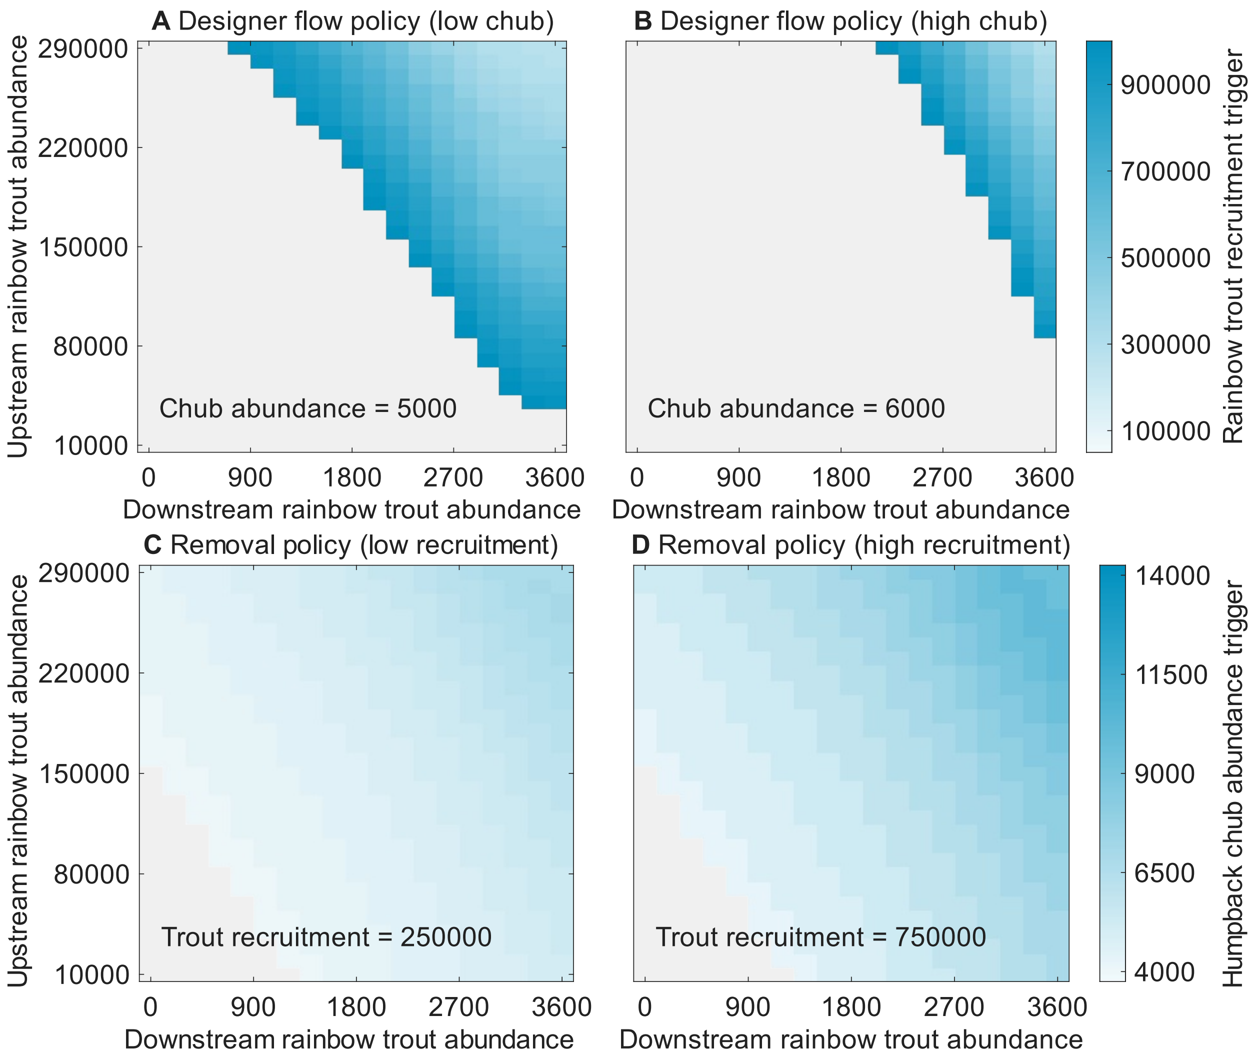

Figure 5: Optimal designer flow policy when trout recruitment is observed, as a function of the adult trout abundances upstream (vertical axes) and downstream (horizontal axes) and chub abundances. In panels A and B, shading indicates the smallest trout recruitment event (and all events above this) where action is undertaken. In panels C and D, shading indicates the largest chub abundance (and all events below this) where action is undertaken. In the gray region, it is never optimal to employ a designer flow.

We find that the recruitment-informed designer flow policy results in virtually the same abundance density as the baseline model (Figure 4) when combined with the new removal policy illustrated in Panels C and D of Figure 5. This implies that designer flows are still used at a low rate even as the optimal policy calls on them in more situations. Further, the recruitment-informed removal policy responds modestly to differences in recruitment information (Figure 5). This is not surprising given that observations of the adult non-native population in later years embody all of the relevant information from an earlier recruitment event. The removal policy relaxes only modestly relative to the baseline (Figure 3B), even in periods with low immediate survival pressure on the chub population and extraordinary trout recruitment (the case least addressed by the removal action). Thus, the expected operating costs are lowered by only 5% with the inclusion of trout recruitment monitoring, showing that the benefit of monitoring is minor.

While recruitment information facilitates the selective use of designer flows and lowers management costs, it is not required for effective control. Further, the removal action impacts both established trout populations downstream and new migrants, whereas trout abundance is immune to designer flow actions the moment they migrate away from Glen Canyon Dam. Therefore, the removal action can be postponed until predation pressure becomes too large to ignore. As mentioned previously, the removal action uniquely presents a recurring management option against each recruitment cohort and therefore reduces the present value of operational costs relative to designer flow implementation through a discounting effect. This benefit of the removal action often outweighs the designer flow's benefit of earlier reduction in trout-derived impacts on chub, even with additional recruitment information to better tailor the designer flow policy.

The relative prioritization of removals versus designer flows depends on the discount rate and energy costs. Increases in the discount rate favor removals over designer flows, given removals' deferrable implementation. However, the binary nature of designer flow implementation requires a large relative change in the costs of the two management actions to motivate a change in the optimal policy. Thus, varying the discount rate within a reasonable range (1-5%) does not meaningfully impact the policy function. In contrast, optimal designer flows change with energy costs in line with expectation: as energy costs increase, designer flow use decreases, and vice versa (Appendix B).

4.3 Designer flows are preferred as threats to conservation worsen

Nonstationary environmental conditions in the Colorado River have led to a significant population increase in another non-native species of concern, brown trout (Yackulic et al., 2020; Healy et al., 2023). Despite their lower numbers, brown trout are much more piscivorous than rainbow trout, with approximately a seven-fold greater per-capita impact on the chub population combined with a longer life expectancy (Runge et al., 2018; Yackulic et al., 2018a). This greater conservation threat, coupled with a lower natural mortality rate, raises the value of the designer flow's effect on an earlier life stage of trout, as it reduces the magnitude of the predation threat by one additional year for a given trout cohort. This section tests the sensitivity of the previously reported preference for the removal action to this critical new development in the ecology of the Colorado River.

We make three parameter changes to our model to reflect brown trout dynamics. First, we update the relationship between trout and chub to match the increase in predation pressure. Second, we set brown trout recruitment to 10% of that of rainbow trout, which leads to lower populations all along the Colorado River. Third, we reduce the natural mortality rate of brown trout, which leads to an adult population that persists longer without management. Otherwise, it is assumed that designer flow efficacy, movement rates, and capture probability for brown trout are similar to those of rainbow trout.

Figure 6 summarizes the optimal policy and long-run population abundances resulting from the updated management problem under the assumption of continued trout recruitment monitoring effort. Panels A through D exhibit similar management patterns as observed previously: the designer flow policy reacts strongly to brown trout recruitment, while the removal policy does not. However, we observe that the removal threshold has relaxed—i.e., is less likely to trigger—in response to more frequent use of designer flows. This is due to the increased natural survival of particularly piscivorous brown trout, which increases the imperative for removing these predators earlier in their life-cycle via EEM. Panel E shows that, despite the elevated threat, the resulting density for the chub is similar to before, which is to be expected given that the viability objective has not changed. Long-run population distributions (panels E and F) together with the designer flow policy function (panels A and B) show that designer flows are used much more frequently relative to the rainbow trout-dominated system, about every other year.

If threats to the humpback chub worsen, designer flows become more favorable, as removals permit a trout cohort to pressure the chub population for an additional year. This result informs the following general conclusion: worsening ecological conditions can reduce the option value advantage of just-in-time traditional actions relative to EEM actions.

4.4 EEM experimentation can produce knowledge at negligible cost

Designer flows are an emerging EEM strategy in our system, so we have limited knowledge concerning their true efficacy. Experimentation in the short run allows our manager to determine their efficacy with greater confidence. This section uses the adaptive management model in Section 3.5 together with the bioeconomic model in Section 3.2 to determine the economic conditions necessary to justify a short-run experimental regime. We apply this adaptive management approach to the brown trout-dominated system in Sec-

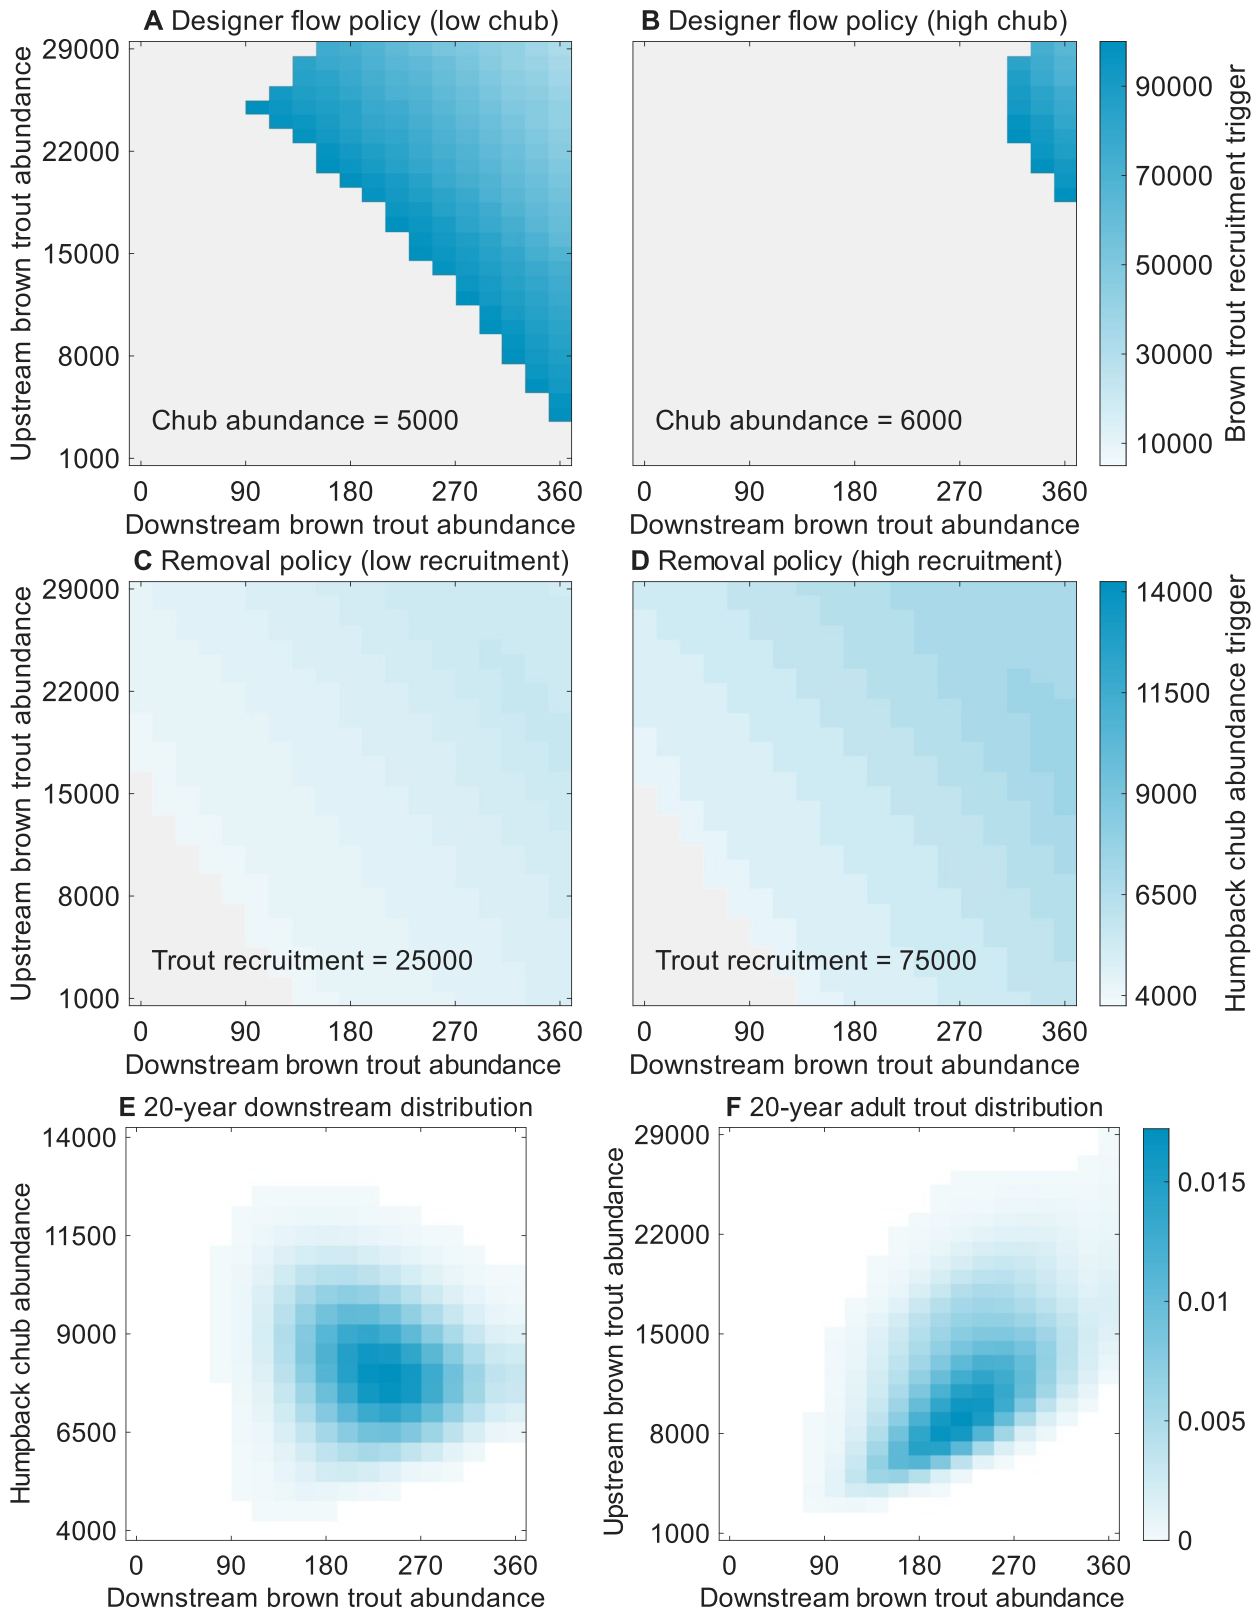

Figure 6: Updated optimal policy and resulting population density for the case of brown trout.

tion 4.3, because the optimal policy for that system is more reliant on designer flows and will therefore be more sensitive to changes in beliefs about designer flow efficacy.

As detailed in Section 3.5, under adaptive management, the manager has the option to employ one of two short-run learning strategies: a more aggressive approach that implements a designer flow every year, and a less aggressive approach that uses them only when brown trout recruitment is high. In either case, we assume that the removal action is suspended due to redundancy and cost. Each designer flow poses an opportunity to gain sufficient confidence about its true efficacy. Upon such a realization, the learning stage ends, and a new optimal policy is designed using the new information.

We consider the point at which the opportunity cost of hydropower is low enough to justify the costs of either of these learning programs over that of a manager who chooses not to implement the adaptive management strategy. Given the significant volatility in energy prices in recent years, we consider how the value of learning changes with changing opportunity costs of hydropower due to exogenous shocks representative of a changing climate or electricity sector. For example, relative to 2020, energy prices at the Palo Verde hub were 162% and 256% higher in 2021 and 2022, respectively, for the months of May-July (California Independent System Operator, 2024). We use this range of prices to determine a range of plausible opportunity costs of implementing a designer flow experiment.

The tradeoff of interest is that these learning regimes are excessively costly in the short run, but experimentation can potentially reduce future management costs. It may be that θ_F is on the higher side of the manager's initial belief distribution, which would result in greater designer flow use and reduce program costs post-learning. However, if θ_F is found to be on the lower end of its theoretical efficacy range, then future costs will likely increase as the manager shifts the post-learning optimal policy back toward the removal action. Thus, the expected change in costs is weighted against the cost of the experiment.

Figure 7 presents the expected program costs for the non-learning scenario, the less aggressive "high recruitment" experimental regime, and the more aggressive "annual flow" experimental regime. Here, there is a linear response to the opportunity costs of designer flows (due to our cost function), and energy costs must fall to 2020 levels in order for a learning regime to produce positive expected net benefits. On this low end of the future energy cost range ($250-350k), the less aggressive experimental regime can promote early adoption of designer flows even when a management plan with lower short-run costs under current and fixed information (the non-learning approach) is available. For all regimes, these lower opportunity costs result in higher designer flow use, which begins to lower the management costs even in the non-learning case. However, lower opportunity costs have a greater effect on the learning regimes, owing to the reduction of the immediate cost of experimentation.

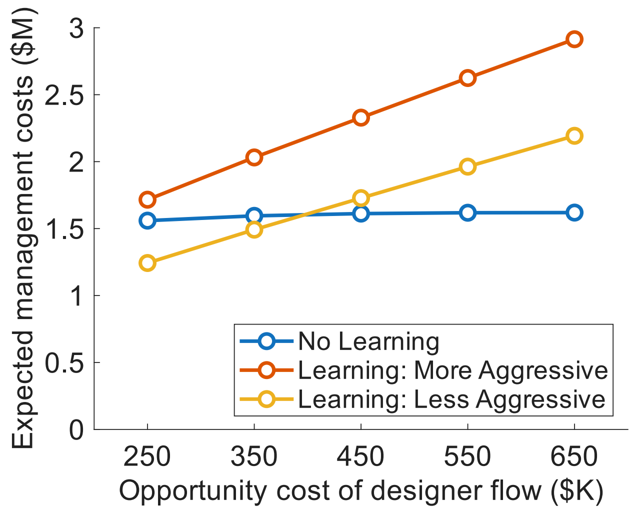

Figure 7: Expected management costs as a function of designer flow opportunity costs under the no learning program and two alternative learning programs.

We find that aggressive learning produces a net loss in all scenarios considered, although this loss is approaching zero for the lowest opportunity cost. This is because the additional designer flow activity in the learning stage is excessive after early experiments quickly bring the system into a favorable state for the chub. Thus, a positive experimentation result (learning efficacy is high) will not substantially lower continuation costs because the post-learning use of the EEM action is still minimal. The learning regime will always produce lower present-valued continuation costs, so our tradeoff is driven by the expected costs of the learning versus the non-learning manager during the learning stage. Thus, a precondition for justifying a learning regime is ”self-subsidization,“ i.e., that the frequency of designer flows is less than annual to keep costs low, yet still sufficient to manage the trout and effectively achieve the viability goal.

The less aggressive learning regime that targets high recruitment events meets the above precondition and appears to be conditionally justified: if future energy costs are not expected to increase, then this regime generates a net benefit relative to a policy that does not employ learning. The optimal non-learning program that manages brown trout costs $1.6M in present valued terms, and potential cost savings from the less aggressive learning regime are 7.5% and 23% of this figure when designer flow opportunity costs are $350K and $250K, respectively. These gains are relatively large, but modest in absolute terms (and similar to the cost of any one management action) due to the lasting impact of a designer flow. These results indicate that, relative to the non-learning policy, a learning regime could be pursued for either a slight reduction or slight increase in costs (depending on electricity prices). Given that designer flow efficacy knowledge is useful for ongoing management—and EEM more generally (conditional on some initial bounds on efficacy beliefs)—these results suggest that an experimental learning regime is sensible to pursue.

5 Discussion

Episodic ecosystem modification actions are an increasingly utilized approach to enhance ecosystem function and services in response to conservation challenges. However, their efficient use in meeting conservation objectives has received limited consideration from economists. In this paper, we consider the cost-effectiveness of a specific form of EEM action, designer flows, in conserving a threatened species in a large dam-regulated ecosystem downstream. We conduct an integrated assessment of designer flows within a shadow value viability framework that incorporates several common implementation challenges associated with EEM actions, including their large opportunity costs, joint deployment with traditional actions, delayed and uncertain effectiveness, and the use of monitoring to improve targeting. This allows us to draw general lessons for utilizing such actions in achieving conservation objectives.

Our application on the Colorado River contains significant economic and biological detail to explore the key mechanisms that inform optimal decision-making and EEM policy design for the stakeholders of the Glen Canyon Dam Adaptive Management Program. We are able to demonstrate the sensitivity of conservation decisions to critical features in our setting, such as the future changes to the cost of foregone hydropower generation, dynamic timing of action, observational uncertainty, interspecific interactions, and opportunities for experimentation to reduce parametric uncertainty regarding the efficacy of designer flows. Our approach provides internal validity for designer flow policy, and our results can additionally inform an ideal course of conservation action in other regulated river systems based on the presence or absence of these critical features.

While our analysis is specific to the Colorado River downstream of Glen Canyon Dam, our study provides at least two insights for the implementation of EEM actions more broadly. First, traditional actions are more likely to be used as controls than EEM actions. Many traditional actions can be deployed in real time as information about the system's focal state is revealed, yielding a cost advantage over the delayed and uncertain effectiveness of EEM actions. For example, traditional actions such as fire suppression offer immediate, visible results by slowing or extinguishing a fire, whereas the benefits of EEM actions, such as prescribed burns, are delayed and uncertain, leading to their underutilization (North et al., 2015). This is particularly true if the opportunity costs of EEM actions exceed those of traditional actions. They become more attractive when conditions are especially dire—in our case, when the abundance of rainbow trout is high and that of humpback chub is low, or when a new threat, such as the more piscivorous brown trout, is present—so that greater overall control is required, particularly if traditional management actions are relatively less effective or are limited by implementation constraints. EEM actions also gain value when monitoring enables targeted, state-contingent deployment when their effectiveness is high. Even so, the economic gains from EEM actions may be modest when the option of implementing traditional actions is present.

Second, experimentation to learn of the efficacy of EEM actions may not be optimal when the short-run cost of traditional actions is low. Experimentation involves a known, high short-run implementation cost with uncertain long-run benefits. While rapid experimentation, with indiscriminate implementation of modification actions without consideration of the system's state, will expedite the learning process, the short-term costs may outweigh those of a tailored, non-learning policy. However, less aggressive state-contingent experimentation has the potential to provide net benefits by implementing modification actions in the most theoretically productive environmental conditions, thereby reducing the short-run cost of experimentation while aiding in achieving long-run conservation objectives. Furthermore, even with targeted experimentation, modification actions in the short run will lower long-run benefits by manipulating a system towards a favorable state that requires less management in the future.

Our study also highlights the benefits of jointly considering EEM actions alongside other regional-scale socio-economic values. For example, if federal hydropower shifts to a larger role in balancing electricity sector grid resources, or if prolonged drought puts additional strain on the Colorado River, our results indicate that the designer flow option may become cost-prohibitive. This increase in problem scope also begets consideration of future variation in many stationary inputs considered here. Electricity price and weather variability will be salient factors in the deployment of designer flows. Shocks to future prices and climate expectations would influence decision making in a similar fashion as observing the juvenile trout recruitment shock, either increasing or decreasing the suitability of deploying designer flows in the given year.

Similarly, future analyses could benefit from a broader scope to consider ancillary benefits and costs arising from designer flows, such as changes in sediment transport dynamics or recreational experiences. While the ancillary benefits and costs arising from designer flows are likely minimal in our setting due to the short duration of the action (Runge et al., 2015), the unintended benefits or costs of other EEM actions, such as altered habitat conditions after prescribed burning or beach nourishment, could be important to consider.

This is especially true if such unintended benefits or costs create institutional and political frictions. In our context, designer flows are used to meet native fish conservation goals but also impact resources of interest to diverse stakeholder groups, such as anglers, hydropower, and Native American Tribes. Similarly, prescribed burns are used to reduce the wildfire risk to nearby communities but may also affect the habitat of protected species under the Endangered Species Act (Ganey et al., 2017). Taken together, these examples illustrate that the desirability of EEM actions ultimately hinges not only on their ecological effectiveness but also on how their ancillary effects intersect with competing stakeholder interests and institutional constraints.

Finally, we note that river-regulating infrastructure and other built environments are constructed for reasons unrelated to conservation objectives. However, creative and flexible dam operation can provide cost-effective approaches to conservation (Bair et al., 2019). Ongoing research drawing on lessons from the present paper involves management of smallmouth bass (*Micropterus dolomieu*), a warm-water non-native predator that is abundant in Lake Powell. Low runoff in the headwaters of the Colorado River basin, combined with only modest decreases in water use, has significantly lowered reservoir elevations in Lake Powell and Lake Mead (Schmidt et al., 2023). Low reservoir elevations in Lake Powell lead to warmer water being passed through hydropower generating intakes, increasing both the passage of smallmouth bass, which are concentrated in these warmer reservoir waters, and their reproduction below the dam, as their spawning and growth depend on warmer water temperatures (Eppehimer et al., 2024). However, by strategically sending water from the bottom of the lake through a hydropower bypass, river temperatures can be reduced to the point where smallmouth bass cannot successfully spawn or grow. These flows were already implemented in 2024 and 2025 at significant cost to hydropower, and modifications were made in both years to meet river temperature targets at minimal cost. While this alternative flow design also comes at a great opportunity cost to hydropower generation, it again highlights how designer flows can strengthen environmental protection, particularly in instances where traditional actions are ineffective and outcomes are dire, as drought conditions become more common and hydrologic variability increases.

References

Amacher, G. S., Malik, A. S., and Haight, R. G. (2005). Not Getting Burned: The Importance of Fire Prevention in Forest Management. *Land Economics*, 81(2):284–302.

Amacher, G. S., Malik, A. S., and Haight, R. G. (2006). Reducing Social Losses from Forest Fires. *Land Economics*, 82(3):367–383.

Bair, L. S., Rogowski, D. L., and Neher, C. (2016). Economic Value of Angling on the Colorado River at Lees Ferry: Using Secondary Data to Estimate the Influence of Seasonality. *North American Journal of Fisheries Management*, 36(6):1229–1239.

Bair, L. S., Yackulic, C. B., Schmidt, J. C., Perry, D. M., Kirchhoff, C. J., Chief, K., and Colombi, B. J. (2019). Incorporating social-ecological considerations into basin-wide responses to climate change in the Colorado River Basin. *Current Opinion in Environmental Sustainability*, 37:14–19.

Bair, L. S., Yackulic, C. B., Springborn, M. R., Reimer, M. N., Bond, C. A., and Coggins, L. G. (2018). Identifying cost-effective invasive species control to enhance endangered species populations in the Grand Canyon, USA. *Biological Conservation*, 220(January):12–20.

Bryan, B. A., Higgins, A., Overton, I. C., Holland, K., Lester, R. E., King, D., Nolan, M., MacDonald, D. H., Connor, J. D., Bjornsson, T., and Kirby, M. (2013). Ecohydrological and socioeconomic integration for the operational management of environmental flows. *Ecological Applications*, 23(5):999–1016.

Bureau of Reclamation (2024). 24-Month Study Projections. Available at: https://www.usbr.gov/lc/region/g4000/riverops/24ms-projections.html (Accessed 13 November 2024).

California Independent System Operator (2024). Open Access Same-time Information System. Available at: http://oasis.caiso.com/mrioasis/logon.do (Accessed 7 August 2024).

Coggins, L. G., Yard, M. D., and Pine, W. E. (2011). Nonnative Fish Control in the Colorado River in Grand Canyon, Arizona: An Effective Program or Serendipitous Timing? *Transactions of the American Fisheries Society*, 140(2):456–470.

Crespo, D., Albiac, J., Kahil, T., Esteban, E., and Baccour, S. (2019). Tradeoffs between Water Uses and Environmental Flows: A Hydroeconomic Analysis in the Ebro Basin. *Water Resources Management*, 33(7):2301–2317.

Crowley, C. S. L., Malik, A. S., Amacher, G. S., and Haight, R. G. (2009). Adjacency Externalities and Forest Fire Prevention. *Land Economics*, 85(1):162–185.

Cutler, E. M., Albert, M. R., and White, K. D. (2020). Tradeoffs between beach nourishment and managed retreat: Insights from dynamic programming for climate adaptation decisions. *Environmental Modelling & Software*, 125:104603.

De Lara, M. and Doyen, L. (2008). *Sustainable management of natural resources: mathematical models and methods*. Springer, Berlin.

Donovan, P., Bair, L. S., Yackulic, C. B., and Springborn, M. R. (2019). Safety in Numbers: Cost-effective Endangered Species Management for Viable Populations. *Land Economics*, 95(3):435–453.

Donovan, P. and Springborn, M. (2022). Balancing conservation and commerce: A shadow value viability approach for governing bycatch. *Journal of Environmental Economics and Management*, 114:102689.

Dzul, M. C. (2021). Life history and population dynamics of humpback chub in the Grand Canyon. *Doctoral Dissertation*, Department of Fish, Wildlife, and Conservation Biology, Colorado State University.

EIA (2023). Electric Power Annual 2022. Technical report, U.S. Energy Information Administration.

Eppehimer, D. E., Yackulic, C. B., Bruckerhoff, L. A., Wang, J., Young, K. L., Bestgen, K. R., Mihalevich, B. A., and Schmidt, J. C. (2024). Declining reservoir elevations following a two-decade drought increase water temperatures and non-native fish passage facilitating a downstream invasion.

Fenichel, E. P., Horan, R. D., and Bence, J. R. (2010). Indirect management of invasive species through bio-controls: A bioeconomic model of salmon and alewife in Lake Michigan. *Resource and Energy Economics*, 32(4):500–518.

Ganey, J. L., Wan, H. Y., Cushman, S. A., and Vojta, C. D. (2017). Conflicting Perspectives on Spotted Owls, Wildfire, and Forest Restoration. *Fire Ecology*, 13(3):146–165.

Garnache, C. (2015). Fish, Farmers, and Floods: Coordinating Institutions to Optimize the Provision of Ecosystem Services. *Journal of the Association of Environmental and Resource Economists*, 2(3):367–399.

Giardina, M., Korman, J., Yard, M. D., Wright, S., Kaplinski, M. A., and Bennett, G. (2024). A literature review and hypsometric analysis to support decisions on trout management flows on the Colorado River downstream from Glen Canyon Dam. Report 2024-1033, United States Geological Survey, Reston, VA.

Gopalakrishnan, S., McNamara, D., Smith, M. D., and Murray, A. B. (2017). Decentralized Management Hinders Coastal Climate Adaptation: The Spatial-dynamics of Beach Nourishment. *Environmental and Resource Economics*, 67(4):761–787.

Harpman, D. A. (1999). Assessing the Short-Run Economic Cost of Environmental Constraints on Hydropower Operations at Glen Canyon Dam. *Land Economics*, 75(3):390.

Hashida, Y., Lewis, D. J., and Cummins, K. (2025). Prescribed fires as a climate change adaptation tool. *Journal of Environmental Economics and Management*, 130:103081.

Healy, B. D., Budy, P., Yackulic, C. B., Murphy, B. P., Schelly, R. C., and McKinstry, M. C. (2023). Exploring metapopulation-scale suppression alternatives for a global invader in a river network experiencing climate change. *Conservation Biology*, 37(1):e13993.

Higgins, S. I., Azorin, E. J., Cowling, R. M., and Morris, M. J. (1997). A dynamic ecological-economic model as a tool for conflict resolution in an invasive-alien-plant, biological control and native-plant scenario. *Ecological Economics*, 22(2):141–154.

Jaeger, W. K. (2025). Accounting for Income and Population Dynamics in Benefit-Cost Analysis: An Application to Dam Removal. *Environmental and Resource Economics*.

Judd, K. L. (1998). *Numerical Methods in Economics*. MIT Press.

Korman, J., Kaplinski, M., Hazel, III, J. E., and Melis, T. S. (2005). Effects of the Experimental Fluctuating Flows from Glen Canyon Dam in 2003 and 2004 on the Early Life Stages of Rainbow Trout in the Colorado River. Technical Report Final Report, U.S. Geological Survey.

Korman, J., Martell, S. J. D., Walters, C. J., Makinster, A. S., Coggins, L. G., Yard, M. D., and Persons, W. R. (2012). Estimating recruitment dynamics and movement of rainbow trout (Oncorhynchus mykiss) in the Colorado River in Grand Canyon using an integrated assessment model. *Canadian Journal of Fisheries and Aquatic Sciences*, 69(11):1827–1849.

Korman, J., Yard, M. D., Dzul, M. C., Yackulic, C. B., Dodrill, M. J., Deemer, B. R., and Kennedy, T. A. (2021). Changes in prey, turbidity, and competition reduce somatic growth and cause the collapse of a fish population. *Ecological Monographs*, 91(1):e01427.

Korman, J., Yard, M. D., and Yackulic, C. B. (2015). Factors controlling the abundance of rainbow trout in the {Colorado} {River} in {Grand} {Canyon} in a reach utilized by endangered humpback chub. *Canadian Journal of Fisheries and Aquatic Sciences*, 73(1):105–124.

Lewis, L. Y., Bohlen, C., and Wilson, S. (2008). Dams, Dam Removal, and River Restoration: A Hedonic Property Value Analysis. *Contemporary Economic Policy*, 26(2):175–186.

North, M. P., Stephens, S. L., Collins, B. M., Agee, J. K., Aplet, G., Franklin, J. F., and Fulé, P. Z. (2015). Reform forest fire management. *Science*, 349(6254):1280–1281.

Oubraham, A. and Zaccour, G. (2018). A Survey of Applications of Viability Theory to the Sustainable Exploitation of Renewable Resources.

Paul, C., Hanley, N., Meyer, S. T., Fürst, C., Weisser, W. W., and Knoke, T. (2020). On the functional relationship between biodiversity and economic value. *Science Advances*, 6(5):eaax7712.

Runge, M. C., LaGory, K. E., Russell, K., Balsom, J. R., Butler, R. A., Coggins, Lewis G., J., Grantz, K. A., Hayse, J., Hlohowskyj, I., Korman, J., May, J. E., O'Rourke, D. J., Poch, L. A., Prairie, J. R., VanKuiken, J. C., Van Lonkhuyzen, R. A., Varyu, D. R., Verhaaren, B. T., Veselka, T. D., Williams, N. T., Wuthrich, K. K., Yackulic, C. B., Billerbeck, R. P., and Knowles, G. W. (2015). Decision analysis to support development of the Glen Canyon Dam long-term experimental and management plan. Report 2015-5176, United States Geological Survey, Reston, VA.

Runge, M. C., Yackulic, C. B., Bair, L. S., Kennedy, T. A., Valdez, R. A., Ellsworth, C., Kershner, J. L., Rogers, R. S., Trammell, M. A., and Young, K. L. (2018). Brown trout in the Lees Ferry reach of the Colorado River—Evaluation of causal hypotheses and potential interventions. Report 2018-1069, United States Geological Survey, Reston, VA.

Schmidt, J. C., Yackulic, C. B., and Kuhn, E. (2023). The Colorado River water crisis: Its origin and the future. *WIREs Water*, 10(6):e1672.

Smith, M. D., Slott, J. M., McNamara, D., and Brad Murray, A. (2009). Beach nourishment as a dynamic capital accumulation problem. *Journal of Environmental Economics and Management*, 58(1):58–71.

U.S. Department of the Interior (2016). *Glen Canyon Dam Long-Term Experimental and Managment Plan*. U.S. Department of the Interior, Bureau of Reclamation, Upper Colorado Region, National Park Service, Intermountain Region.

U.S. EPA (2024). Guidelines for Preparing Economic Analyses, 3rd Edition. Technical Report EPA-240-R-24-001, U.S. EPA, Washington, DC.

U.S. Fish and Wildlife Service (2002). Humpback chub (Gila cypha) Recovery Goals: amendment and supplement to the Humpback Chub Recovery Plan. Technical report, U.S. Fish and Wildlife Service, Mountain-Prairie Region (6), Denver, Colorado.

Waldo, S., Deemer, B. R., Bair, L. S., and Beaulieu, J. J. (2021). Greenhouse gas emissions from an arid-zone reservoir and their environmental policy significance: Results from existing global models and an exploratory dataset. *Environmental Science & Policy*, 120:53–62.

Yackulic, C., Bair, L., Eppehimer, D. E., Salter, G. L., Deemer, B., Butterfield, B. J., Kasprak, A., Caster, J., Fairley, H. C., Grams, P., and others (2024). Modeling the impacts of Glen Canyon Dam operations on Colorado River resources. Technical report, Bureau of Reclamation.

Yackulic, C. B., Dodrill, M., Dzul, M., Sanderlin, J. S., and Reid, J. A. (2020). A need for speed in Bayesian population models: a practical guide to marginalizing and recovering discrete latent states. *Ecological Applications*, 30(5):e02112.

Yackulic, C. B., Korman, J., and Coggins, L. (2018a). Population dynamics of humpback chub, rainbow trout and brown trout in the Colorado River in its Grand Canyon Reach: modelling code and input data.

Yackulic, C. B., Korman, J., Yard, M. D., and Dzul, M. (2018b). Inferring species interactions through joint mark-recapture analysis. *Ecology*, 99(4):812–821.

Yackulic, C. B., Yard, M. D., Korman, J., and Haverbeke, D. R. (2014). A quantitative life history of endangered humpback chub that spawn in the {Little} {Colorado} {River}: variation in movement, growth, and survival. *Ecology and Evolution*, 4(7):1006–1018.

Yard, M. D., Coggins, L. G., Baxter, C. V., Bennett, G. E., and Korman, J. (2011). Trout Piscivory in the Colorado River, Grand Canyon: Effects of Turbidity, Temperature, and Fish Prey Availability. *Transactions of the American Fisheries Society*, 140(2):471–486.

Yard, M. D., Yackulic, C. B., Korman, J., Dodrill, M. J., and Deemer, B. R. (2023). Declines in prey production during the collapse of a tailwater Rainbow Trout population are associated with changing reservoir conditions. *Transactions of the American Fisheries Society*, 152(1):35–50.

Yoder, J. (2004). Playing with Fire: Endogenous Risk in Resource Management. *American Journal of Agricultural Economics*, 86(4):933–948.

A Substitutability of designer flow and removal actions

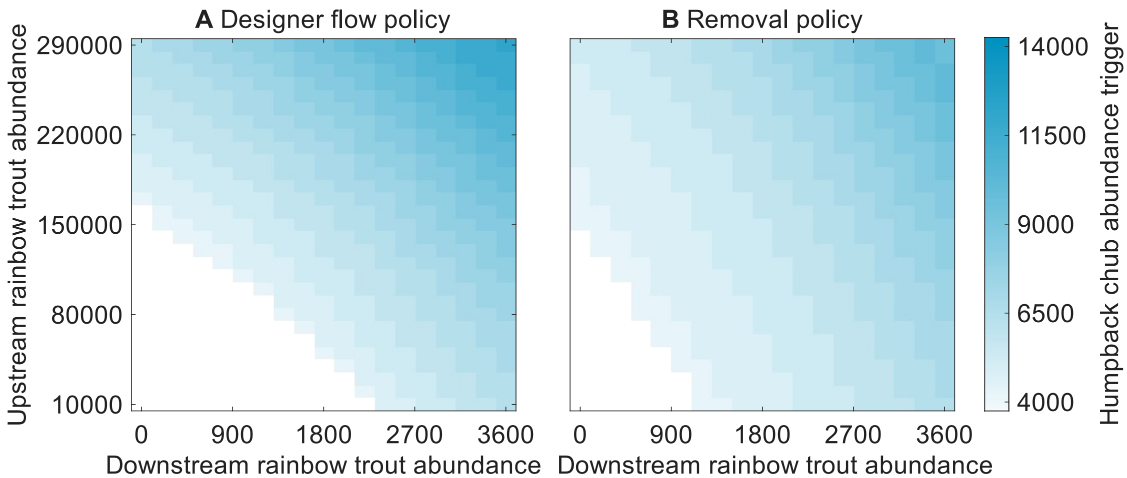

Figure A1: A: designer flow policy when removals are unavailable. B: removal policy when designer flows are unavailable. Either policy can achieve the joint abundance distribution in Figure 4, although the designer flow policy is more costly (and thus not chosen in Figure 3).

B Sensitivity of the optimal policy to energy costs

Figures B1 and B2 show the optimal policy function and resulting density as the opportunity costs of a designer flow fall to \$250k and rise to \$650k per implementation, respectively. As expected, this large change in the relative costs of the two management actions motivates a change in the policy function, with a greater impact on the designer flow policy than the removal policy. Smaller relative cost changes, like those arising from a change in discount rate, are not displayed due to the optimal program's insensitivity to changes at that scale. The resulting density does not change meaningfully because optimal action quickly improves the biological state, so further action is not frequently needed, as seen in Section 4.1.

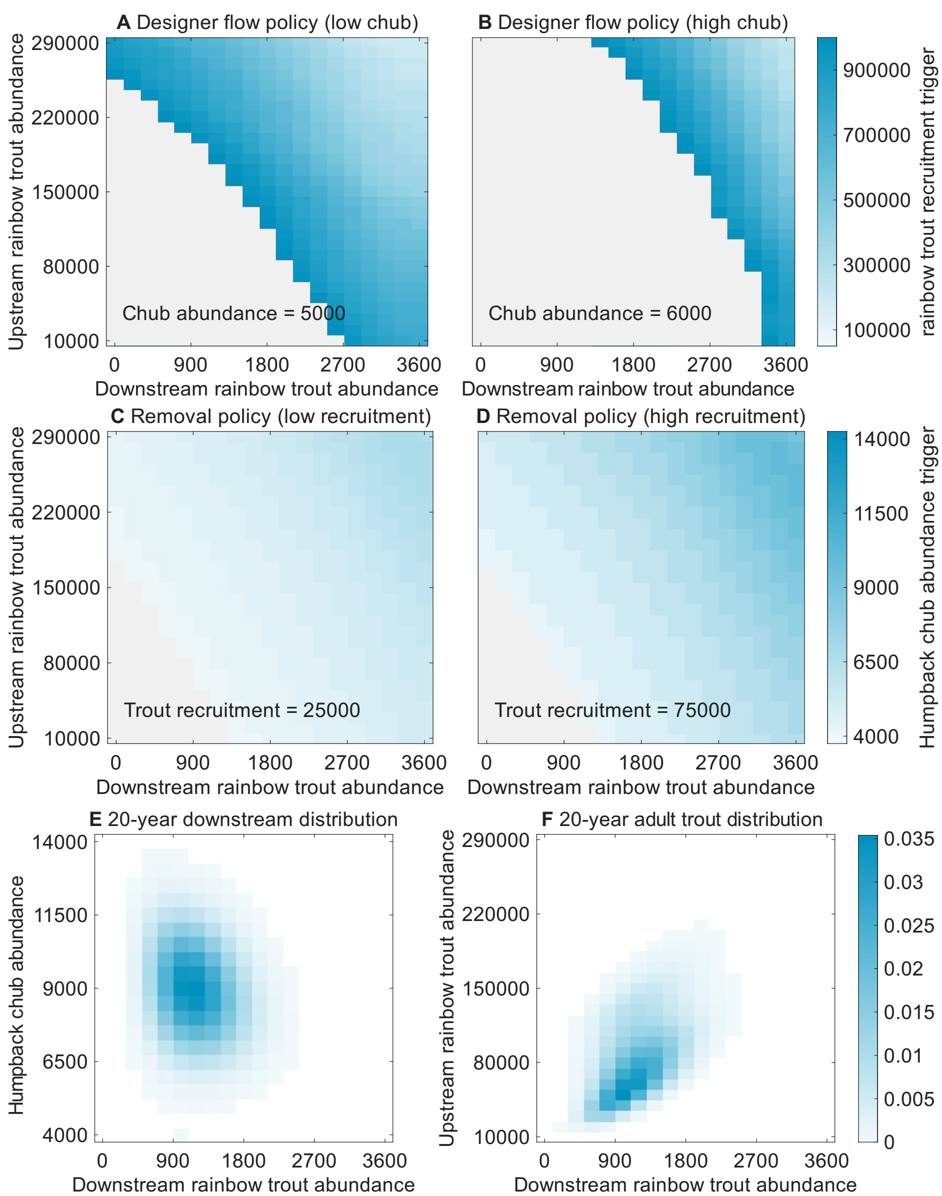

Figure B1: Optimal policy and resulting density when opportunity costs of a designer flow = \$250k.

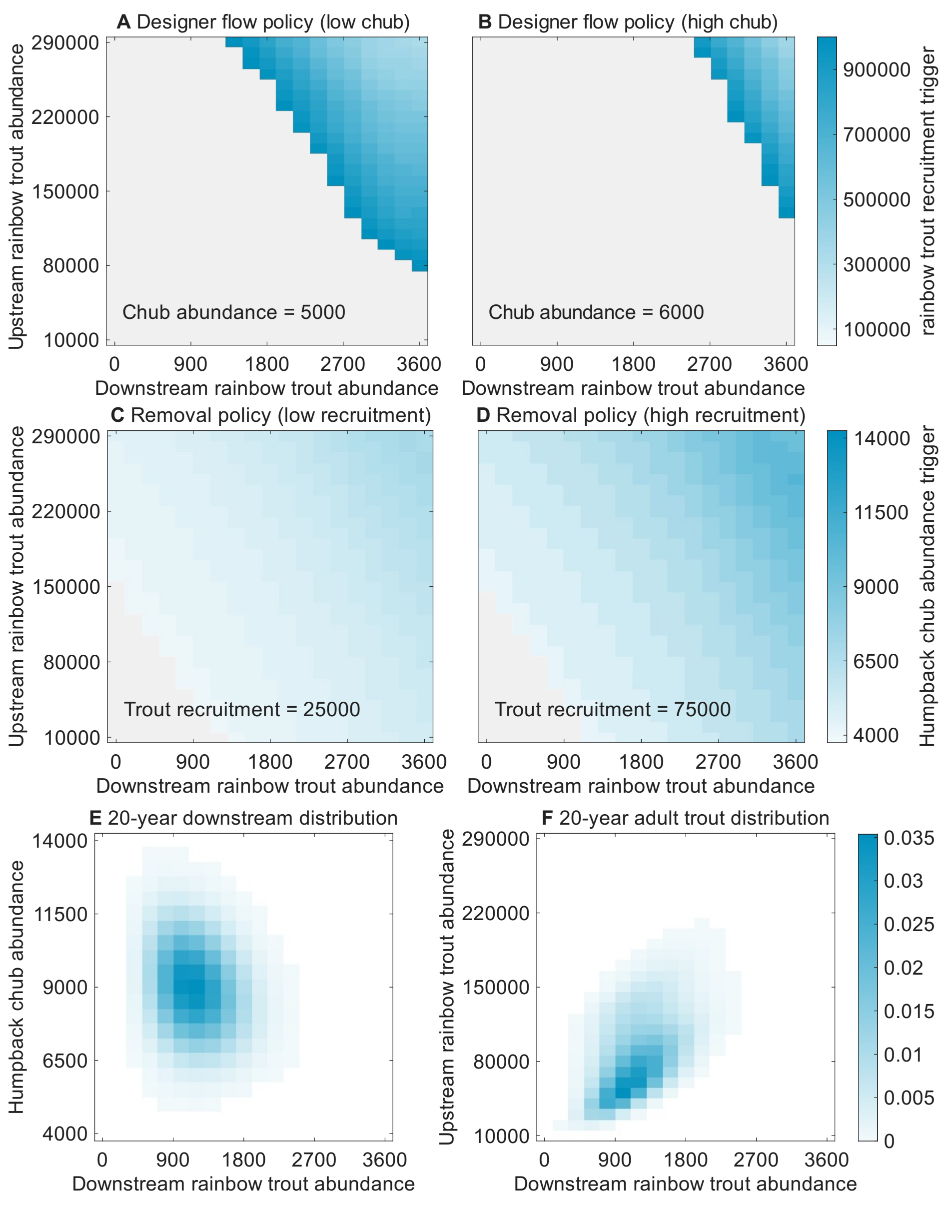

Figure B2: Optimal policy and resulting density when opportunity costs of a designer flow = $650k.