UC Davis

UC Davis Previously Published Works

Title

Safety in numbers: Cost-effective endangered species management for viable populations

Permalink

<https://escholarship.org/uc/item/4vz9j7tj>

Journal

Land Economics, 95(3)

ISSN

0023-7639

Authors

Donovan, P
Bair, LS
Yackulic, CB
et al.

Publication Date

2019-08-01

DOI

10.3368/LE.95.3.435

Peer reviewed

Safety in numbers: cost-effective endangered species management for viable populations

Pierce Donovan¹, Lucas Bair³, Charles B. Yackulic³, and Michael Springborn²

¹Corresponding Author (email: donovan@ucdavis.edu),
University of California, Davis, Agricultural and Resource Economics

²University of California, Davis, Environmental Science and Policy

³U.S. Geological Survey, Southwest Biological Science Center,
Grand Canyon Monitoring and Research Center

October 31st, 2018

Abstract

We develop a bioeconomic model to identify the cost-effective control of an invasive species (rainbow trout) to achieve a population viability goal for an endangered species (humpback chub) in the Grand Canyon of the U.S. southwest. Solving the population viability problem is difficult since avoiding a threshold with a given confidence imposes a probabilistic restriction on joint outcomes (survival) over many periods. We develop a novel dynamic programming solution approach that is fast and forgoes the simulation method requirement of imposing structure on the policy function. We also investigate an adaptive management model that incorporates learning about uncertain biological dynamics.

Keywords— population viability, margin of safety, chance-constrained dynamic programming, adaptive management, endangered species conservation, invasive species, value of information

---

The Glen Canyon Dam Adaptive Management Program financially supported 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

A general class of management problem, both within and outside economics, is concerned with staying above (or below) a particular threshold over time with a specified margin of safety. Notable examples include safely maintaining endangered species away from extinction, global temperatures below a maximum increase, and diseases away from outbreak levels. These examples share two key attributes which make the application of standard expected net benefit objectives problematic: dynamic uncertainty is central, and the consequences of crossing the threshold are dire and difficult to quantify. In these cases, one natural way forward is to substitute a “chance constraint” for consequences—whereby a threshold is avoided with a given confidence—and focus on minimizing management costs.

Our application is based in the Colorado River in the U.S. southwest and concerns the control of an invasive species—rainbow trout (*Oncorhynchus mykiss*) immediately below the Glen Canyon Dam—in order to maintain a viable population of an endangered species, humpback chub (*Gila cypha*) downstream in the Grand Canyon. Control of rainbow trout relieves the competitive and predatory pressure on humpback chub. Efficient management of interacting invasive and endangered species populations is a pressing conservation issue (Brown and Shogren, 1998; Lampert et al., 2014). Invasive species are important drivers of decline for endangered species across the U.S. (Wilcove et al., 1998). Ongoing invasive species population control is frequently employed as a threat reduction strategy in endangered species programs, particularly when full eradication is prohibitively costly. Even if eradication is feasible, in some cases it may not be desirable because the invasive species provides other benefits (e.g. Lampert et al. 2014). In our setting, rainbow trout support a valuable recreational fishery below the Glen Canyon Dam (GCD).

We incorporate population viability analysis (PVA), which has been used to identify threats faced by a population, the risk of population decline, and the probability of meeting a recovery target (Coulson et al., 2001; Rout et al., 2009; Doyen et al., 2012; Pe’er et al., 2013; Finseth and Conrad, 2014). Optimal management for viability presents unique bioeconomic modeling challenges. Rather than focusing on maximizing a measure of social surplus, conservation objectives for endangered species in this context typically involve meeting population-oriented targets. In practice, since population dynamics are stochastic, achievement of a target is not deterministic but rather a question of likelihood. Thus the focus in the decision framework shifts from justifying conservation ends (e.g. economic value of a species) to achieving conservation goals with a given confidence (e.g. Newbold and Siikamäki 2009; Sagoff 2009). Such management focused on a critical levels or thresh-

olds is similar to the logic of safe minimum standards, detailed thoroughly in Margolis and Nævdal (2008).

Significant economic value is associated with threatened and endangered species (Boyle and Bishop, 1987; Loomis and White, 1996). Estimated values range from \$6 - \$95 per household and are dependent on species abundance and characteristics and type of respondent (Loomis and White, 1996; Kotchen and Reiling, 2000; Richardson and Loomis, 2009). Values for humpback chub in the Colorado River in Grand Canyon National Park have been estimated through a series of non-market valuation surveys (Welsh et al., 1995; Duffield et al., 2016). Sampling both regional and national populations, Duffield et al. (2016) estimated marginal values of \$1.75 per U.S. household for a 1% increase in humpback chub abundance. While the valuation is substantial, additional information is required for decision making since the marginal value almost certainly increases drastically as the population declines.

Previous work involving the economic cost of species conservation plans has focused either on direct conservation costs or opportunity costs, e.g. from forgone revenues from the working landscape. Montgomery et al. (1994) develop a cost curve for the probability of spotted owl survival, and analyze the trade-offs between timber sector revenues and increased species survival. Haight (1995) balances site use for conservation or timber by maximizing timber revenues while meeting a viability target over a given horizon for a species. Marshall et al. (2000) simulate the population response of Kirtland's warbler to timber rotation length. They then seek to minimize timber revenue losses with respect to a viability goal, using the distributions found during simulation. Finseth and Conrad (2014) focus on minimizing the direct costs of conservation (transporting species or constructing habitat) for the red-cockaded woodpecker.

This work also relates to the literature on applying *viability theory* to sustainable renewable resource management, which involves meeting a combination of biological, economic and social constraints (Oubraham and Zaccour, 2018). Viability problems share many of the same features found in optimal control (dynamically evolving states subject to control) but lack a traditional objective function (e.g. profit) to be optimized. The key objects of interest in viability theory are often *sets*, the most salient being the *viability kernel*, a subset of the initial state-space that admits a viable solution path that meets constraints, either with certainty or with a given confidence. In the analysis to follow, we derive the least-cost policy to maintain a system within a [stochastic] viability kernel. For further discussion of this literature see De Lara and Doyen (2008), Baumgärtner and Quaas (2009), Doyen and De Lara (2010), and Oubraham and Zaccour (2018).

In this paper we illustrate a new chance-constrained dynamic programming solution

approach to the species viability management problem, which is also applicable to the broader class of joint-chance constraint problems. Because the chance constraint is conditional on joint probabilistic outcomes over many periods, it is non-linear and notoriously difficult to handle. One common approach is to use simulation. For example, Bair et al. (2018) solve for the cost-minimizing policy to achieve viability of an endangered species over a given time horizon with a specified confidence. This involves generating a large block of simulations for every candidate policy in a large set. Thus, it is computationally expensive and necessarily imposes parametric structure on the policy function.

Multiple authors have addressed the problem using a hybrid, simulation-programming approach. Haight (1995) uses a large block of stochastic simulations to evaluate whether the viability constraint is satisfied under a given action sequence. This constraint is enforced using a penalty function (violation of the constraint generates a high cost and is thus avoided). Conditional on a starting state, Haight uses nonlinear programming to identify the optimal harvest action (at one to three time points). While successfully addressing the chance constraint, this approach relies on simulation, solves only for a single starting state, and requires particular care to avoid convergence to local optima. Finseth and Conrad (2014) also apply a simulation-programming hybrid but in the opposite order. First they solve the viability problem using dynamic programming but for a deterministic population target sought by the end of the time horizon, again enforced via an arbitrarily chosen penalty function. They then use simulation *ex post* to explore the effect of stochasticity on population performance relative to the target. While the application of dynamic programming reliably optimizes over a range of potential states, the solution and full treatment of uncertainty are separated.

In this paper we implement a solution approach to the species viability problem that integrates uncertainty in the dynamic programming solution from start to finish. The dynamics and optimal policy endogenously determine the distribution of population outcomes. We use a penalty approach to address the constraint. Instead of applying the penalty on violation of the probabilistic constraint, we apply the penalty conditional on hitting the undesirable threshold itself. We solve for the penalty that is just severe enough to achieve satisfaction of the manager's probabilistic viability goal. Thus the penalty level is determined endogenously. This method generates additional economic intuition for the viability problem by returning an estimate of the shadow value of failing to avoid the threshold, i.e. the economic loss—implied by the chance constraint—associated with crossing the threshold. Relative to the simulation-only solution approach, our method does not require any policy functional form assumptions. Since no simulations are required the method is fast. Our approach is most similar to two recent applications. Ono

et al. (2015) consider a joint chance constraint problem of safely landing a rover on Mars without hitting obstacles. However, while Ono et al. simplify (linearize) the key constraint to ease the solution, we maintain the original. Alais et al. (2017) solve for the optimal management of a hydroelectric dam amid competing generation and recreation values without simplifying the key constraint. Our solution approach also incorporates the unsimplified constraint but is less computationally intensive and holds more generally over the state space as detailed in Section 4.

Using data on rainbow trout control measure effectiveness and population dynamics from long-term study at the U.S. Geological Survey, Grand Canyon Monitoring and Research Center (Yackulic et al., 2014; Bair et al., 2018; Yackulic et al., 2018), we solve for the cost effective rainbow trout removal policy which achieves a given humpback chub viability goal. We find that the resulting policy rule is a non-linear function of the population of both species, a form that is not obtainable with a standard, parsimonious parametric model (as commonly used in simulation-based solution methods). The resulting estimated distribution of population states under optimal management is maintained far above the threshold of interest in order to achieve the stated margin of safety. The cost of the optimal program is significantly less than that from previous analysis of the problem using a simulation-based method (Bair et al., 2018).

In our analysis, we provide new intuition on the scale of economic loss associated with a population falling below a given threshold that is implied when managers assert a given viability goal. We show how this “current” shadow value of violating the threshold propagates to nearby states as a “present” shadow value (the discounted, probability-weighted value of falling to the threshold). The likelihood of deteriorating to a non-viable population increases as the population state degrades, and thus so does the present shadow value. In a relatively safe state, non-viability is not an immediate issue and the present shadow value is low.

After our baseline analysis, we extend the model to incorporate uncertainty in a central parameter and the opportunity to reduce this uncertainty with learning. The speed of the dynamic programming solution method allows us to explore the implications of [reducible] parametric uncertainty. We focus on the impact of rainbow trout on humpback chub, which is imperfectly understood. We consider an adaptive management scenario in which the manager learns (imperfectly) about the true nature of this relationship. This information arrives exogenously in a manner consistent with previous and planned research on these species.¹ To our knowledge, optimal adaptive management in joint-chance con-

---

¹This type of learning and response is known as passive (versus active) adaptive management. In the passive case, the manager does not seek to optimize over the level and timing of learning but rather simply

strained problems has yet to be studied. See LaRiviere et al. (2018) for a general review of uncertainty and learning in natural resource management and a focused discussion of the economic value of reducing environmental uncertainty.

In our results we decompose the value of learning into two components. First, a pre-learning “prospective” value arises from the anticipated arrival of new information, allowing for a temporary relaxation of control policy (shifting costly action into the future). Second, a post-learning “congruity” value stems from less error in optimal management, as beliefs are more closely related to the truth. We find that the value of information reduces the expected present cost of the optimal policy by up to 20%, and is relatively invariant to the starting state.

2 The Humpback Chub Management Program

2.1 Humpback Chub

The focus of this study is on managing interactions between adult rainbow trout (defined as greater than 300 mm in length) and juvenile humpback chub (defined as between 40 and 99 mm in length) in the Colorado River (mainstem) near its confluence with the Little Colorado River (LCR) in Grand Canyon National Park as depicted in Figure 1. Adult rainbow trout eat and compete with juvenile humpback chub lowering survival and, to a lesser extent, growth rates of juvenile humpback chub (Yard et al., 2011; Ward et al., 2016; Yackulic et al., 2018). Both species migrate into this area after spawning elsewhere in the system—humpback chub in the lower 13.6 km of the LCR and rainbow trout further upstream in the mainstem. Humpback chub rear in both the LCR and mainstem and all size classes can be found in both rivers, however, the size of an adult humpback chub population that can be supported by rearing in the LCR alone is insufficient to meet manager’s adult humpback chub abundance goals (Yackulic et al., 2014). Rainbow trout were introduced, and now naturally reproduce, in the tailwater that was created in the 25 kilometers (km) below GCD. Some portion of rainbow trout produced in the tailwater move down into upper Marble Canyon, and an even a smaller portion of these recruits eventually make their way down to the confluence of the mainstem and LCR, 125 km downstream from GCD (see Korman et al. 2012, 2015 for more details).

In this study we focus on population dynamics in the mainstem of the Colorado River around the confluence with the LCR. Regular monitoring of humpback chub and rainbow

---

anticipates and incorporates new knowledge as it arrives exogenously to improve decisions (LaRiviere et al., 2018).

trout occurs in the Juvenile Chub Management Reach (JCMR), a 3 km reach just below the mainstem-LCR confluence. Estimates from this reach are intended to inform triggers for management actions, including the removal of rainbow trout (Yackulic et al., 2014). Figure 1 displays both the confluence of the two rivers, as well as the approximate extent of the humpback chub population and monitoring reaches (Yackulic et al., 2014).

Figure 1: Geographic area of interest. Glen Canyon Dam (GCD) on the Colorado River (mainstem) is located in northern Arizona (upper right panel). The rainbow trout fishery below GCD extends 25 kilometers to Lees Ferry and rainbow trout disperse 125 kilometers downstream to the confluence of the mainstem and Little Colorado River (left panel). The locus of management is the confluence where rainbow trout and the humpback chub aggregation overlap (lower right panel).

2.2 Management Methods

The primary management tool available to managers to aid humpback chub is to mechanically remove rainbow trout in the management reach using electro-fishing techniques. Given the remote nature of the the confluence between the mainstem and the LCR, re-

removals can be costly. Moreover, the area is considered sacred by many local Native American tribes, some of which object to removals. For these reasons, an optimal policy would make sparing use of removals, while also maintaining the chub population above a threshold level.

While biological models of humpback chub used in the system frequently consider multiple size classes in addition to the juvenile size class, we collapse all sub-adult and adult age classes into an “adult chub equivalent”.² Managers are interested in maintaining an aggregate adult chub abundance in the mainstem above 5,000 adult humpback chub (>200 mm). Over space, this aggregate adult chub population is split between the LCR and mainstem. Our focus is on the mainstem, the portion of the population that is impacted by rainbow trout. Converting the abundance goal from adult to “adult equivalent” and accounting for our focus on the confluence of the mainstem and LCR, we are concerned with staying above a threshold of population of 4,000 residing in the mainstem.³ For the remainder of the paper, we simply use humpback chub population or level to signify the adult chub equivalent in the mainstem.

3 Baseline Model

The viability goal is to maintain the chub population level (Y) above a viability threshold ($\bar{Y} = 4,000$) over a given time horizon ($T = 20$) with a given level of confidence ($\Delta = 90\%$). The optimization problem involves selecting the level of trout control (A) based on the population of trout (X) and chub to achieve the viability goal at the lowest expected present cost. The decision-maker’s viability horizon of T years is nested in an infinite horizon optimization problem since as time advances we continue to be concerned with survivorship over the next T years. The problem below characterizes a unique optimal solution, one that traces out a policy function detailing the “best” action $A(X, Y)$ for any combination of trout and chub.

The optimal policy minimizes the expected present cost of trout control such that our

---

²Chub age classes include two sub-adult classes (100 mm -149 mm; 150 mm - 199 mm) and two adult classes (200 mm - 249 mm; 250+ mm). We use an average survival rate based on the two adult size classes and apply expected sub-adult growth and survival to transition juveniles to adult equivalents.

³Maximum likelihood estimates of survival, growth and population size structure from Yackulic et al. (2014), suggest that in the mainstem Y “adult equivalents” equate to $0.75Y$ adults. If we make the further assumption that 2,000 adults are always present in the LCR portion of the population, then adult equivalent can be mapped to an aggregate abundance, N , using the equation: $N = .75Y + 2000$.

joint chance constraint (or viability goal) is satisfied over the time-horizon T :

$$\min_{\{A_t\}_{t=0}^{\infty} \in \mathcal{A}} E \left[\sum_{t=0}^{\infty} \delta^t C(A_t) \right] \quad (1)$$

$$\text{s.t. } Pr \left\{ \bigcap_{t=1}^T Y_t > \bar{Y} \right\} \geq \Delta. \quad (2)$$

The contemporaneous cost is given by $C(A_t) = cA_t$, where c is the cost per removal trip. Within a year, each additional trip removes a lower absolute number of trout since trips already completed in the season leave fewer trout to target. Because the season suitable for control is limited, the number of removal trips is selected from the bounded set $\mathcal{A} = \{0, 1, \dots, 6\}$.

Our biological model primarily follows previous biological and bioeconomic analyses in this system (Bair et al., 2018; Yackulic et al., 2014, 2018). Population dynamics for the mainstem adult trout (X) and chub (Y) stocks are given by

$$X_{t+1} = (X_t + x_t) \cdot s_X(A_t) \cdot \gamma, \quad X_t \geq 0 \quad (3)$$

$$Y_{t+1} = \begin{cases} Y_t \cdot \eta + y_t \cdot s_Y(X_t), & Y_t > \bar{Y} \\ \bar{Y}, & Y_t \leq \bar{Y} \end{cases}, \quad (4)$$

where a chub population that falls below the viability threshold ($\bar{Y}$) is considered “collapsed.” Natural annual survivorship shares for trout and chub are given by γ and η , respectively. The stochastic recruitments to each stock, x_t and y_t , are drawn from the following distributions

$$x_t = \psi_X \cdot \exp(\varepsilon_t^X), \quad \varepsilon_t^X \stackrel{\text{iid}}{\sim} \text{unif}(\alpha_X, \beta_X), \quad (5)$$

$$y_t = \psi_Y \cdot \varepsilon_t^Y, \quad \varepsilon_t^Y \stackrel{\text{iid}}{\sim} \text{unif}(\alpha_Y, \beta_Y). \quad (6)$$

The functions $s_X(A_t)$ and $s_Y(X_t)$ give post-removal survival of trout and survival of juvenile chub as a function of the rainbow trout abundance, respectively:

$$s_X(A_t) = (1 - \theta)^{p \cdot A_t}, \quad (7)$$

$$s_Y(X_t) = \left(\text{logit}^{-1}(\mu + \lambda X_t) \right)^{12}. \quad (8)$$

Parameter descriptions, values, and sources are given in Table 1. A timeline summarizing the dynamics reflected in Equations 3 and 4 is given in Appendix A.

Parameters are estimated using data collected at the Juvenile Chub Management Reach

and the LCR (Yackulic et al., 2014). We make some important adjustments, as past simulation-based modeling efforts used a monthly timestep and a more finely resolved population model (Bair et al., 2018). For example, in the trout survivorship Equation 8, monthly survivorship is converted to annual using an exponent of 12.

Table 1: Variable and parameter definitions with values and sources.

State and Choice Variables			
t	timestep (annual)		
X_t	population of trout at time t		
Y_t	population of chub at time t		
A_t	number of removal trips at time t		
Parameter	Description	Value	Source
α_X, β_X	stochastic trout recruitment bounds	11,14	Korman et al. (2012)
α_Y, β_Y	stochastic chub recruitment bounds	4000, 35000	Interior (2016)
ψ_X	trout outmigration rate to JCMR ⁴	0.0035	Korman et al. (2015)
ψ_Y	share of chub recruits to JCMR ⁵	0.1	Yackulic et al. (2014)
γ	trout survival after natural mortality	0.61	Korman et al. (2015)
η	adult chub survival after natural mortality	0.83	Yackulic et al. (2014)
c	cost of each trout removal trip	\$75,000	Bair et al. (2018)
p	passes per trip	5	Bair et al. (2018)
θ	removal efficacy	0.011	Korman et al. (2012)
μ, λ	trout viability effects (monthly)	5, -0.0009	Yackulic et al. (2018)
T	viability time horizon	20 years	Interior (2016)
δ	discount factor	0.97	Chosen
$\bar{Y}$	lower abundance limit of chub	4000	Chosen
Δ	viability confidence	0.90	Chosen

4 Dynamic programming with a joint chance constraint

The manager’s problem, specified in Equations 1-2 above, is to minimize the present expected cost of trout removal over an infinite horizon, given a finite-horizon viability (abundance) goal. In dynamic programming, the optimal policy is often found via value function iteration (VFI, see e.g. Judd 1998). Our policy of interest, $A(x, y)$, is one that

⁴More specifically, outmigration is calculated as (migration share) · (1-year survival) · (arrival in JCMR).

⁵Similarly, (migration share) · (1-year juvenile survival).

motivates the fixed point of the Bellman equation

$$V(X_t, Y_t) = \min_{A_t \in \mathcal{A}} \overbrace{C(A_t)}^{\text{cost of current removals}} + \delta \cdot \underbrace{E_\varepsilon[V(X_{t+1}, Y_{t+1}) | A_t, X_t, Y_t]}_{\text{expected cost of future removals}} \quad (9)$$

$$\text{s.t. } Pr \left\{ \bigcap_t^{t+T} Y_t > \bar{Y} \right\} \geq \Delta. \quad (10)$$

This problem is also subject to population dynamics specified in Equations 3-8, which we suppress here for simplicity.

To our knowledge, there is no direct solution to the dynamic programming problem expressed above. Straightforward application of VFI is not feasible due to the viability goal, which imposes a joint chance constraint spanning outcomes over the viability horizon (T). Concretely, the manager is concerned with the likelihood of avoiding the population threshold *jointly* across T periods over which risk may be traded off for cost-effectiveness.

The difficulty imposed by this multi-period, non-additive constraint has led many modelers to use Monte Carlo simulation to solve the optimization problem (e.g. Bair et al. 2018). However, this approach carries two significant drawbacks. First, simulation methods are time-intensive and only converge to the optimal policy if it happens to be considered in the set of candidate solutions. More critically, simulation requires a modeler to specify a particular functional form for the policy function. For example, to allow for the simplest non-linear policy function—a quadratic defined over chub population levels—would involve generating a large block of simulations for each candidate parameter triplet. As one might expect (and as we show in our results) the optimal policy is unlikely to nicely follow a quadratic form. Furthermore, we are interested in solving the problem conditional on more than just the chub population level; the state of trout as a chub predator is also of central concern. However, to parametrically capture a second state variable would entail a second set of policy parameters (and chub-trout state interaction terms). Such a brute force grid search over seven or more dimensions of the policy parameter space would be computationally prohibitive.

As we note in the introduction, others that have used programming methods for such problems have made a number of simplifications. For example, Ono et al. (2015) solve the above problem by linearizing the chance constraint using Boole's inequality, therefore allowing for a Lagrangian approach. In our framework this would entail removing the

“joint-ness” in Equation 10. Boole’s inequality gives us

$$1 - Pr \left\{ \bigcap_t^{t+T} Y_t > \bar{Y} \right\} = Pr \left\{ \bigcup_t^{t+T} Y_t \leq \bar{Y} \right\} \leq \sum_t^{t+T} Pr(Y_t \leq \bar{Y}). \quad (11)$$

Ono uses the right-hand side of Equation 11 to form a new constraint,

$$\sum_t^{t+T} Pr(Y_t \leq \bar{Y}) \leq 1 - \Delta. \quad (12)$$

Inequalities 11 and 12 imply that our original constraint will also be met. The optimal policy to this new problem, Equations 9 and 12, is more conservative than the original under the constraint in Equation 10 (Ono et al., 2015). This results in a more costly-than-necessary policy and forces a departure between the “true” dynamics of the system and that of the model.

Our solution approach is summarized in the modified dynamic programming problem in Equations 14-16. To overcome the challenge presented by the joint chance constraint, we impose an *estimable* “penalty” (Ω) in value terms that is incurred by the manager if the chub population falls to the threshold level ($\bar{Y}$). Once reached, the threshold population level is treated as irreversible (see Equation 4), thus the penalty is conceptualized as the present value perpetuity cost of having violated the threshold. This provides incentive for the manager to take costly action to avoid the threshold.

Once the optimal policy is found for any particular penalty, $A(x, y|\Omega)$, we assess the “risk-to-go” function (Ono et al., 2015),

$$r(X_0, Y_0|A(x, y|\Omega)) := 1 - Pr \left\{ \bigcap_{t=1}^T (Y_t > \bar{Y}) | X_0, Y_0, A(x, y|\Omega) \right\}, \quad (13)$$

which designates the probability of hitting the chub population threshold given an initial state, policy function $A(x, y)$, and time horizon, T (see Appendix C for details on calculation). We iterate over penalty levels until solving for the lowest penalty that induces the manager to meet (or exceed) the viability goal ($r(\cdot) \leq 1 - \Delta$) *everywhere it is feasible*.⁶ This feasible region is commonly referred to as the [stochastic] *viability kernel* in the viable control literature, (see De Lara and Doyen (2008) and Oubraham and Zaccour (2018)). Appendix B details the application of our penalty function method to the PVA problem. In

---

⁶As an alternative we could solve the model to meet the viability goal from a given point in the state space (as opposed to everywhere it is feasible). However, this policy would not be stationary but rather would have to be re-solved for each new point in the state space that is reached.

summary, we solve the dynamic programming problem given by

$$\tilde{V}(X_t, Y_t) = \min_{A_t \in \mathcal{A}} \overbrace{C(A_t)}^{\text{cost of current removals}} + \delta \cdot \underbrace{E_\varepsilon[\tilde{V}(X_{t+1}, Y_{t+1}) | A_t, X_t, Y_t]}_{\text{expected cost of future removals}} \quad (14)$$

$$\text{s.t. } \tilde{V}(X_t, \bar{Y}) = \Omega, \quad (15)$$

$$r(X_t, Y_t) \leq 1 - \Delta. \quad (16)$$

For any generic penalty, the resulting policy function will be cost-effective. However, there exists a unique minimum penalty that “just” motivates enough action to satisfy the viability constraint. In addition, this penalty provides a “current” shadow value of the viability constraint. The penalty detailed above (Ω) represents the cost of crossing the threshold implied by the viability objective. While the current value of the penalty is actually incurred only at the non-viable population level, the expected present cost of incurring the penalty, or “present” shadow value, propagates in the value function to neighboring states, which are undesirable given their proximity to the threshold and the penalty incurred there, thus providing an incentive to avoid degraded population levels:

$$\omega(X_0, Y_0) = \Omega \cdot \sum_{t=0}^{\infty} \delta^t \cdot Pr(Y_{t+1} \leq \bar{Y} | X_t, Y_t > \bar{Y}). \quad (17)$$

In other words, $\omega(X_0, Y_0)$ represents the “pain” of potentially realizing the undesirable threshold outcome, propagated out to nearby states depending on how risky they are. The probability, $Pr(\cdot)$, represents the marginal probability of reaching the threshold for the first time in period $t + 1$ (i.e. $Y_t > \bar{Y}$), given trout and chub levels in period t . This probability, multiplied by the discounted penalty, characterizes the marginal contribution to the present shadow value at the given state.

The present shadow value in Equation 17 emerges from the dynamic programming solution without actually computing the right hand side of Equation 17. As described above, we impose the penalty solely at the threshold and allow it to propagate across the state space via the VFI solution process. After calculating the value function, we can difference out the expected present cost of implementing the optimal level of control. This difference gives us the present shadow value in dollar terms at any given time $t = 0$ state:

$$\omega(X_0, Y_0) = \tilde{V}(X_0, Y_0) - E \left[\sum_{t=0}^{\infty} \delta^t C(A^\Omega(X_t, Y_t)) \right] \quad (18)$$

In our numerical application we verify that the implicit calculation in Equation 18 pro-

duces the same values as the explicit expression in Equation 17.

The present shadow value plays the same role of the penalty function seen in other approaches, as it incentivizes the avoidance of an undesirable state, however, like our resulting population outcomes it is also endogenously determined within the model. Thus we avoid the common ad-hoc assumption pertaining to the penalty function’s form, usually a linear or quadratic function with arbitrarily defined parameters. In our case, neither of these forms can sufficiently “fit” the non-parametric result.

The approach of Alais et al. (2017) is quite similar to ours. In contrast to the model above, they use a more-familiar Lagrange multiplier approach that directly penalizes a violation of the chance constraint itself (i.e. any distance between the expected likelihood of avoiding the threshold and the target confidence level). Like us, they first “choose” a penalty and then solve the dynamic programming problem given that value, iterating on the penalty until the chance constraint binds. While the approach of Alais et al. (2017) is elegant and clear advance over earlier approaches there are two key limitations. First, their algorithm solves only for a single starting state. We solve for the optimal policy simultaneously across all states. Second, Alais et al. (2017) solve for a fixed (non-rolling) horizon such that after stepping one period ahead, viability over the *next* T years is not assured.⁷ Furthermore, by period $T - 1$, the policy is only concerned with viability over a *single* period. In contrast, our “rolling window” implementation of the finite-horizon chance constraint ensures that for every time period (e.g. the current year, next year, etc.) the chance constraint is always satisfied over the ensuing T periods. Additionally, their solution approach involves higher computational burden. For a given penalty level, their risk-to-go function must be calculated for every potential policy during their optimization step. We need calculate ours only once, post-optimization, conditional on the optimal policy. Finally, the optimal penalty in Alais et al. (2017) provides economic intuition about the constraint, but only for a particular starting state. Insight from the penalty term in our model holds across state space within the viability kernel.

5 Results

5.1 Optimal Policy and Population Outcomes

We present the optimal policy function, mapping any point in the state space to a level of control, in Figure 2A. Here we consider 100 trout and chub states each, for a total of

---

⁷Of course the model can be resolved in the next period. However, for this approach to be fully optimal such policy updating would need to be acknowledged in expected dynamics, which is not straightforward.

10,000 possible states. Superimposed on the policy function are two additional pieces of information. First, the dashed line stretching from the top left to the right bisects the region where meeting the joint chance constraint is feasible (below) and not feasible (above) under both a policy of maximum action everywhere and the equivalent (in viability terms) optimal policy, depicted here. This feasible region is commonly referred to as the *viability kernel* (De Lara and Doyen, 2008). Second, the concentric curves show level sets of the density function after 20 years under the optimal policy, a density which is shown in full in Figure 2B. The level sets lie at cumulative densities of 10%, 25%, and 50% respectively as they expand out from the mode at $(x, y) = (1400, 7400)$. The fan-shaped region in the upper right of Figure 2A is defined in the caption and discussed further below.

Figure 2: Panel A shows the policy function (underlying shading) indicating how many mechanical removals of trout are optimal in a year given the current populations of trout (horizontal axis) and chub (vertical axis) in the management reach. The dashed line spanning the figure from left to right delineates the upper boundary of the viability kernel. The fan-shaped region in the top right indicates states at which the chub population declines to the threshold with certainty. Panel B depicts the joint density of chub and trout levels after 20 years under the optimal policy. The concentric curves in panel A represent level sets of the joint density from panel B, centered at the mode and depicting cumulative densities of 10%, 25%, and 50%. (Chub level = adult chub equivalent population in the mainstem.)

As expected, maximum possible control (six removal trips) is optimal at very low levels of chub. Maximum control is also ideal at very high levels of trout, unless the chub population is very high. When chub are abundant and/or trout are less abundant (towards the bottom-left), the optimal policy specifies no action.

Currently, in the mainstem the population of chub at around 12,000 is very strong⁸ while trout levels are very low (approximately 100) (Korman and Yard, 2017; Yackulic et al., 2018). Here the optimal policy prescribes no action. This locus is also safely within the feasible region (below the spanning dashed line), indicating that current conditions for the chub are not dire. This finding is in agreement with Bair et al. (2018)—who set removal triggers based on trout abundance alone—and consistent with removal triggers set in a U.S. Fish and Wildlife Service Biological Opinion (Interior, 2016). However, we illustrate how ideal management should be sensitive to evolving joint chub-trout populations rather than solely trout (Bair et al., 2018). Our analysis suggests that cost-effective control can maintain a viable population. Even so, the possibility of transitioning north of the dashed line (and eventually to the threshold) is precisely what drives the optimal level of action depicted in Figure 2A, even at a substantial distance below this line.

Recall that when population levels lie outside of the viability kernel (above the spanning dashed line), it is not feasible to meet the joint chance constraint with the specified confidence. In Figure 2A we have assumed that maximum action would still be applied in this circumstance. However, in this region managers would have a strong incentive to expand the set of management options beyond what is modeled in our analysis. This is especially true in the top-right of Figure 2A where the fan-shaped area delineates a region in which populations are so extreme that the resident chub population will fall to the threshold level with certainty, even under maximum trout control. Fortunately, both current populations levels and the population density at the end of the viability horizon are safely to the southwest of this zone.

Because the policy function depicted in Figure 2 is represented by a lookup table, we avoid imposing a specific functional form. We find that the optimal policy function is non-linear in both trout and chub population levels and also exhibits a clear interaction between these two state variables. Given the nonlinearity, interaction, and relatively rapid transition from minimum to maximum action in the state space, this function would be difficult to represent in a parsimonious parametric model, especially before this general shape is even identified. This illustrates an advantage of the dynamic programming approach used here relative to simulation-based solution techniques that require a functional form assumption.

In Figure 2B we show the state space density—characterizing the joint likelihood of chub and trout levels—after 20 years under the optimal policy. Although such a density depends on the starting chub-trout population for shorter time ranges, we find that after

---

⁸Approximately 11,000 aggregate adults correspond to 12,000 adult equivalents in the mainstem (Yackulic et al., 2014)

20 periods any dependence on any starting state (in the feasible region) is visually imperceptible. The threshold level of chub the manager seeks to avoid lies at the lowest shown level of chub, along the top of both panels in Figure 2. To avoid this threshold, the manager maintains the chub stock fairly far away from this minimum, typically in the 6,000 to 8,500 range.⁹ Likely population levels are much safer than the level specified by the viability goal: the risk-to-go, $r(\cdot)$, is much lower than $1 - \Delta$ in most of the feasible region. See Appendix C for a full plot of the risk-to-go function.

The density in Figure 2B arises from a combination of stochastic population dynamics and policy responses. In the southwest region, under zero control, the joint population will tend to the northeast, i.e. increasing trout and then declining chub states. This movement will bring the population into a zone of active control which will first push expected trout levels down (to the west) which then leads to expected expansion in the chub population (to the south).

5.2 Shadow Value

In the baseline specification of our model described above, we find that the penalty for hitting the threshold that is just sufficient for inducing the manager to meet the viability constraint wherever feasible is $\Omega^* = \$380\text{M}$. Under our discount rate of 3% this is equivalent to an annual perpetuity penalty (cost) of \$11M, which significantly exceeds the annual cost of maximum action (\$450K). The intuition of our solution approach is that in order to justify the action necessary to meet the viability goal wherever feasible, the manager acts *as if* the one-time cost of hitting the threshold level of chub is \$380M (or, equivalently, \$11M per year forever). We would not necessarily expect this figure to match a non-market valuation of this degraded population level. Rather it is the valuation that is consistent with the given viability goal. It does provide the basis for an important check: while absolute valuation for a particular species may be costly to conduct, a simpler (and useful) question is whether the “actual” value lies demonstrably above or below the implicit value that is consistent with the viability goal. If above, then the policy identified here is economically justified (at minimum).

As described in Section 4, the penalty (current shadow value) associated with hitting the threshold propagates across the state space conditional on the optimal policy and the stochastic dynamics. To illustrate this, in Figure 3 we show the value function decomposed into two components: (1) the expected present cost of optimal control, EPC and (2) the present shadow value function, $\omega(X_0, Y_0)$. The figure shows variation in these com-

---

⁹After 20 years under the optimal policy, the mean population levels for trout and chub, respectively, are (1700, 7400) with standard deviations of (700, 1300).

ponents of the value function over chub levels given the trout population is fixed at its modal level under the optimal policy. Mathematically, the present shadow value component is equal to the present expected penalty, i.e. the penalty from hitting the threshold diminished by both discounting and less-than-certain chance of hitting the threshold. As shown in Figure 3, at low chub levels this present shadow value rapidly expands with further population decline, eventually dominating the EPC. In the other direction, as the possibility of non-viability becomes less salient, with increasing chub and/or decreasing trout, the present shadow value ebbs.

The influence of the present shadow value on policy is evident. In the example case presented in Figure 3 this variable accelerates as chub levels decline past 6,000. Returning to Figure 2A we see that (at the modal trout level) the optimal policy cycles quickly from no action to intermediate and then maximum action as chub levels decline from 6,000 to 5,000.

Figure 3: Decomposition of the value function into (1) expected present costs of optimal control and (2) the present shadow value. The function varies over chub levels (Y) with the trout level fixed at its mode under the optimal policy ($X = 1400$). (Chub level = adult chub equivalent population in the mainstem.)

Similar to Montgomery et al. (1994), we explore how elements of the solution change with respect to a change in the confidence level (Δ). The direction of the effect is not straightforward due to two competing effects: while we expect more effort as a result of a

stricter constraint, the viability kernel also shrinks since space that was just viable under a lower confidence level no longer is. All else equal, we would expect less effort with a smaller viability kernel. Taking these two effects together, we would not necessarily expect the optimal penalty (present shadow value) nor the optimized control cost to change monotonically as the confidence increases. Indeed we find that these outcomes oscillate as confidence increases. However, the overall trend is as expected as we increase the required confidence level from 0.5 to near certainty. At the median level of chub and trout, we find that the [absolute] present shadow value (ω) increases on average by \$5.7K per percentage point increase in confidence. The expected present cost of control increases by \$10K per percentage point increase in confidence. In general this cost effect tends to be higher for less favorable population levels (lower chub and higher trout) and lower for more favorable population levels.

6 Parametric Uncertainty

In the baseline model analysis above, we incorporate stochastic population dynamics but model parameters are assumed to be fixed and known. In this section we examine the implications of uncertainty in λ , a parameter that determines the impact of trout abundance on juvenile chub survivorship and thus is central to joint chub-trout dynamics. We first consider a model in which uncertainty in the parameter is unchanging. Later, we consider a manager who expects to reduce this parametric uncertainty through learning. We are concerned with (1) how initial uncertainty affects current management, (2) how the anticipated arrival of information affects current management, and (3) the value of information (VOI) with respect to a reduction in the uncertainty of λ .

In our setting, λ is estimated from an ongoing mark-recapture study of the chub and trout populations. Rainbow trout abundance and juvenile chub survival are both imperfectly observed because only a fraction of the population is observed in each research trip and because various other factors, including environmental factors like river temperature and turbidity affect both capture probability and survival rates of humpback chub (Yackulic et al., 2018). Over time, the expectation is that estimates of λ will become more precise, but not necessarily perfect.

We represent the manager's beliefs (B^λ) over the uncertain parameter with a truncated normal distribution, $\lambda \sim tN(E[\lambda], Var[\lambda], [\underline{\lambda}, \bar{\lambda}])$. We set the mean equal to the deterministic value used in the baseline model in Section 5, $E(\lambda) = \lambda_0$. The variance is approximately based on the uncertainty estimated in Yackulic et al. (2018) from field data. Specifically, λ span from zero to twice the mean: $[\underline{\lambda}, \bar{\lambda}] = [0, 2\lambda_0]$. The manager's problem

is then modified to determine removal action given expected dynamics (and thus, payoffs) that integrate over possible values of this predation parameter.

6.1 Management under Fixed Parametric Uncertainty

In Figure 4 we present the policy function given fixed parametric uncertainty over λ . We find several differences relative to the deterministic λ base case from above in Figure 2. First there is no longer a region in the upper right corner in which it is impossible to avoid the threshold, since there is some chance that λ may be lower than previously assumed. Second, there is an expansion to the state space above the spanning dashed line in which the meeting the joint chance constraint is *not* feasible. This is because population dynamics are generally *less favorable* to chub given symmetric uncertainty over λ . This stems from the nonlinearity in the juvenile chub survival function in Equation 8: a given increase in λ has a more deleterious effect on chub survival than the advantage conferred by a decrease in λ of the same magnitude.

These less favorable chub dynamics drive a key change in the policy function: optimal control becomes more precautionary. The transition zone from no control to maximum control (the multi-shaded band in Figure 4) shifts to the southwest. Commensurately the chub-trout population density after 20 years (concentric circles) shifts in the same direction.¹⁰ Overall, while accounting for uncertainty in the effect of trout on chub leads to *less favorable* expected dynamics for chub, precautionary optimal management leads to maintenance of *greater* chub levels than are maintained in the baseline model.

6.2 Management with Learning and the Value of Information

In our learning scenario, we assume that ongoing research will provide a more precise understanding of the true parameter value after a specific number of years have passed. The arrival of information serves to tighten the distribution over plausible values for λ . This information is generated by a research program such that the timing of information arrival is predetermined and exogenous. This learning structure is motivated by and parameterized to be consistent with monitoring program for chub and trout in this system.

Initial beliefs, B_0^λ , are the same as in the fixed uncertainty in Section 6.1. The manager also knows that information to update these beliefs will arrive at some known number of periods in the future, T_{info} . We assume that learning is not complete, i.e. there will still be some uncertainty (though reduced) about the true parameter value. Upon learning, beliefs evolve to $B_{T_{info}}^\lambda$. This involves updates to both the expected value, $E_{T_{info}}[\lambda]$,

---

¹⁰The cumulative density at 50% (outer ring) no longer dips below 6,000 chub nor exceeds 2,400 trout.

Figure 4: The policy function under uncertain λ showing the optimal number level of control in a year given the current state (chub and trout levels) for the case of $E(\lambda) = 0.0009$ and $\sigma(\lambda) = 0.00045$. The concentric curves show level sets of the joint density (10%, 25%, and 50%) after 20 years under the optimal policy. The dashed line spanning the figure from left to right delineates the upper boundary of the viability kernel. (Chub level = adult chub equivalent population in the mainstem.)

and the variance, $Var_{T_{info}}[\lambda]$. We assume that learning drives a decrease in the variance from $\lambda_0^2/4$ to $\lambda_0^2/9$. We assume new information arrives in $T_{info} = 5$ years. These assumptions are motivated by annual monitoring of juvenile chub and consideration of the interannual variability of environmental (temperature and turbidity) and population factors (intraspecific density dependency and rainbow trout abundance) influencing juvenile chub survival (Yackulic et al., 2018).

The updated, post-learning mean, $E_{T_{info}}[\lambda]$, is not known at $t = 0$. Because the management model must be solved for *every* potential value of $E_{T_{info}}[\lambda]$, we specify a discrete set of possibilities. We assume that updated beliefs are also symmetric-truncated-normal with a range that lies within the domain of the initial belief nodes. Combining an assumed number of nodes to represent initial beliefs with the assumptions above, returns a set of possible values for $E_{T_{info}}[\lambda]$.¹¹

¹¹We represent initial beliefs with 21 nodes. Post-learning beliefs shrink to 15 nodes given the decline in variance. Given pre- and post-learning variance levels, this results in seven possible levels for $E_{T_{info}}[\lambda]$.

6.2.1 A Solution Algorithm under Learning

We solve the optimization with learning model in two stages:

1. First, we use value function iteration (see Section 4) to solve the problem from time T_{info} forward, given any potential state (X, Y) and *for every possible new set of beliefs about λ* given by the new information: $V_{T_{info}}(X, Y|B_{T_{info}}^\lambda)$. See Appendix D for examples of optimal policy functions under a range of updated belief states.
2. Second, conditional on the optimal responses at time T_{info} determined above, we use backward induction from T_{info} to solve the remaining problem, $V_0(X, Y)$ (toward the present period), given initial beliefs B_0^λ . Starting from T_{info} , we calculate $E_B[V_{T_{info}}(X, Y|B_{T_{info}}^\lambda)]$, where each contribution to the expectation, $V_{T_{info}}(X, Y|B_i)$, is weighted by the prior likelihood, $Pr(B = B_i)$. Then we use repeated application of the Bellman operator in order to find $V_0(X, Y)$ and then the optimal policy at time $t = 0$.

6.2.2 Learning Results

Here we explore the value and policy impacts of learning to improve our understanding of λ . Figure 5 shows the optimal policy function for the manager anticipating some information event at T_{info} periods in the future.¹² The dashed line, marking the feasibility of our viability constraint, has not changed from the non-learning case (i.e. the “defensible” region has not changed). Once the learning manager averages over all possible future beliefs (necessary when considering the continuation value at T_{info}), the expected dynamics under the learner’s policy are identical to that of the non-learner. This can be seen, for example in the case of applying maximum action, by evaluating the risk-to-go (i.e. likeli-

---

The probability of transitioning to any of these levels is given by the cumulative probability (under initial beliefs) that lies in the range of any given updated belief. We also considered more dense discretization of beliefs; this did not change the results but increased computing time. Numerically, initial beliefs are characterized by $E_0(\lambda) = 9.0 \times 10^{-4}$ and $Var_0[\lambda] = 2.0 \times 10^{-7}$, while post-learning beliefs have a variance of $Var_{T_{info}}[\lambda] = 9.0 \times 10^{-8}$ with a mean in the range $[6.3 \times 10^{-4}, 1.2 \times 10^{-3}]$.

¹²More precisely, it is the policy function at time zero, the beginning of the program. The reader might recognize that finite-time horizon problems will admit a policy function for each year before T_{info} . In our numerical exercise, this policy function changes very little from periods 0-5. Of course, post-learning, the policy function may dramatically change. Appendix D provides the post-learning policy function for three possible posterior beliefs.

hood of having hit the population threshold by time T):

$$\begin{aligned}
 r(Z_0 | \max\{\mathcal{A}\}) &= 1 - \Pr \left\{ \bigcap_{t=1}^T (Y_t > \bar{Y}) \middle| Z_0, \max\{\mathcal{A}\} \right\} \\
 &= E_Z \left[1_{Y_T = \bar{Y}}(z) \middle| Z_0, \max\{\mathcal{A}\} \right] \\
 &= E_Z \left[E_B \left[1_{Y_T = \bar{Y}}(z) \middle| Z_0, B_{T_{info}}^\lambda, \max\{\mathcal{A}\} \right] \right],
 \end{aligned} \tag{19}$$

where z represents the state space $\{x, y\}$, $1_{Y_T = \bar{Y}}(z)$ is the indicator function and the last line follows from the Law of Iterated Expectations. Further detail on calculating the risk-to-go function is provided in Appendix C.

Figure 5: The policy function under uncertain λ for the case of the learning manager. Shading indicates how many mechanical removals of trout are optimal in a year given the current populations of trout (horizontal axis) and chub (vertical axis) in the management reach. The dashed line spanning the figure from left to right delineates the upper boundary of the viability kernel. The concentric curves show level sets of the joint density (10%, 25%, and 50%) after 20 years under the optimal policy. The solid line spanning the figure from left to right (towards the bottom) shows the southern extent of the state space over which the non-learner imposes a nonzero level of control. (Chub level = adult chub equivalent population in the mainstem.)

The solid line spanning from left to right—a new addition to Figure 5 not appearing in previous policy figures—marks the southern extent of the region in which the *non-*

learning manager (from Figure 4) applies a positive level of control. Thus, in Figure 5, the zone stretching from the spanning solid line up to where the learner reaches maximum control is where the learner and non-learner policies differ. Specifically this is where the learner exerts less control, experiences lower costs, and yet still meets the same viability goal. The learner's relaxed approach is consistent with the finding from Section 6.1 that uncertainty in λ induces a more precautionary policy. Thus, an anticipated reduction in uncertainty reduces the incentive for that elevated precaution.

Lastly, we are interested in the value of information to the manager, i.e. the expected benefit from the anticipated and then realized reduction in the uncertainty over λ . This value can be evaluated by taking the difference in the expected present cost of optimal control for two programs, one with improved knowledge, and the other without:

$$EVOI = E \left[\sum_{t=0}^{\infty} \delta^t C(A_t) \right]_{no\ info} - E \left[\sum_{t=0}^{\infty} \delta^t C(A_t) \right]_{info}. \quad (20)$$

We expect the *EVOI* to be non-negative since the manager should do no worse with additional information. We present the *EVOI* in Figure 6A as a function of any given current population state. The *EVOI* reaches a high of around \$600K in the region where the learning manager's policy differs from the non-learner's policy (bounded by the spanning solid line in Figure 5). The *EVOI* is zero at the population threshold and low in the northeast corner of Figure 6A where chances of declining to the population threshold are high (thus forcing maximum action regardless of the information set).

Figure 6: Panel A shows the *EVOI* (\$100Ks). Panels B and C decompose the *EVOI* into the prospective and congruity VOI, respectively. (Chub level = adult chub equivalent population in the mainstem.)

We decompose the *EVOI* into two components. The first is the change in expected management cost before learning takes place while the manager is anticipating learning, which we call the “prospective” VOI. In general this value could be positive or negative. In our setting, across the state space it is likely that this value will be either close to zero (where the learning and non-learning managers policies are the same) or positive (where the learning manager is less aggressive as depicted in Figure 5). When the manager anticipates learning there is less need in the short run to hedge against the worst-case scenario, which can be addressed if revealed to be the true state of world. This is indeed the case for the prospective VOI shown in Figure 6B.

The second component, which we call the “congruity” VOI, is the expected post-learning value of information that stems from less error in optimal management (since beliefs are more tightly concentrated on the truth). In general this value will be non-negative. The congruity VOI, shown in Figure 6C, is the larger portion of the *EVOI* (roughly 80%) since the learning horizon is short relative to the viability horizon ($T_{info} = 5$ years). The congruity VOI is also relatively uniform over population levels since this value accrues after some delay (T_{info} years) and over the long run, in which we would expect to move substantially around the state space. The exception is the northeast corner of the figure where chances of declining to the population threshold are high.

7 Discussion

In this paper, we demonstrate a new dynamic programming algorithm for solving chance-constrained problems. Unlike previous approaches, our method does not rely on simulation. We provide a solution for the joint-chance constrained optimization problem without the need to approximate the chance constraint or segregate stochasticity. The use of dynamic programming allows us to recover a policy function that is too complex to obtain via a parametric simulation approach. The distribution of population outcomes is endogenous with respect to this policy, rather than chosen ex ante. We also gain insight into the implied economic value of a viable chub population. Previous penalty-based approaches have assumed penalty values for violating a viability objective. However our method endogenizes the penalty level, revealing an implied valuation of the chub threshold given the manager’s risk preferences (from the viability goal) and cost-effective trout control.

An important simplification in our model is the condensing of two sub-adult and two adult chub size classes into a single state variable. While this is common in the broader fisheries literature it is also known that accounting for age classes can be important for

management (Tahvonen, 2009). Modeling chub as proceeding from juveniles directly to the adult population of interest sets aside additional sources of uncertainty in chub dynamics that could be important. This approach also means that control choices are quickly represented in the adult population. Incorporating the chub size class model of Yackulic et al. (2014)—i.e. delayed returns from control and additional uncertainty—would likely lead to a more conservative policy function in which control is initiated sooner as a robust chub population declines or a low trout population increases. Our solution method would still serve in such a model though the state space would expand to five continuous variables. Recent advancements in dynamic programming methods for resource management problems with such high numbers of state variables (see e.g. Springborn et al. 2018) should allow for such a model.

We find that the present expected cost of control of maintaining a viable chub population over an infinite time horizon is \$4.6M, 45% lower compared to previous findings by Bair et al. (2018). This difference arises partially because our optimal policy is fully flexible (not parametrically constrained) and depends not only on trout abundance, but chub abundance as well. To illustrate, when chub are in great numbers and trout are at a moderate 1,500 individuals, their policy prescribes maximum action, while ours suggests none.¹³ Our policy of comparison is that of a non-learning manager; further reducing the expected cost of the removal program is the added benefit of new information.

We find that the learning manager’s current shadow value—the penalty incurred for violating the viability threshold of 4,000 chub—is \$330M. Under the optimal management strategy, this outcome is unlikely to occur. Valuing a different marginal change, Duffield et al. (2016) estimate the aggregate U.S. household marginal value for only a 1% increase in humpback chub abundance is \$130M. If we set the trout level to its modal value under the optimal policy, we find a 1% increase in chub abundance implies a change in the present shadow value of \$170M, \$62K, and \$5K at initial chub levels of 4,000, 7,000, and 12,000, respectively. Thus the value implied by the viability goal only reaches the existing valuation estimate as we approach the threshold.

We provide the first study of optimal adaptive management in a joint-chance constrained problem. There are substantial returns to learning. The expected present cost of the viability program is reduced by 11% compared to a program run by a manager fac-

---

¹³An extensive quantitative comparison between the results of Bair et al. (2018) and our results should be considered with caution due to key differences in modeling. First, we update parameters of population dynamic equations to reflect the latest science. Second, Bair et al. only tracks the trout population while we track both chub and trout. Finally, and most importantly, the specification of the viability goal—the central constraint—is quite different. Bair et al. focus on the average annual survival rate of juvenile chub. Here we instead focus explicitly on the likelihood of adult chub remaining on the safe side of the population threshold.

ing reducible uncertainty but not receiving any new information. Most of this cost savings is due to finer-tuning of trout management in response to new information, rather than the anticipation of the information itself. A similar idea has been explored by Costello et al. (1998), who consider the value of El Niño forecasts in the west-coast coho salmon fishery versus a world without them. They find returns to accurate forecasts at around 1.5% of net present benefits of the fishery. The discrepancy in returns is likely due to the relative importance of the source of reducible uncertainty in either model. In our case, the effect that trout have on chub survival is much more salient than the effect of El Niño on the coho salmon fishery.

The economic assessment developed here provides an initial introduction to chance-constrained dynamic programming and environmental decision making. Further opportunities for applying this method include climate change mitigation, and invasive species and disease control. An optimal emissions schedule or disease control program would reflect the desire to reduce the chance of further warming over a dangerous threshold or an outbreak of disease, respectively. Fundamental components of these problems are the centrality of jointly uncertainty outcomes and the difficulty of estimating the cost of a disaster scenario.

Future research in this vein could incorporate viability constraints in larger bioeconomic problems. Immediately related to our setting, the management of water flows in the Colorado River as an action to control invasive species involves hydropower and recreational values in addition to the impact on species viability. In order to assess the greater management problem, the development of an optimal chub management strategy would be considered as a subset of the broader management strategy. More relevant to marine resources is the management of commercial harvest of healthy stocks in the presence of by-catch of degraded species, contrasting market-derived values with concrete, yet invaluable conservation obligations.

Appendix

A Timeline of biological events

The timeline of biological events in each decision period is given in Table A.1. While this simplified sequence of events is stylized, it includes the essential features of joint-trout-chub dynamics in the mainstem.

Table A.1: Timeline of biological events in each decision period.

Order	Event	Updates
1	Trout and chub stocks are observed	X_t, Y_t
	Trout removal decision is made, cost is realized	$A_t, C(A_t)$
2	New recruitment (& migration) is generated	x_t, y_t
3	Existing trout stock affects juvenile chub survival	$s_Y(X_t)$
4	Trout removals undertaken	$s_X(A_t)$
5	Total trout stock natural survivorship	γ
	Adult chub stock natural survivorship	η
6	Complete updating of the two stocks	X_{t+1}, Y_{t+1}

B The dynamic programming solution algorithm

1. Define parameters, functional forms and discretization of state and action space
  - (a) biological and economic parameters, state transition equations
  - (b) viability goal given by the triplet $\{\bar{Y}, \Delta, T\}$
  - (c) state and control sets $\{x, y, A\}$, $(\inf(y) = \bar{Y}, \text{an absorbing state})^{14}$
2. Specify or initialize dynamic programming routine arrays and functions
  - (a) Markov transition matrix $P(X_{t+1}, Y_{t+1} | X_t, Y_t, A_t)$
  - (b) “risk-to-go” function $r(X_0, Y_0 | A(x, y))$
  - (c) identify the *viability kernel* (De Lara and Doyen, 2008), i.e. the subspace where viability goal is feasible at desired confidence Δ , $\{x, y\}_f$ s.t. $r(\{x, y\}_f | \max\{A\}) \leq 1 - \Delta$

---

¹⁴Generically, even if $\bar{Y}$ is not truly an absorbing state, it is useful to model it as such in the solution algorithm since we are concerned with any obtainment of $\bar{Y}$ (failure) in a time path and thus capturing the first such obtainment is sufficient and simple.

- (d) initial (large) bounds for cost of non-viable population (penalty), $[\Omega_L, \Omega_H]$
- 3. Guess a value for penalty, $\Omega = \text{mean}(\Omega_L, \Omega_H)$. Fix the value function at the absorbing state to the penalty level, $V^\Omega(x|\bar{Y}) = -\Omega$, and solve for the value function $V^\Omega(x, y)$ and optimal policy $A^\Omega(x, y)$ via value function iteration (dynamic programming step)
- 4. Check if viability constraint is met within the viability kernel (see e.g. Figure C.1)
  - (a) if $r(\{x, y\}_f | A^\Omega(x, y)) \leq 1 - \Delta$ everywhere, then $\Omega_H := \Omega$, since $\Omega^* \leq \Omega$
  - (b) if $r(\{x, y\}_f | A^\Omega(x, y)) > 1 - \Delta$ somewhere, then $\Omega_L := \Omega$, since $\Omega^* \geq \Omega$
  - (c) if within desired tolerance, $[\Omega_H - \Omega_L] < \epsilon$, then $\Omega^* := \Omega_H$ and terminate
  - (d) otherwise, repeat 3-4

C Calculation of the risk-to-go function

Having solved for the optimal policy $A^\Omega(x, y)$, we can produce a transition matrix conditional on this optimal policy, $P^\Omega(X_{t+1}, Y_{t+1} | X_t, Y_t)$, which for this section we will simply call P . Further, the $\{x, y\}$ state space can be rolled into the symbol z .

For a row-stochastic (rows sum to 1) matrix P , an initial distribution Z_0 (1-by- n) and real-valued function $f(z)$ (n -by-1), we can describe the expectation of f in period T as

$$E[f_T(z)] = Z_0 P^T f(z) \quad (\text{C.1})$$

Decomposing the right-hand side of equation (C.1), $Z_0 P^T$ represents the distribution of states (1-by- n) at time T , Z_T , which will prove useful in showing the 20-year joint probability mass function earlier in the text. $P^T f(z)$ describes the expected evolution of the function (n -by-1) after the same number of steps, $E[f_T(z) | Z_0]$.

We want to evaluate the probability of an event (non-viability) happening, conditional on starting state: $E[f_T(z) | Z_0] = 1 - \Pr\left\{\bigcap_{t=1}^T (Y_t > \bar{Y}) | X_0, Y_0, A(x, y)\right\}$, which we will simplify to $\Pr(Y_T = \bar{Y} | Z_0)$.¹⁵ Taking the expectation of the indicator function $1_{Y_T = \bar{Y}}(z)$ provides us with the probability that our threshold is hit within T periods:

$$E[f_T(z) | Z_0] = E[1_{Y_T = \bar{Y}}(z) | Z_0] = \Pr(Y_T = \bar{Y} | Z_0) \quad (\text{C.2})$$

---

¹⁵Recall that $\bar{Y}$ is an absorbing state, which eliminates the possibility of double-counting a crossing of the threshold.

By choosing $f(z) = 1_{Y=\bar{Y}}(z)$, we can evaluate the degree of confidence in which we achieve the viability target, conditional on the optimal policy and the starting state.

In Figure C.1 we show the risk-to-go, $r(\cdot)$, i.e. the probability of reaching the chub population threshold level given any current chub-trout state. Here again, the spanning dashed line traces the states where we can just meet the joint chance constraint (viability goal): $r(\cdot) = 1 - \Delta$. The risk-to-go is quite low in most of the state space, where $r(\cdot)$ is substantially below $1 - \Delta$. For example, towards the bottom-left of Figure C.1 we have $r(x = 1000, y = 8,000) = 0.1\%$.

Figure C.1: The risk-to-go—probability of reaching the chub population threshold level—within 20 years from any initial starting state and conditional on the optimal policy. For most states, the viability goal is met or exceeded. The dashed line spanning the figure from left to right delineates the upper boundary of the viability kernel at the specified confidence level, $\Delta = 90\%$.

D Post-learning policy functions

Figure D.1: Optimal policy functions under different post-learning distributions of λ where variance falls to $Var_{T_{info}} = \lambda_0^2/9$. The updated mean, $E[\lambda]$, varies across panels, decreasing from low to high from left to right. The dashed line spanning each figure from left to right delineates the upper boundary of the viability kernel. The concentric curves represent level sets of the joint density function after 20 years under the optimal policy and depict cumulative densities of 10%, 25%, and 50%.

References

- Alais, J.-C., Carpentier, P., and De Lara, M. (2017). Multi-usage hydropower single dam management: chance-constrained optimization and stochastic viability. *Energy Systems*, 8(1):7–30.
- 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*, 220:12–20.
- Baumgärtner, S. and Quaas, M. F. (2009). Ecological-economic viability as a criterion of strong sustainability under uncertainty. *Ecological Economics*, 68(7):2008–2020.
- Boyle, K. J. and Bishop, R. C. (1987). Valuing wildlife in benefit-cost analyses: A case study involving endangered species. *Water resources research*, 23(5):943–950.
- Brown, G. M. and Shogren, J. F. (1998). Economics of the endangered species act. *The Journal of Economic Perspectives*, 12(3):3–20.
- Costello, C. J., Adams, R. M., and Polasky, S. (1998). The value of El Niño forecasts in the management of salmon: a stochastic dynamic assessment. *American Journal of Agricultural Economics*, 80(4):765–777.
- Coulson, T., Mace, G. M., Hudson, E., and Possingham, H. (2001). The use and abuse of population viability analysis. *Trends in Ecology & Evolution*, 16(5):219–221.
- De Lara, M. and Doyen, L. (2008). *Sustainable management of natural resources: mathematical models and methods*. Environmental science and engineering Environmental science. Springer, Berlin. OCLC: 263412377.
- Doyen, L. and De Lara, M. (2010). Stochastic viability and dynamic programming. *Systems & Control Letters*, 59(10):629–634.
- Doyen, L., Thébaud, O., Béné, C., Martinet, V., Gourguet, S., Bertignac, M., Fifas, S., and Blanchard, F. (2012). A stochastic viability approach to ecosystem-based fisheries management. *Ecological Economics*, 75:32–42.
- Duffield, J., Duffield, C., and Patterson, C. (2016). Colorado river total value study. *Final Report, Prepared for the National Parks Service*.
- Finseth, R. M. and Conrad, J. M. (2014). Cost-effective Recovery of an Endangered Species: The Red-cockaded Woodpecker. *Land Economics*, 90(4):649–667.
- Haight, R. G. (1995). Comparing extinction risk and economic cost in wildlife conservation planning. *Ecological Applications*, 5(3):767–775.
- Judd, K. L. (1998). *Numerical Methods in Economics*. MIT Press.

- Korman, J., Martell, S. J., 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. and Yard, M. D. (2017). Effects of environmental covariates and density on the catchability of fish populations and interpretation of catch per unit effort trends. *Fisheries Research*, 189:18–34.
- 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.
- Kotchen, M. and Reiling, S. (2000). Environmental attitudes, motivations, and contingent valuation of nonuse values: A case study involving endangered species. *Ecological Economics*, 32:93–107.
- Lampert, A., Hastings, A., Grosholz, E. D., Jardine, S. L., and Sanchirico, J. N. (2014). Optimal approaches for balancing invasive species eradication and endangered species management. *Science*, 344(6187):1028–1031.
- LaRiviere, J., Kling, D., Sanchirico, J. N., Sims, C., and Springborn, M. (2018). The Treatment of Uncertainty and Learning in the Economics of Natural Resource and Environmental Management. *Review of Environmental Economics and Policy*, 12(1):92–112.
- Loomis, J. B. and White, D. S. (1996). Economic benefits of rare and endangered species: summary and meta-analysis. *Ecological Economics*, 18(3):197–206.
- Margolis, M. and Nævdal, E. (2008). Safe Minimum Standards in Dynamic Resource Problems: Conditions for Living on the Edge of Risk. *Environmental and Resource Economics*, 40(3):401–423.
- Marshall, E., Homans, F., and Haight, R. (2000). Exploring Strategies for Improving the Cost Effectiveness of Endangered Species Management: The Kirtland’s Warbler as a Case Study. *Land Economics*, 76(3):462–473.
- Montgomery, C. A., Adams, D. M., and others (1994). The marginal cost of species preservation: the northern spotted owl. *Journal of Environmental Economics and Management*, 26(2):111–128.
- Newbold, S. C. and Siikamäki, J. (2009). Prioritizing conservation activities using reserve site selection methods and population viability analysis. *Ecological Applications: A Publication of the Ecological Society of America*, 19(7):1774–1790.
- Ono, M., Pavone, M., Kuwata, Y., and Balaram, J. (2015). Chance-constrained dynamic programming with application to risk-aware robotic space exploration. *Autonomous Robots*, 39(4):555–571.

- Oubraham, A. and Zaccour, G. (2018). A Survey of Applications of Viability Theory to the Sustainable Exploitation of Renewable Resources. *Ecological Economics*, 145:346–367.
- Pe'er, G., Matsinos, Y. G., Johst, K., Franz, K. W., Turlure, C., Radchuk, V., Malinowska, A. H., Curtis, J. M., Naujokaitis-Lewis, I., Wintle, B. A., and Henle, K. (2013). A Protocol for Better Design, Application, and Communication of Population Viability Analyses. *Conservation Biology*, 27(4):644–656.
- Richardson, L. and Loomis, J. (2009). The total economic value of threatened, endangered and rare species: An updated meta-analysis. *Ecological Economics*, 68(5):1535–1548.
- Rout, T. M., Hauser, C. E., and Possingham, H. P. (2009). Optimal adaptive management for the translocation of a threatened species. *Ecological Applications: A Publication of the Ecological Society of America*, 19(2):515–526.
- Sagoff, M. (2009). Regulatory Review and Cost-Benefit Analysis. *Philosophy and Public Policy Quarterly*, 29(3/4):21–26.
- Springborn, M. R., Faig, A., Baskett, M., and Dedrick, A. (2018). *Beyond biomass: valuing genetic diversity in natural resource management*. Published: UC Davis working paper.
- Tahvonen, O. (2009). Economics of harvesting age-structured fish populations. *Journal of Environmental Economics and Management*, 58(3):281–299.
- U.S. Department of the Interior, National Park Service, Intermountain Region (2016). Biological assessment for the Glen Canyon long-term experimental and management plan (LTEMP).
- Ward, D. L., Morton-Starnner, R., and Vaage, B. (2016). Effects of Turbidity on Predation Vulnerability of Juvenile Humpback Chub to Rainbow Trout and Brown Trout. *Journal of Fish and Wildlife Management*, 7(1):205–212.
- Welsh, M., Bishop, R., Phillips, M., and Baumgartner, R. (1995). GCES non-use value study. *Report prepared for the Glen Canyon Environmental Studies Non-Use Value Committee, US Bureau of Reclamation*.
- Wilcove, D. S., Rothstein, D., Dubow, J., Phillips, A., and Losos, E. (1998). Quantifying Threats to Imperiled Species in the United States. *BioScience*, 48(8):607–615.
- Yackulic, C. B., Korman, J., Yard, M. D., and Dzul, M. (2018). 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., Jr, L. G. C., 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.