How and When to Measure Renewable Resources under State Uncertainty

(Previous title: “Measure, harvest, learn: renewable resource management under state uncertainty”)

David M. Kling[a,*], Paul L. Fackler[b], Michael R. Springborn[c]

December 20, 2019

Preliminary and Incomplete – Please Do Not Cite or Circulate

Abstract

Natural resource management involves acquiring and using information about the stock being managed. Assessment of available stocks is subject to multiple forms of uncertainty, even when budgetary and technical resources are available. Here the focus is on uncertainty in the level of population of a harvested species. We allow for learning to reduce uncertainty where information may be generated through resource use as well as assessment independent of resource use. From an economic perspective, investment in stock assessment should be weighed against its expected benefit in terms of improved management performance. The manager chooses the level of resource use and independent assessment in order to maximize the expected net present value of rents. In the process, the manager takes into account the value of information in the form of more precise stock level estimates. We find that the dynamics of learning and opportunities for policy experimentation in our model differ substantially from simpler cases where information involving harvesting data is unused and in which independent assessments are always available at no cost. Our model provides insight into the economic value generated by different levels of investment in learning, in particular the return on investment in information. We illustrate the model with a numerical example of the joint harvest and stock assessment problem in fisheries.

Keywords: State uncertainty, learning

[*] david.kling@oregonstate.edu

[a] Department of Applied Economics, Oregon State University

[b] Department of Agricultural and Resource Economics, North Carolina State University

[c] Department of Environmental Science & Policy, University of California, Davis

1. Introduction

Renewable resource management involves acquiring and using information about managed stocks. For example, groundwater reserves and recharge rates must be measured. Current abundance estimates for many threatened species, even well-known species like African elephants, are subject to considerable uncertainty (Chase et al. 2016). In fisheries management, a standard practice is to coordinate regulations like harvest limits with a stock assessment, the aim of which is to predict how management choices will affect future yield and stock abundance. Even when budgetary and technical resources exist to produce a cutting-edge assessment, results are subject to substantial uncertainty (Ralston et al. 2011). Disease management focuses on early detection and measurement of infectious disease outbreaks, in which the dynamics of susceptible and infected populations can be challenging to track (MacLachlan and Springborn 2016).

State uncertainty describes a decision problem involving a key state variable, like the abundance of a wild population, which is either observed with error or not observed at all in more or more periods. While state uncertainty has been recognized as being fundamental to most renewable resource management problems, methods for modeling optimal management that account for state uncertainty have been slow to develop (LaRiviere et al. 2018). Recently progress has been made so that standard bioeconomic models can be analyzed after introducing state uncertainty (MacLachlan et al. 2016, Kling et al. 2017, Memarzadeh et al. 2019, Sloggy et al. 2019). Among the insights from these models are findings on optimal investment in learning by measuring the uncertain state variable, as well as the potential cost of applying policies based on conventional perfect-observability models.

Despite considerable recent progress, state uncertainty in most analyses is considered in a limited sense. One common simplification is that state variable measurements are received for free in every decision period, perhaps from an established monitoring program with negligible non-sunk costs. Also, only one means of measurement is typically considered – i.e., at most one type of observation is available to use. In practice, there may be multiple methods available to obtain information about the status of a resource. A well-known example is marine fisheries stock assessment, which may incorporate “fishery-dependent” catch data and fishery independent measurements to produce an estimate of stock abundance. For problems that involve costly measurement through potentially more than one approach, cost-effective learning jointly with other management activities is not well understood. (Nichols and Williams 2006).

We study the when and the how of learning about an uncertain renewable resource in this paper. Our bioeconomic optimal harvest model features two types of investment in population measurement: acquiring and learning from catch per-unit effort (CPUE), or direct monitoring of population abundance. Neither measurement is freely received by the decision maker in our model, and it is possible to pay to learn from both in the same period. The resulting model is a continuous-state partially observable Markov decision process (POMDP). We develop and apply a novel numerical approximation approach that provides a tractable means of modeling how information available to the decision maker develops over time.

Our results show that optimal learning about an uncertain resource state can involve irregular investment in obtaining measurements. Depending on cost assumptions, alternative solution concepts that assume investment in measurement occurs every decision period may entail substantial over-investment in acquiring information. To the best of our knowledge, our analysis is the first to characterize optimal investment in both catch-based and catch-independent monitoring information as part of a single management strategy.

The paper is structured as follows. The review of our methods includes a description of the process, observation, and learning components of our model. Following a description of our approximation strategy and numerical solution procedure in Section 2, we present results on the optimal policy. We compare properties of our model under different assumptions for costs with alternative solution concepts in Section 3. Our concluding discussion discusses opportunities for future research.

2. Methods

We model the problem of a decision maker–the “resource manager”–who chooses allowable harvest and investment in measurement of a population. Dynamics of the population are stochastic and never perfectly observed, which introduces state uncertainty. Our model relaxes two key restrictions imposed in most related models that deal with state uncertainty. First, rather than always being free and/or exogenously given, optional population observations are costly to acquire. Second, the decision maker has two ways of investing in population measurement: learning from catch per-unit effort (CPUE) and/or learning from a monitoring activity that is independent of CPUE.

While our analysis nests a number of related dynamic optimization models, it involves a number of important simplifying assumptions. Like nearly all dynamic optimization models in resource economics that consider state uncertainty, we assume the bioeconomic process model (including the population state transition model and objective function) are known to the resource manager. In particular, we do not consider parameter uncertainty. Second, in order to connect our analysis to related work, we focus on a stylized management problem concerning a valuable harvested species such as wild game or marine fish, and do not consider related applications like invasive species management (Kotani et al. 2011). In our concluding discussion we preview ongoing work to generalize the model presented here to simultaneously consider state and parameter uncertainty.

Next, we present the model by introducing, in turn, the population model, observation (learning) model and the integrated decision problem. Finally, we describe the numerical solution method and present our application, including functional form assumptions and base parameter values.

2.1. Population Model

In each period t, the resource manager chooses the expected share of harvest of a stock X_t by selecting the level of a harvest effort control U_t. The realized amount of harvest H_t, is determined by a standard Gordon-Schaffer assumption (Gordon 1953, 1954) via the catch equation:

$$H_t = q\left(\xi_t^q; U_t\right)U_t X_t \tag{1}$$

Effort is normalized so that $U_t \in [0,1]$. Catch per-unit effort when U_t is nonzero therefore takes the following form:

$$CPUE_t = q\left(\xi_t^q; U_t\right)X_t \tag{2}$$

We follow the conventional assumption that realizations of harvest conditional on a non-zero effort level and the current stock level are stochastic by including a stochastic catchability term $q\left(\xi_t^q; U_t\right)$. CPUE is unlikely to have the same coefficient of variation (CV) (conditional on the true stock level) when effort is very low as it does when effort is very high. We build flexibility into this component of the model by representing realizations of catchability as draws from a parameterized density that is bounded on the unit interval, but shifts as a function of effort. In Eq. (1), ξ_t^q is an iid uniform draw, and $q(\cdot)$ is an inverse CDF function that admits U_t as a parameter restriction. In our application, we assume that $E\left[q\left(\xi_t^q; U_t\right)|U_t = U_i\right] = \bar{q}$ for all feasible effort choices U_i, and $CV\left[q\left(\xi_t^q; U_t\right)\middle|U_t = U_i\right]$ is a decreasing function of U_i. To simplify notation, in what follows we write $q\left(\xi_t^q; U_t\right) \equiv q(U_t)$. We include additional details on our numerical implementation of catchability and other densities involved in the model, as well as base parameter values, in Section 2.5 below. Overall, this structure encodes the notion that information gleaned from observing CPUE is less "noisy" the great the level of harvest effort.

Population growth is stochastic and is produced from escapement S_t following harvest: $S_t = X_t - H_t$. We adopt a stochastic Beverton-Holt model (Beverton & Holt, 1957) for the population state equation:

$$X_{t+1} = \left(\frac{(1+r)S_t}{1 + \left(\frac{r}{K}\right)S_t}\right)Z_t^g \tag{3}$$

In Eq. (3), r is a discrete time deterministic growth rate, K is the environmental carrying capacity (where expected future population equals escapement) and Z_t^g is an iid growth shock with positive support.

2.2. Observation Model

Observations of the stock are not free nor are they perfect. Instead, two types of noisy measurements of the stock are available. Each period the resource manager decides whether to incur the cost of obtaining one or both observations.

The first is CPUE (Eq. (2)). A control $L_t \in \{0,1\}$ is an indicator for the manager choosing to acquire and learn from CPUE in period t. The catchability shock $q_t(U_t)$ acts as an observation shock that prevents the resource manager from observing the population perfectly. Our assumption that the CV of $q_t(U_t)$ is decreasing in U_t means that higher levels of harvest effort produce a more informative estimate of X_t, assuming the manager also chooses to invest in learning from CPUE. When $L_t = 0$, the CPUE measurement is not received.

The second method of observing the population is monitoring. A continuous monitoring effort control variable $M_t \in [0,1]$, when non-zero, produces an observation Y_{t+1} received at the beginning of the next period:

$$Y_{t+1} = Z^m(\xi_{t+1}^m; M_t)X_{t+1} \tag{4}$$

We label Eq. (4) the monitoring observation equation. The strictly positive stochastic term $Z^m(\xi_t^m; M_{t-1})$ behaves similarly to catchability in terms of its relationship to the corresponding control variable choice. The function $Z^m(\cdot)$ is an inverse CDF function for a parametrized density that maps the iid uniform shock ξ_t^m to a draw from the corresponding density, subject to the constraint on the parameters imposed by the choice of M_{t-1}. We assume that $E\left[\frac{1}{Z^m(\xi_{t+1}^m; M_t)} \middle| M_t = M_i\right] = 1 \; \forall M_i \in (0,1]$ and that $CV[Z^m(\xi_{t+1}^m; M_t) | M_t = M_i]$ is decreasing in M_i over $(0,1]$. These restrictions serve two purposes. First, Y_t is an unbiased estimator of X_t (it can be shown that $E[X_t | Y_t, M_{t-1}] = Y_t$). Second, the informativeness of the monitoring observation increases as monitoring effort increases, all else being equal. For parsimony we define $Z_{t+1}^m(M_t) \equiv Z^m(\xi_{t+1}^m; M_t)$ in what follows.

In the next section we describe how both types of learning (via CPUE and/or monitoring) are incorporated into the resource manager's decision problem.

2.3. Decision Problem

The resource manager's objective is to maximize the expected discounted sequence of rents from harvest, net of costs of harvest effort and population measurements. The per-period reward function is:

$$\Pi(U_t, M_t, L_t) = R(q_t U_t X_t) - C_U(U_t) - C_M(M_t) - c_L L_t \tag{5}$$

Our assumptions for harvest rent, $R(\cdot)$, and costs are standard: rent is concave in $H_t = q_t U_t X_t$; $C_U(\cdot)$ and $C_M(\cdot)$ are each convex in the corresponding control variable; and c_L is a nonnegative constant. The resource manager applies a constant positive discount factor β that is strictly less than one.

Unlike most dynamic optimization problems in economics, the decision maker in our model does not observe the physical state, X_t, perfectly. Moreover, depending on the control choices made in the current and prior period, she may choose not to receive any measurements of the population in a given period at all. Instead, decisions must be made using an information state I_t. The information state includes the record of actions and observation received. Let an action be the control choices made in a particular period: $A_t \equiv (L_t, M_t, U_t)$. Then in period t, $\{A_\tau\}_{\tau=0}^{t}$ is included in I_t by definition. Observations $CPUE_t$ and Y_t also are included in I_t, but for any particular decision period they will not appear if the resource manager does not choose to invest in acquiring them (e.g., if $L_t = 0$, then $CPUE_t$ does not enter into the update of I_t to I_{t+1}).

Summing up, the resource manager's problem may be stated as follows:

$$\max_{\{U_t, M_t, L_t\}} \sum_{t=0}^{\infty} E_{I_t}[\Pi(U_t, M_t, L_t)]\beta^t \qquad (6)$$

subject to the process (Eqs. (1), (3), and (5)) and observation models (Eqs. (2) and (4)), along with our definition of I_t and an initial information condition I_0. A well-known property of this type of decision problem is that, in general, an analytical or exact numerical solution is unavailable. We describe our solution approach in the next section.

2.4. Solution Method

The resource manager's problem is an example of a continuous-state partially-observable Markov decision process (POMDP). Given the Markov structure of the process and observation models, optimal decisions may be made using a belief state $\varpi_t(X_t)$, which is a continuous univariate density over the range of the true population level, rather than the full information state I_t. With $\varpi_t(X_t)$, a belief-state dynamic programming problem may be posed and in principle solved to obtain an optimal feedback policy.

Unfortunately, for decision problems with one continuous uncertain state variable, the exact $\varpi_t(X_t)$ for the corresponding continuous-state POMDP may be any continuous univariate density. This presents a severe curse of dimensionality problem. Our approximate solution method builds on ideas introduced by Zhou, Fu & Marcus (2010) (*hereinafter* ZFM) and applied in several recent papers (MacLachlan et al. 2017, Kling et al. 2017, Sloggy et al. 2019). We approximate the belief state with $b_t(X_t)$, a univariate lognormal density that may be summarized in terms of its parameters. The resource manager's problem involves multiple types of observations, which may not be received at all in a given period. As a result, modeling the dynamics of $b_t(X_t)$ from one period to the next in a tractable way for the purpose of dynamic optimization is a challenging numerical problem. We describe what we believe is a new approach to modeling belief dynamics below.

Table 1. Qualitative description of the belief update cases

Control choices	Description	Outcome
$M = 0, L = 0$ (any U)	Do not pay to use a population measurement to update beliefs.	Prior belief $b(X)$ is propagated through the population dynamics conditional on U to obtain posterior $b^+(X^+)$.
$M = 0, L = 1$ $(U > 0)$[a]	Pay to measure CPUE.	1. Prior belief is updated based on $CPUE$ observation (Eq. (2)) conditional on U. 2. Updated beliefs are propagated through population dynamics conditional on U to obtain posterior belief state.
$M > 0, L = 0$ (any U)	Pay to monitor population.	1. Prior belief is propagated through the population dynamics conditional on U. 2. Monitoring observation Y^+ (Eq. (4)) is used in update to posterior belief state conditional on M.
$M > 0, L = 1$ $(U > 0)$	Pay to measure CPUE and monitoring population.	1. Prior belief is updated based on $CPUE$ observation. 2. Updated beliefs are propagated through population dynamics conditional on U. 3. Monitoring observation Y^+ is used in update to posterior belief state.

Notes: (a) The set of feasible controls excludes $L > 0$ and $U = 0$.

2.4.1. Learning model

The belief state is updated conditional on actions taken and observations received. The multiple measurement controls and the "Bayesian tracking" nature of the model makes the update possibilities more complex than prior related work. Table 1 provides a brief qualitative description of the possible belief update cases in our model. To simplify the notation, we adopt the "+" convention for next-period values and suppress the time subscript.

Each path for beliefs begins with a prior $b(X)$ that determines control choices $\{U, M, L\}$. There are four cases for updating beliefs, summarized above in Table 1. First, when $M = L = 0$, the resource manager predicts the next period population using the population model (Eqs. (1) and (3)) without making use of new information, other than her choice of harvest effort. New information is available to update beliefs when either M or L are positive. When $M = 0$ and $L = 1$, CPUE is measured. CPUE is a noisy observation of X, and is therefore used to update $b(X)$ directly (conditional on U, which determines the density of $q(U)$) before propagating the belief state through the population dynamics to produce the posterior $b^+(X^+)$. When $M > 0$

and $L = 0$, a monitoring observation is used to update beliefs conditional on M (given the dependence of the observation shock $Z^{m+}(M)$ on M).

A two-stage belief update occurs when $M > 0$ and $L = 1$. First, $b(X)$ is updated based on $CPUE$ conditional on U. This updated belief is then propagated through the population dynamics conditional on U. Beliefs are updated a second time given Y^+ conditional on M to produce the posterior belief state.

2.4.2. Approximating Belief Updates

With a lognormal approximate belief state, beliefs regarding the true population level may be summarized by the lognormal location and scale parameters (μ_t, σ_t). Unfortunately, none of the belief update possibilities described above produce a posterior belief state that has a lognormal posterior. This is a well-known feature of continuous-state POMDPs that is not unique to our approximation strategy. To address this we find the parameters of an approximate posterior distribution that match the true posterior most closely in the sense of minimizing the Kullback-Leibler divergence. As shown by ZFM, this so-called projection approach is equivalent to matching the moments of the true and approximate posterior when the belief distribution is in the exponential family, of which the lognormal is a member. For the lognormal distribution, this involves matching the first two moments of the log population.

To implement moment matching we build on an approach to particle filtering coupled with density projection first proposed by ZFM. In what follows, we again adopt the "+" notation. To simplify the discussion we also suppress the dependence of observation shocks $q(U)$ and $Z^m(M)$, instead writing q, Z^m, while also not explicitly noting the dependence of the Y and $CPUE$ conditional densities on control choices. Lastly, although we describe here computations involved in producing the approximately optimal policy function, we perform closely-related calculations when simulating belief state dynamics conditional on a pre-computed policy (see Section 3.3).

The basic idea is as follows. First, we discretize the actions and the parameters of the approximate belief state. For each of these discrete values we generate samples of both the target variables (X or X^+) and, where applicable, measurements ($CPUE$ and Y). For each observation of the measurement variables, we determine the conditional probability of obtaining each sampled value of the corresponding target variable. These probabilities are used to obtain weighted averages of the target samples to estimate moments of the target variables that can be interpreted as expectations conditioned on the measurements. Using these averages, we determine interpolation weights on the discrete values of the approximate belief state parameters, which are averaged to obtain transition probabilities over discrete values of posterior (μ^+, σ^+). These transition probabilities represent columns in a column-stochastic belief state Markov transition matrix (MTM).

We begin by selecting a discrete set of nodal levels for the belief parameters (μ, σ) and actions (L, M, U). As described in further detail in the next section, a theoretically consistent and computationally convenient strategy is to divide the transition computations first between those conditional on $L = 0$ and $L = 1$, and then into 2 stages involving first learning from CPUE and then from monitoring. Doing so produces separate sets of MTMs. One set corresponds to (μ, σ, L, U) (consisting of two MTMs for each node of U, with one conditional on $L = 0$ and another conditional on L=1). The second corresponds to (μ, σ, M) (one MTM for each node of M). This collection of MTMs is then combined to generate an approximation of the full MTM for (μ, σ, L, M, U).

We first consider the 2 stage transition conditional on $L = 1$ (meaning $U > 0$ by necessity) and $M > 0$. Given any specific value of (μ, σ) and positive (M, U) we use the particle filter to obtain two separate MTMs that map (μ, σ, M, U) into probabilities associated with levels of the updated belief parameters (μ^+, σ^+). In the first stage we obtain an MTM mapping (μ, σ, U) into (μ^+, σ^+), which includes the observation update for (μ, σ) (resulting from information in $CPUE$) and the time update that maps the updated (μ, σ) into (μ^+, σ^+) via the stochastic population dynamics. The second stage correspond to the monitoring update, using (μ^+,σ^+, M) to update (μ^+,σ^+) (through receiving Y^+).[1] This produces a second MTM.

It is useful to first describe the second-stage monitoring update in further detail. The approach to generating a monitoring update is similar to that used by Kling et al. (2017). We generate n_d samples of X_i^+ using (μ^+, σ^+) from which associated values of $Y_i^+ = Z_i^{m+} X_i^+$.[2] In other words, the target variable is the beginning of next period population X^+, and the measurement is the monitoring observation Y^+. For each monitoring observation Y_j^+ the conditional probability of each population observation S_i^+ is proportional to[3]

$$w_{ij}^m \propto \frac{f_{Z^m}\left(\frac{Y_j^+}{X_i^+}\right)}{X_i^+} \tag{7}$$

with the particle weights w_{ij}^m are normalized to sum to 1: $\sum_{i=1}^{n_d} w_{ij}^m = 1$.

Notice that the sample values of X_i^+ do double duty. First, they provide a set of samples of the values of Y_i^+. Second, they serve as a set of possible population levels each with a conditional probability weight (w_{ij}) that depends on the level, X_i^+ conditioned on Y_j^+. The weights can

[1] In the case of a 2-stage transition involving an update conditional on both $CPUE$ and Y^+, it would perhaps be more accurate to describe an intermediate posterior $(\tilde{\mu}, \tilde{\sigma})$ produced from the first stage update, which is then mapped to the final posterior (μ^+, σ^+) by the monitoring transition. However, since the same MTMs are used to model belief transitions when one of L or M is zero (e.g., the first-stage transition MTMs for U nodes used to compute expected posterior beliefs when monitoring is not chosen), we avoid introducing an intermediate prior in order to make the description applicable to other belief update cases.
[2] Note here as described above the density of Z_i^{m+} depends on the conditioning value of M.
[3] Dividing by zero does not occur since all sampled X_i are strictly positive.

then be used to obtain samples of values of the mean and standard deviation of the log population levels conditional on observations of Y_j^+:

$$\mu_j^+ \approx \sum_{i=1}^{n_d} w_{ij}^m \log(X_i^+) \tag{8}$$

$$\sigma_j^+ \approx \sqrt{\sum_{i=1}^{n_d} w_{ij}^m \left[\log(X_i^+) - \mu_j^+\right]^2} \tag{9}$$

Samples (μ_j^+, σ_j^+) and are used to compute the MTM that maps the unconditional belief parameters (μ^+, σ^+) into parameters of the conditional belief distribution. Sampled values will not generally fall on nodal points, so we assign probability values using the interpolation approach detailed in the next section.

The belief update for the case where learning from a CPUE measurement is more complicated. To the best of our knowledge, our technique for modeling the belief update using projection particle filtering is a novel. The target variable this time is the beginning of period population X, and the measurement is $CPUE$. Using the prior belief parameters (μ, σ), we first generate n_d samples of the population level X_i and, separately, the catchability shock q_i.[4] This allows us to compute samples $CPUE_i = q_i X_i$ Eq. (2). Next, given U, we obtain a sample of escapement values S_{ij} for each $CPUE_j$. Directly computing $S_{ij} = X_i - CPUE_j U$ by reusing the samples X_i and $CPUE_j$ values can lead to two numerical problems. First, some S_{ij} may be negative, especially when q_i or U are large. Such zero-probability escapement values would need to be removed, leading to a second problem of small useable samples that may fall well below the intended number (n_d).

To produce useful escapement samples, we generate j sets of population samples from the prior belief truncated from below by $CPUE_j$:

$$\log(X_{ij}) = \Phi^{-1}(\bar{P}_j + (1 - \bar{P}_j)\xi_i^P; \mu, \sigma^2) \tag{10}$$

where $\Phi^{-1}(\cdot)$ is the inverse normal CDF, ξ_i^P is a sample from the uniform distribution, and $\bar{P}_j = \Phi(\log(CPUE_j); \mu, \sigma^2)$. With this resampling approach, each X_{ij} satisfies the intuitively reasonable condition that $X_{ij} \geq CPUE_j$ (i.e., there must have been at least sufficient population available to generate the CPUE measurement).

Similar to Eq. (7), the conditional probability of X_{ij} given $CPUE_j$ is:

[4] We suppress here the dependence of $q_i(U)$ on U.

$$w_{ij}^C \propto \frac{f_q\left(\frac{CPUE_j}{X_{ij}}\right)}{X_{ij}} \tag{11}$$

where again we normalize setting $\sum_{i=1}^{n_d} w_{ij}^C = 1$. Next, each X_{ij} is simulated forward through the population dynamics equation using a sample of growth shocks Z_i^g to produce samples X_{ij}^+. With these, we produce samples of (μ_j^+, σ_j^+) by performing calculations with w_{ij}^C and X_{ij}^+ analogous to Eqs. (8)-(9). An MTM is produced looping over the belief state space and performing interpolation over nodal values of the approximate belief state parameters.

A third type of MTM is necessary for control combinations involving $L = 0$. In that case, the first stage transition does not involve a Bayesian update. Instead, population samples produced from nodal values of (μ, σ) are propagated through the population dynamics conditional on U using draws of Z_i^g. Moments of the sampled $X_i^+$are used directly to match the posterior to the lognormal distribution parameters (μ^+, σ^+) without weighting. Note that this assumption means that the manager knows and uses the effort but not the harvest information. As with the other transition stages, interpolation is carried out to construct an MTM.

2.4.3. Numerical Implementation

Once our approximation approach is used to produce Markov transition matrices characterizing belief dynamics, the problem may be solved as a belief state dynamic programming problem using conventional methods. Aside from the more complicated belief dynamics, our numerical procedure is similar to recent analyses of bioeconomic models involving state uncertainty (Kling et al. 2017, Sloggy et al. 2019). We choose nodal values of the belief state variables by selecting a regular grid over the arithmetic mean and CV $(mean_i, CV_i)$ of the population stock X and convert these to obtain a discretization of the mean and standard deviation of (μ, σ) of $\log(X)$.

Like Sloggy et al., we use Halton draws to generate samples and apply linear interpolation in the MTM computations. We also exploit a technique described by Fackler (2018) for carrying out dynamic programming computations by calculating the expected continuation value in the Bellman equation without forming the full problem transition Markov matrix. The full transition matrix has the dimension $n_{mean}n_{CV} \times n_{mean}n_{CV}n_L n_M n_U$ (see Table 2 for definitions). By splitting up the problem we instead generate 2 $n_{mean}n_{CV} \times n_{mean}n_{CV}n_U$ matrices (for with- and without CPUE learning) and an $n_{mean}n_{CV} \times n_{mean}n_{CV}n_M$; this is proportional to $2n_U + n_M$ rather than $2n_U n_M$. This results in significantly faster processing speeds.

Table 2. Details on solution approximation: sizes

Symbol	Definition	Value	Notes
n_d	Samples	2,500	Halton draws are used to generate samples
n_{mean}	Number of belief mean nodes	201	
n_{CV}	Number of belief CV nodes	51	
$(mean_L, mean_U)$	Mean node bounds	(1e-4, 200)	
(CV_L, CV_U)	CV node bounds	(1e-4, 0.5)	
n_L	Learn from CPUE options	2	$L \in \{0,1\}$
n_M	Number of monitoring nodes	5	
n_U	Number of harvest effort nodes	100	

Once formed, we use the modified policy iteration (MPI) to solve the dynamic program. The solution is found relatively quickly. By far the most time-consuming aspect of this method is computing the MTMs. The run time depends on the number of samples and of state and action nodes chosen. For reasonably accurate approximations, it is not unusual for the MTM computation to take one or more days, while MPI takes minutes to converge. Fortunately, the MTM calculations may be sped up to some degree with parallel processing.

2.5. Parameter values

In this section we report parameter values underlying our numerical analysis. We develop a base set of parameter values and distributional assumptions designed to be useful while also allowing the model to match stylized facts about stochasticity in renewable resource management problems. Table 3 reports model parameters held fixed across cases we consider in our analysis.

All cases considered adopt a simplified version of the per-period reward function (Eq. (5)) with only linear variable costs for monitoring and fixed costs for using CPUE data:

$$\widetilde{\Pi}(U_t, M_t, L_t) = pq_t(U_t)X_tU_t - c_M M_t - c_L L_t \tag{12}$$

Parameters c_M and c_L vary across cases. In addition the marginal cost of harvest is constant and is netted out of profit.

Table 3. Base parameter values

Symbol	Definition	Value	Notes
r	Intrinsic growth rate	0.5	
K	Carrying capacity	100	
$\overline{CV}_g$	Growth shock CV	0.15	Z_t^g ~Burr-3 and $E(Z_t^g) = 1$
γ_g	Growth shock skewness	-0.1	
$\bar{q}$	Mean catchability	0.6	$q_t(U_t)$ ~Kummaraswamy
$\overline{CV}_q^U$	Maximum catchability CV	0.425	
η_q	Catchability CV calibration coefficient	2.75	
$\overline{CV}_m^U$	Maximum monitoring noise CV	0.425	$Z_t^m(M_t)$ ~Burr-3 and $E\left(\frac{1}{Z_t^m(M_t)}\right) = 1$
γ_m	Monitoring noise skewness	-0.1	
η_m	Monitoring shock CV calibration coefficient	2.25	
p	Price per-unit harvest	1	
β	Discount factor	0.952	$\beta \approx 1.05^{-1}$

We require specific functional forms for the shocks included in the process and observation models. We start by assuming realizations of $q_t(U_t)$ are described by a Kumaraswamy distribution. This is a parameterized density with support over (0,1) with similar properties to the Beta distribution but with an easily inverted CDF:

$$F(x; a, b) = 1 - (1 - x^a)^b \tag{13}$$

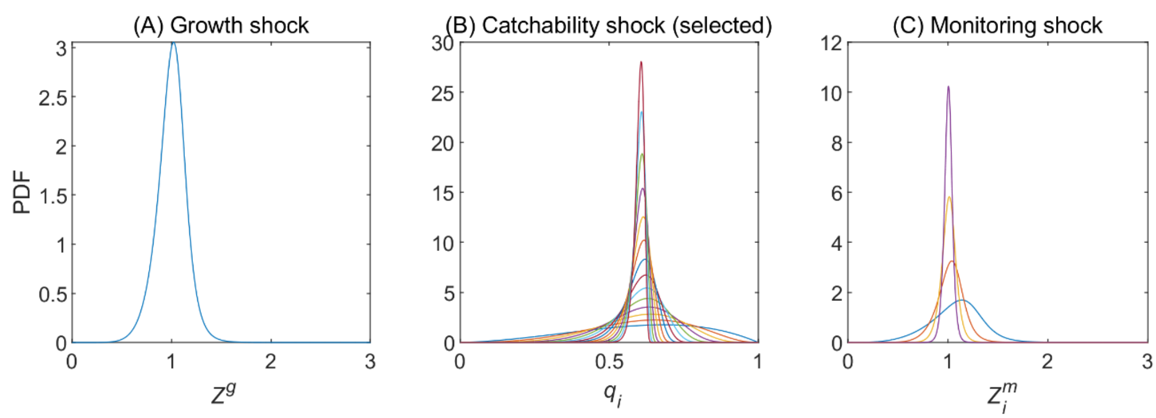

Fig. 1. Densities for process model shocks

For each node of U_i used in our solution method, we numerically solve for Kumaraswamy parameters (a_i, b_i) such the mean of the density is $\bar{q}$ and the CV is determined by the following decreasing function of U_t: $CV[q_t(U_t)|U_t = U_i] = \overline{CV}_q^U \exp(-\eta_q U_i)$. This functional form results in the coefficient of variation of CPUE declining in U within the range $[\overline{CV}_q^U \exp(-\eta_q), \overline{CV}_q^U)$. The Kumaraswamy distribution is ideally suited for numerical analysis since it has closed forms for both the cumulative distribution function (CDF) and its inverse. Selections of the U_i node-specific densities generated from this procedure are shown in Fig.1, Panel B, with the most flat distribution corresponding to $U = 0$ and the most narrow corresponding to $U = 1$.

We assume that both the growth shock Z_t^g and multiplicative monitoring measurement error shocks follow Burr-3 distributions. The Burr-3 CDF is given by:

$$F(x; a, b, c) = \left[1 + \left(\frac{x}{a}\right)^{-c}\right]^{-b} \tag{14}$$

The Burr-3 shares some similarities with well-known distributions like the lognormal and gamma, while having a much simpler and easily invertible CDF involving no special functions. As with harvesting effort, increasing monitoring effort decreases the coefficient of variation of multiplicative measurement error; specifically: $CV[Z_t^m(M_t)|M_t = M_i] = \overline{CV}_m^U \exp(-\eta_m M_i)$. Similar to the catchability shock, this bounds the CV of the monitoring shock on $[\overline{CV}_m^U \exp(-\eta_m), \overline{CV}_m^U)$. We adjust distribution parameter values numerically so that the mean, skewness, and CV of the Burr-3 densities in the model fit assumed values reported in Table X or given by the assumed relationship for conditional CV. Resulting densities corresponding to nodal values of M_i used in our numerical analysis are shown in Fig. 1 Panel C.

3. Results

Our presentation of numerical results begins with a brief review of the model solution in the special case of perfect observability. Next, we review three cases of the solution to our full model, which differ in terms of assumptions about management costs. We close by describing findings from a simulation exercise that compares dynamics associated with the full model with a selection of counterfactuals.

3.1. Stochastic model solution

To anchor the discussion that follows we first describe the solution to a special case of the model when the resource manager does not face state uncertainty. Assuming that the beginning-of-period population is always observed perfectly results in a standard stochastic bioeconomic optimization problem with one state (X_t) and one control (U_t). The optimal policy and value functions are shown in Fig. 2.

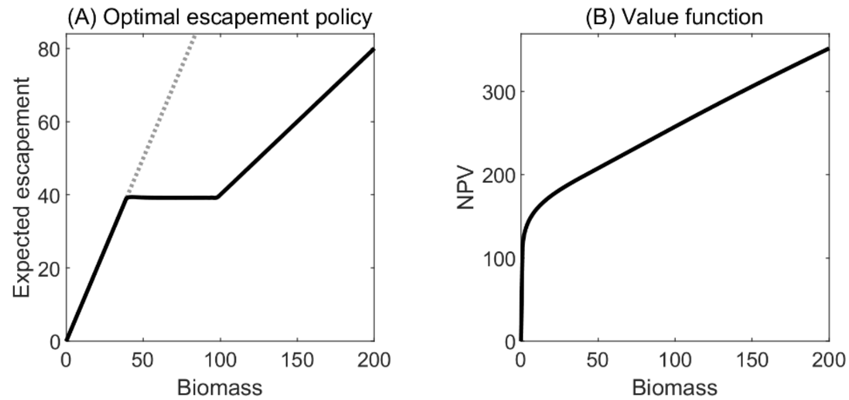

Fig. 2. Stochastic model solution

A standard way to show the optimal feedback policy is in the form of expected escapement: $E[S_t|X_t, U_t] = (1 - \bar{q}U_t)X_t$. The pattern up until roughly the value of $X_t = K = 100$ is a standard constant escapement (Fig. 2, Panel A).[5] For larger values of X_t, this pattern is broken and escapement is allowed to rise, albeit at a slower rate than the dashed 45° line shown in the panel. The reasons is that we adopt what we regard as a realistic assumption that it is not feasible to exhaust the stock in expectation. The maximum expected harvest yield is $\bar{q}X_t$. We assume $\bar{q} < 1$. While humans have extirpated many wild populations, our view is that it is unrealistic to allow this to occur in expectation within one decision making period. The model still allows for asymptotic extirpation of the population, however.

Table 4. Full model solution cases

Case	Parameter Values
Base	$c_L = 0.01,\ \ c_M = 0.4$
High c_L	$c_L = 1,\ \ \ \ \ \ c_M = 0.0001$
High c_M	$c_L = 0.001,\ \ \ \ \ c_M = 1.25$

3.2. Full model solution

A solution to the full model is an optimal value function and feedback policy defined over approximate belief state parameters, or equivalently in our application the lognormal statistics $(mean_t, CV_t)$. The feedback policy maps the belief state to an optimal action comprised of the three controls: whether or not to learn from CPUE (L_t), monitoring (M_t), and harvest effort (U_t). We explore the behavior of the full model solution by recomputing it for different cases

[5] To fully characterize the constant escapement nature of the stochastic model solution, we use interpolation during value iteration for this version of the model in order to achieve a finer approximation with respect to the control variable. While this is practical for a simple problem like this special case, we have not yet attempted to apply this method to the much larger full POMDP model.

of the objective function parameters (Eq. (12)).[6] Table 4 above summarizes these cases, which we consider in turn.

3.2.1. Base case

Our strategy for developing a base case was to find values of the cost of measuring and learning from a CPUE observation (c_L) and the marginal cost of monitoring effort (c_M) such that both controls are applied with appreciable frequency in the ergodic set of belief dynamics. The base case parameters (Table 4) produces a solution with complex behavior with respect to the control choices in the belief state space (Fig. 3 below). Effort for the most part takes on a familiar profile, staying at zero until the expected population biomass reaches a threshold value (Panel A & B). This is overall consistent with constant escapement.

Two features of the optimal effort policy depart from the standard constant escapement shape. First, effort is slightly elevated as confidence falls (CV of the belief state rises).[7] This seemingly anti-precautionary feature of the policy should be interpreted while noting that the "medium confidence" iso-CV value is relatively rare and the low confidence value is outside of the ergodic set for this parameterization of the model. Still, the fact that effort is rising slightly rather than falling is perhaps surprising. We attribute this to two effects. First, greater harvest can reduce harvest in the next period by narrowing the range of likely populations following growth. A second feature of our model is that some level of harvest is necessary in order to learn from a CPUE measurement. The connection between effort and learning from CPUE creates a complex but intuitive pattern in this solution (compare Panels A and D). A narrow band of low effort choices emanates out into low-to-moderate expected stock sizes and moderate to high CV values (Panel A). This appears to be driven by the simultaneous choice to learn from CPUE over those values (Panel D).

[6] A noteworthy feature of this solution approach is that solving the model with different objective function parameters is *much* faster than re-solving after changing parameters or other assumptions in the process or observation models (Table 3), or discretization choices (Table 2). This is due to the fact that the same MTMs may be used for different cases of the reward function, while one or more MTMs must be re-computed when other parameters are changed. Re-solving the dynamic programming problem when the same MTMs may be used is relatively fast (on the order of minutes).

[7] In Figs. 3-5, the iso-CV values chosen are also node values in the discretization and are approximately: 0.0701 ("high confidence"), 0.2401, and 0.41 (low confidence).

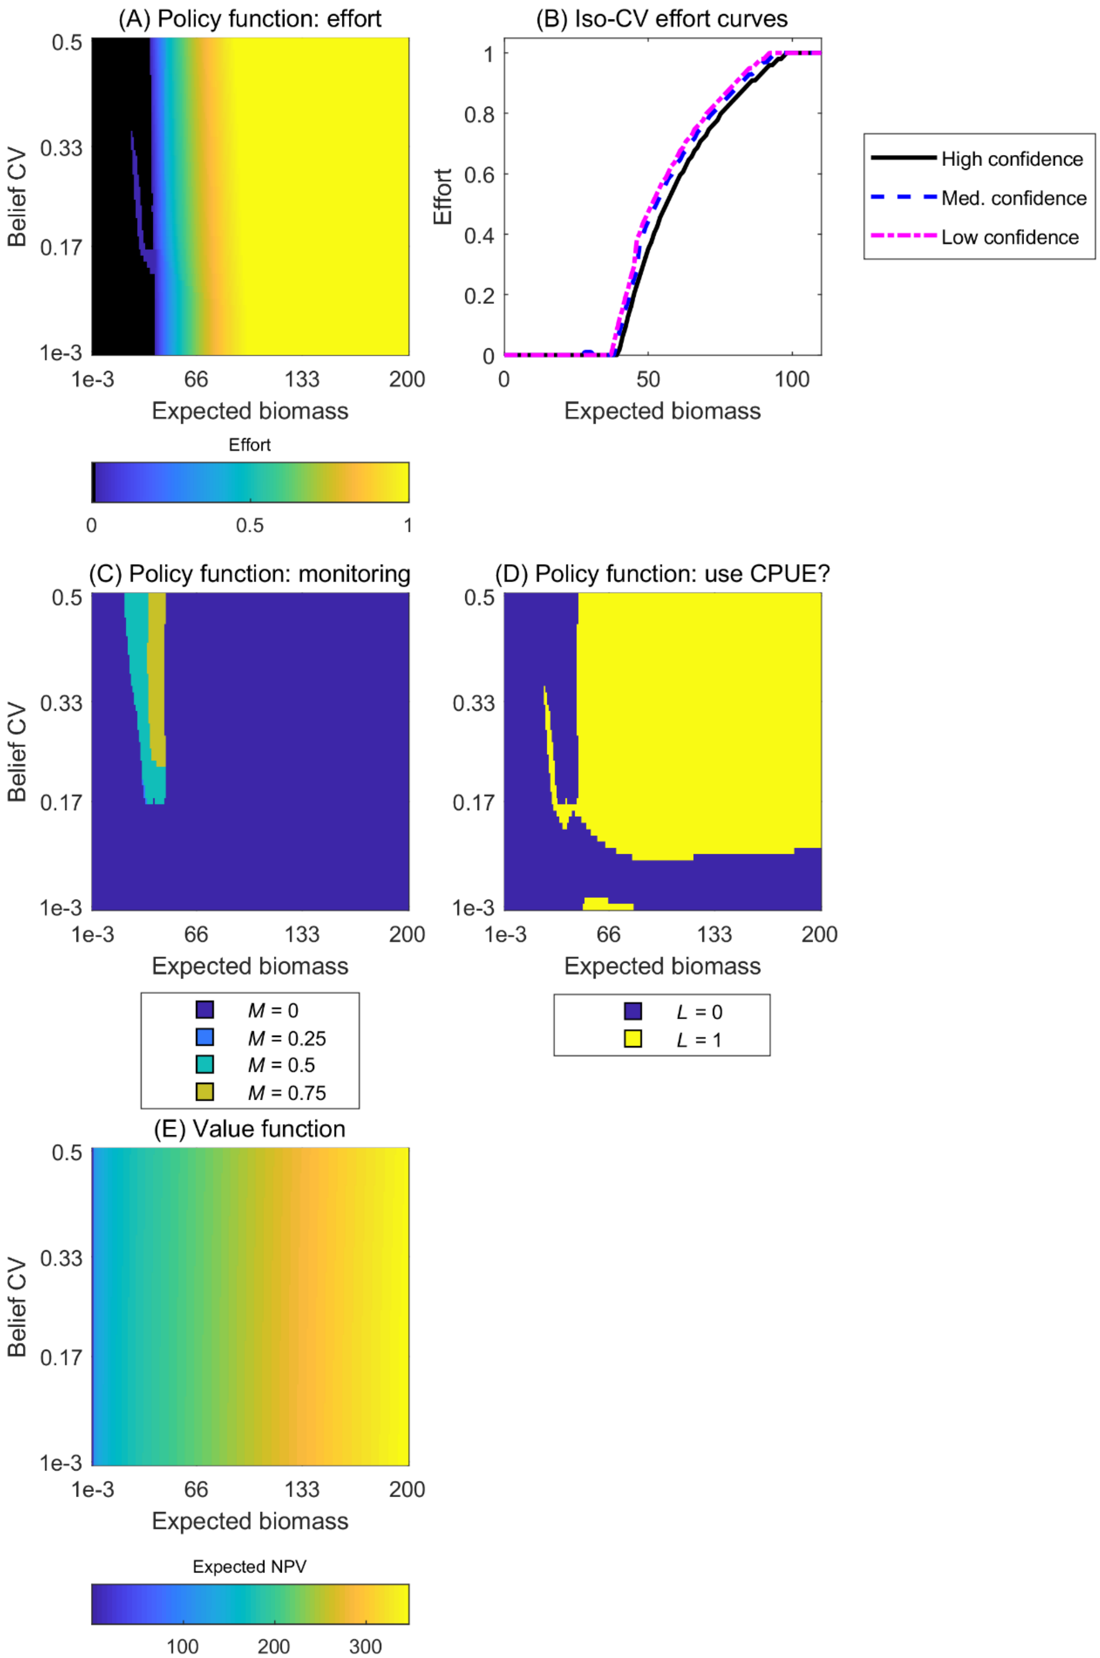

Fig. 3: Model solution: base case

Monitoring effort is optimal over a range of the state space near the positive effort threshold (Panel C). At moderate to high belief CV levels (lower confidence), when it is not optimal to harvest it is optimal to monitor to determine if the stock is sufficiently recovered to allow harvesting to be resumed. When harvesting is undertaken, it is useful to use the CPUE data so long as current uncertainty levels are high enough that their use will reduce uncertainty in a cost effective manner. This result is intuitive: a more accurate estimate of the current population is useful when it may change the decision of whether to harvest the population or not. Interestingly, for some low CV levels a zone where harvesting and learning from CPUE bounds the monitoring zone from below. This produces a pattern where, fixing the mean, it can be optimal to collect no measurements, then at higher CV values switch to learning from CPUE and then to monitoring effort. In this solution L and M are substitutes: no nodal values of the belief state involve investing in both learning from CPUE and monitoring in the same period, although the model does allow for that possibility.

The optimal value function follows a pattern similar to the model with no state uncertainty, rising quickly from low expected biomass levels (Panel E). The cost of initial uncertainty is possible to characterize using the value function. Under this set of parameter values, a higher initial belief CV reduces the expected value of management, but the effect is small: taking the mean % difference of the value function along the low confidence iso-CV transect used in Panel B, relative to the value function over high confidence iso-CV transect yields an average of 0.6% loss. The same exercise using the boundary values (CV_L, CV_U) (Table 2) produces an average of 0.9%. We attribute the small differences to the transient nature of high belief CV values. Beliefs are rapidly pulled from initial conditions to a neighborhood of long-run beliefs in this model. Calculations are also sensitive to parameter values, and we provide additional results on dynamics realizations of management in the next section. Due to the subtle visual changes in the value function across cases, we do not produce separate panels for the value function for the remaining cases.

3.2.2. High cost of learning from CPUE, low cost of monitoring

Next, we consider a case when the cost of monitoring is negligible, and the cost of learning from CPUE is high enough so that the solution does not apply that control anywhere in the state space (Fig. 4 Panel D). In this sense, this brings the full model closest to previous bioeconomic POMDP models that involve free monitoring in every period. It turns out that it is enough to set $c_L = 1$ to produce this result. Compared to the average magnitude of the value function (~253), this apparently small cost suggests that learning from CPUE is relatively uninformative about the future post-growth population level. Results from our dynamic analysis presented in the next section confirm this property of the model.

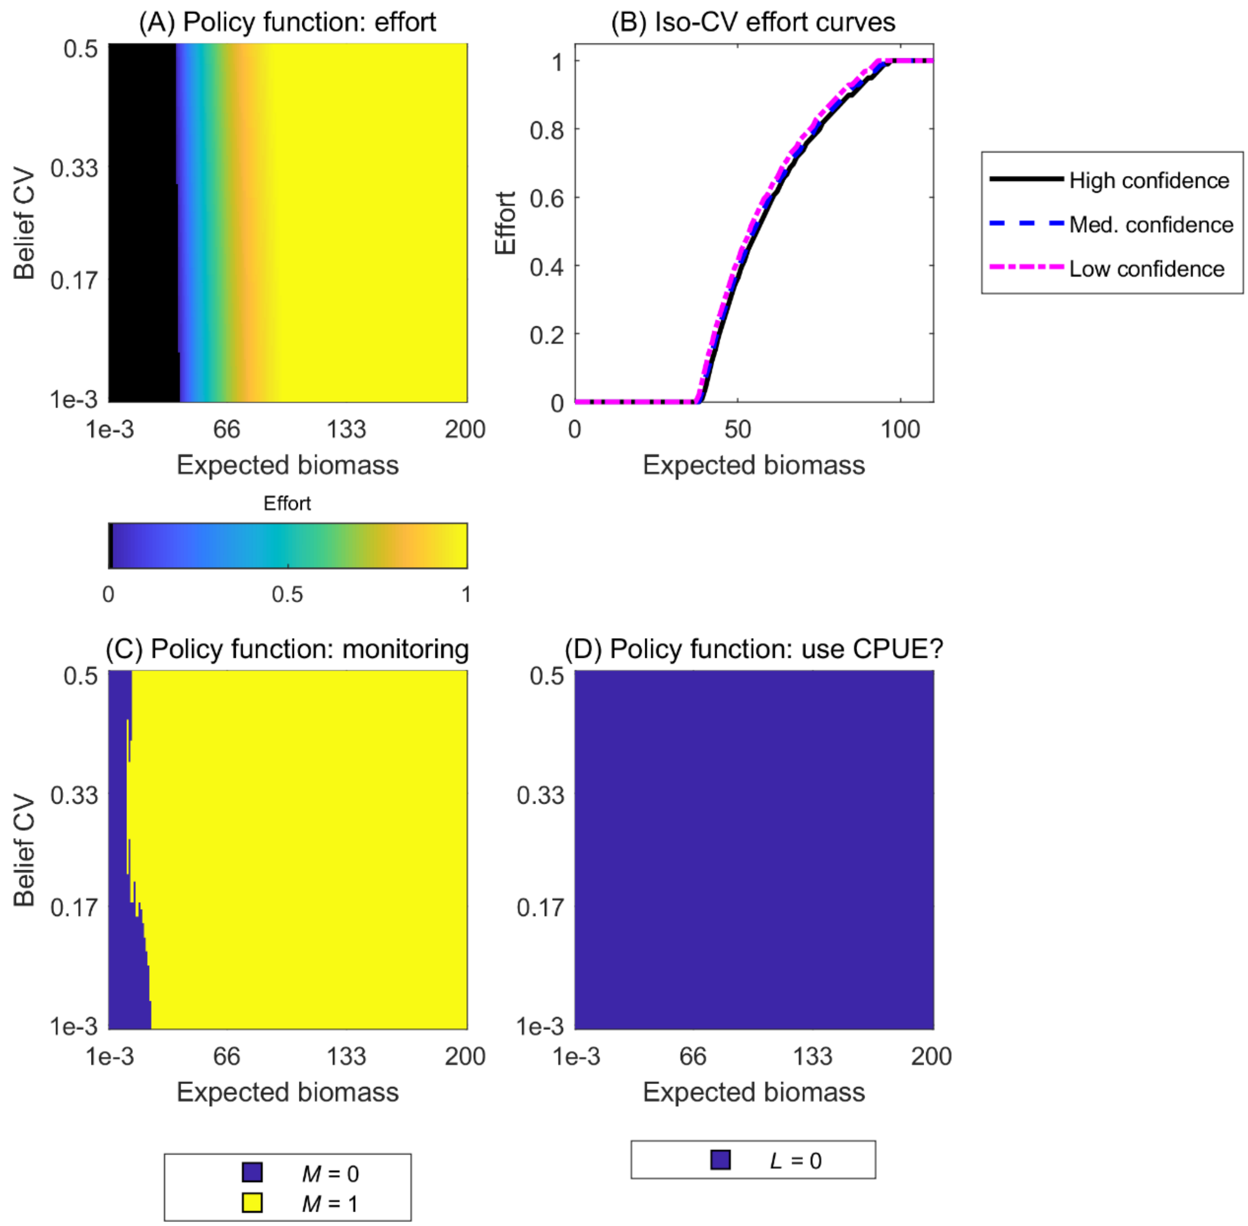

Fig. 4: Model solution: high cost of learning from CPUE, low cost of monitoring (high c_L case)

Effort tends to increase slightly as a function of the belief CV, however this effect is more muted than the base case (Panel B). This case also involves negligible monitoring costs, so monitoring is conducted nearly everywhere. Interestingly, for low enough expected biomass levels monitoring is not conducted despite its low cost. There are a few possible explanations for this result. First, our solution perhaps involves too few samples (n_d) to differentiate the expected value of monitoring vs. not monitoring in this area of the belief state. Second, the zone where monitoring is not chosen overlaps with the zero effort zone of beliefs (compare Panels A and C). The strong tendency for beliefs to rapidly revert to the ergodic set, coupled with the distance in belief space to the harvest threshold likely makes even a small cost of monitoring tip the balance toward inaction.

3.2.3. Low cost of learning from CPUE, high cost of monitoring

When the marginal cost of monitoring effort is relatively expensive c_M = 1.25, it is not chosen anywhere in the belief state space (Fig. 5 Panel C). Substitution to learning from CPUE is evident in the pattern of effort (compare Panel A and D). One consequence is a more pronounced but still modest anti-precautionary effect, with sufficiently high belief CV values triggering harvest (Panel B).

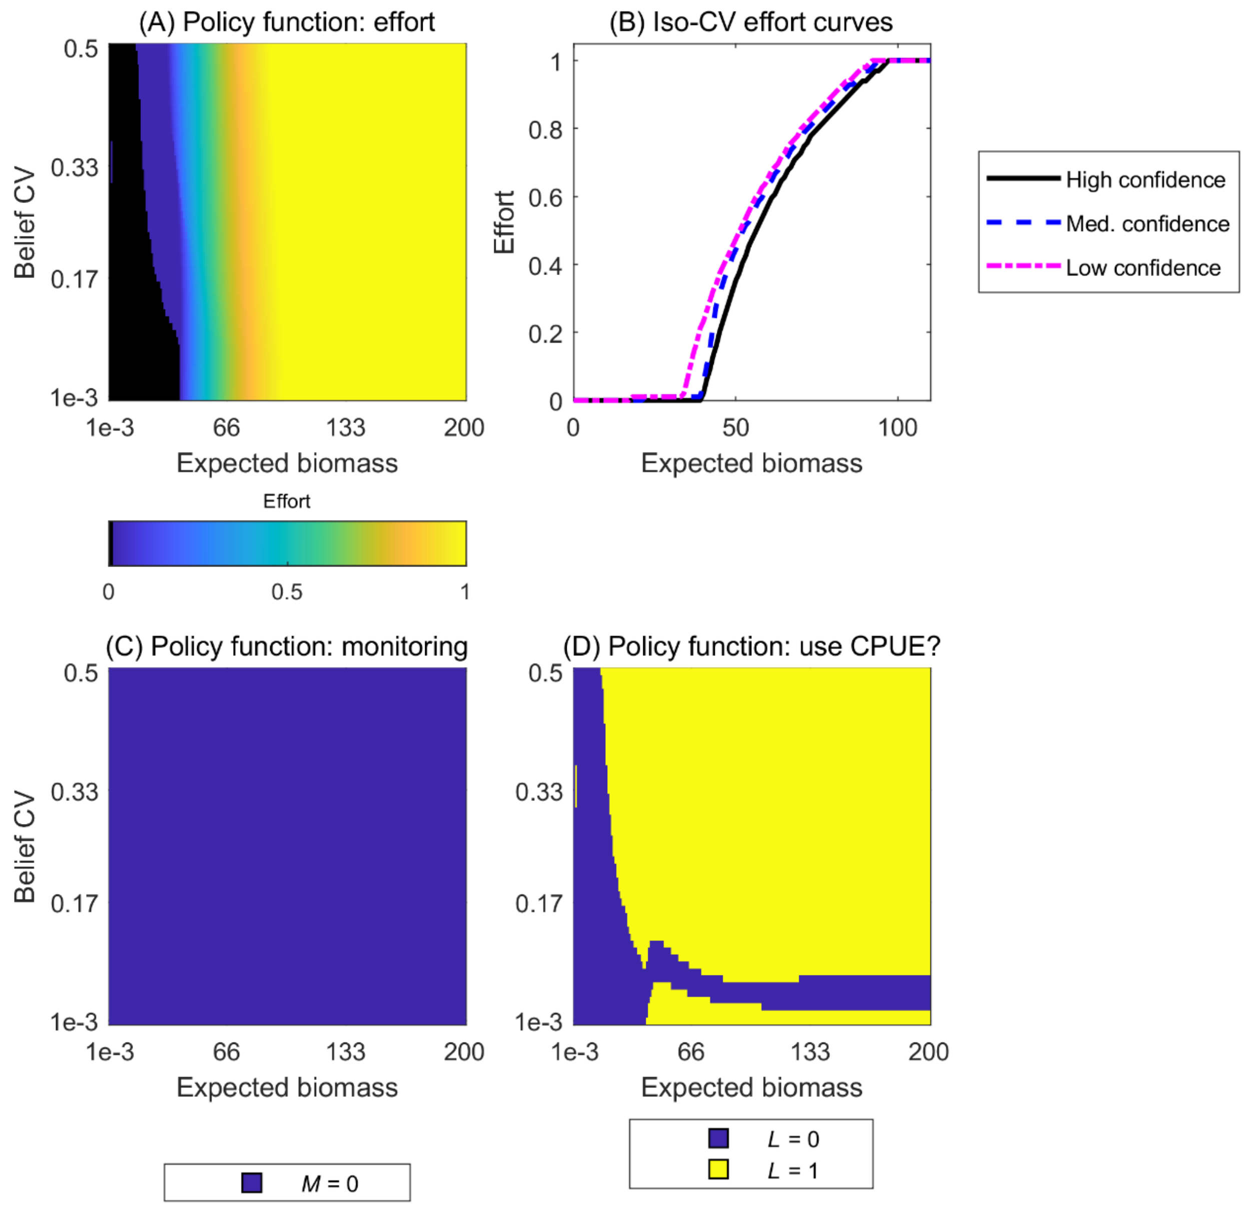

Fig. 5: Model solution: high cost of monitoring (high c_M case)

The pattern of the L control over the state space is complex (Panel D). We suspect numerical precision may be the cause of the blue channel of inaction through low CV values.[8]

3.3. Dynamics

We conduct a series of simulations to supplement insights from the case-specific value and policy functions reviewed in previous section. To do so, we generate dynamic realizations of the true underlying stock, corresponding beliefs, and resulting action choices. Simulations are produced over T periods for each case of the full model presented above, as well as additional counterfactuals for comparison. Importantly, we produce true matched counterfactual dynamic runs. For each time series k, realizations of $\{\xi_{tk}^q, Z_{tk}^g, \xi_{tk}^m\}_{t=1}^T$(Eqs. (1), (3)-(4)) are pre-computed and applied to each policy, and each time series is initialized from the same initial belief and true stock level. As a result, differences among policy-specific realizations indexed by k are due solely to differences in controls selected or available to the respective policies. Finally, like other dynamic models involving state uncertainty we require two initial conditions in order to simulate dynamics of the controlled process: one for the belief state and one for the true underlying population. For all simulations reported in this section we use $X_0 = 100$ for the true population initial condition, meaning the population is in the range of unexploited population levels near K, and $(mean_0, CV_0) = (100, 0.14)$, which suggests the resource manager has an accurate initial estimate characterized by moderate confidence.

Table 5. Summary of supplemental policies

Policy	Description
Sethi POMDP	Equivalent to the full model with monitoring effort always chosen at it's second-highest level without any cost incurred.
Perfect observability	State uncertainty is eliminated from the model and the conventional stochastic dynamic programming solution is applied to the true population level (Fig. 2).
Ignore state uncertainty	Measurements are generated from Eq. (4) assuming the lowest level of positive monitoring effort ($M = 0.25$). The observations are taken as the true population biomass and effort is chosen according to the perfect observability solution.

The additional policies we consider are summarized in Table 5. The perfect observability case provides an expected upper bound on the net present value of management, all else being equal. The "Sethi POMDP" alternative is perhaps more correctly attributed to Memarzadeh and Boettiger in that it is a POMDP solution to an earlier optimal harvesting problem involving state

[8] For this case the difference in the value function between $L = 0$ and $L = 1$ is very small, and even with near indifference (less than 0.01% gap), the dynamic programming algorithm will default to $L = 0$. We plan to investigate this issue with additional work using a more refined grid over belief state values and more samples. Unfortunately, producing these results with a 42 core server can entail a run time of several days.

uncertainty developed by Sethi et al. (2005).[9] It involves no learning from CPUE and free monitoring at a fixed monitoring effort level which we specify as being the second-most informative monitoring control choice in the full model. As a plausible lower bound on performance from managing while acting on monitoring observations, we introduce a counterfactual where the optimal control for the model with no population uncertainty is applied to the most recent (and informative but relatively inaccurate in this specification) measurement.

Table 6. Selected features of dynamics

Feature[a]	Model					
	Base case	High c_L	High c_M	Sethi POMDP	Ignore state uncertainty	Perfect observability
	Mean Value [Standard Deviation]					
Average belief state mean	47.87 [0.48]	48.69 [0.63]	47.43 [0.4]	48.6 [0.61]	-	-
Average belief state CV	0.174 [0.0013]	0.044 [0.0004]	0.212 [0.003]	0.072 [0.0007]	-	-
Average true population biomass	47.92 [1.22]	48.69 [0.682]	47.43 [1.44]	48.6 [0.747]	45.09 [0.938]	48.93 [0.642]
True population biomass CV	0.227 [0.016]	0.178 [0.011]	0.245 [0.019]	0.184 [0.011]	0.241 [0.013]	0.175 [0.011]
Average number of periods with zero harvest	4.67 [2.64]	22.87 [5.24]	0.0004 [0.02]	18.6 [5.03]	60.72 [5.957]	22.61 [5.22]
Per-period average realized harvest	9.94 [0.535]	10.14 [0.518]	9.84 [0.547]	10.12 [0.52]	9.75 [0.49]	10.16 [0.518]
Average effort level	0.322 [0.012]	0.316 [0.014]	0.325 [0.01]	0.316 [0.014]	0.313 [0.013]	0.314 [0.015]
Average number of periods with learning from CPUE	104.17 [5.17]	0 [0]	199.99 [0.02]	-	-	-
Average number of periods with monitoring	53.78 [2.95]	199.74 [0.534]	0 [0]	-	-	-
Average number of periods with learning from both CPUE and monitoring	0 [0]	0 [0]	0 [0]	-	-	-

Notes: (a) Results summarize 5,000 time series of policy-specific dynamics over 200 years. All dynamic runs begin a true initial stock of $X_0 = 100$ and an initial belief state with $(mean_0, CV_0) = (100, 0.14)$.

[9] We keep the Sethi label to acknowledge the structural similarities between our bioeconomic model and theirs.

Table 6 provides a selection of diagnostics calculated from the dynamics. Differences among the counterfactuals are often subtle, but some intuitive results emerge. The base case involves average confidence falling between the case where maximum monitoring effort is nearly always applied, and the case where the resource manager always learns from a CPUE measurement (Table 6, row 2). Ignoring state uncertainty and acting on inaccurate measurements leads to a relatively large number of periods where harvest is forgone, while this essentially never occurs when monitoring costs are high and the resource manager learns from CPUE in nearly every period (row 5). Unsurprisingly, when monitoring is very low cost and highly accurate observations are received each period, the dynamics closely parallel those of the perfect observability case (compare column 2 and column 6).

Table 7. Selected performance measures

Feature	Model[a]					
	Base case	High c_L	High c_M	Sethi POMDP	Ignore state uncertainty	Perfect observability
	Mean Value [Standard Deviation]					
Realized NPV	252.75 [22.91]	257.32 [22.92]	252.51 [22.92]	256.96 [22.9]	252.62 [21.59]	257.56 [23.01]
Realized PV of harvest rent	253.93 [22.83]	257.34 [22.92]	252.54 [22.92]	-	-	-
Realized PV of non-harvest management costs	1.23 [0.133]	0.21 [0.0001]	0.021 [2e-10]	-	-	-
Realized harvest CV	0.725 [0.044]	0.769 [0.046]	0.724 [0.041]	0.769 [0.045]	1.1 [0.05]	0.776 [0.047]
Realized NPV: Base case rewards	-	-	-	250.72 [22.89]	246.37 [21.59]	-
Realized NPV: high c_M rewards	-	-	-	237.43 [22.89]	246.15 [21.59]	-

Notes: (a) Results summarize 5,000 time series of policy-specific dynamics over 200 years. All dynamic runs begin a true initial stock of $X_0 = 100$ and an initial belief state with $(mean_0, CV_0) = (100, 0.14)$.

Performance measures for the same set of dynamic runs are reported in Table 7. Differences are generally quite small, however with a known data generating process and the ability to generate arbitrarily many time series, signing differences among policy counterfactuals is possible (conditional on parameter values and discretization scheme).[10] One surprising finding is the closeness of average NPV among all policies relative to the perfect observability policy. We conjecture that, compared to a naïve alternative like our version of ignoring state uncertainty, Bayesian solutions will do particularly well when "process error" affecting stock dynamics via catchability and the growth shock is low and measurement error is high. To investigate this issue, we plan to pursue sensitivity analyses over other bioeconomic model parameters outside of the objective function in forthcoming updates to this research.

One interesting use of these results is to reevaluate the Sethi POMDP solution by recomputing realized NPV of management assuming that monitoring is not free. Recall that the Sethi POMDP policy is computed assuming a fixed level of monitoring is carried out in each period. Since the solution involves constant monitoring, any costs of monitoring are excluded from the associated dynamic optimization problem. When these costs are instead accounted for at values consistent with their levels in either the base or high c_M scenarios, the gap between the Sethi POMPD counterfactual and the performance of our model increases in both cases (rows 5-6). This result quantifies the gain in performance from optimally investing in monitoring, and planning harvest effort jointly with the monitoring strategy.

The same recalculation of NPV may be performed for the policy that ignores state uncertainty. The surprising closeness of the NPV between that alternative and the high c_M case of the full model is in this light understood as being a function of the management costs: they are computed for the full model but left out of the ignore state uncertainty counterfactual in keeping with the original assumption of myopia for the latter policy. When monitoring costs are accounted for, a small gap emerges. However, the surprisingly competitive performance associated with acting on the most recent measurement (albeit less predictable in terms of realized CV of harvest) points to the disadvantage of having to rely exclusively on CPUE measurements, which occurs in the high c_M case. CPUE is much less informative about the next period population, and also requires harvesting when it would not be optimal given a high-quality monitoring observation.

Finally, Figure 6 presents some long run results given 2 alternative policies. In the top panels the beliefs about the population level are displayed when no harvesting occurs.[11] In this case mean beliefs (shown in the left panel, are located near the carrying capacity K which is the long-run mean population level. The uncertainty around this level, shown in the right panel, reflects the long-run uncertainty due to environmental noise. The bottom panels display the long run distribution of the mean and CV of the belief distribution with optimal actions. In this case the mean is distributed in a fairly narrow band around 50. The CV on the other hand displays a curious bi-modality. Closer examination of the pattern of results suggests that

[10] Specifically, standard deviations reported in Tables 6 and 7 are across-realization standard deviations of the corresponding calculation from a full known data generating process, and not standard errors of statistical estimates.

[11] Due to time constraints these results are from a different analysis than other results presented and hence may not be completely consistent with other presented results – this will be corrected in the draft that is posted online for the meeting.

monitoring results where the CV is at its upper level of about 0.12 whereas the CPUE (or no information) is used when the CV is at its lower level of about 0.04. In the former case the CV is reduced in the next period and the latter case it is increased. Thus, for these parameters, a different policy tends to be selected in alternate years.

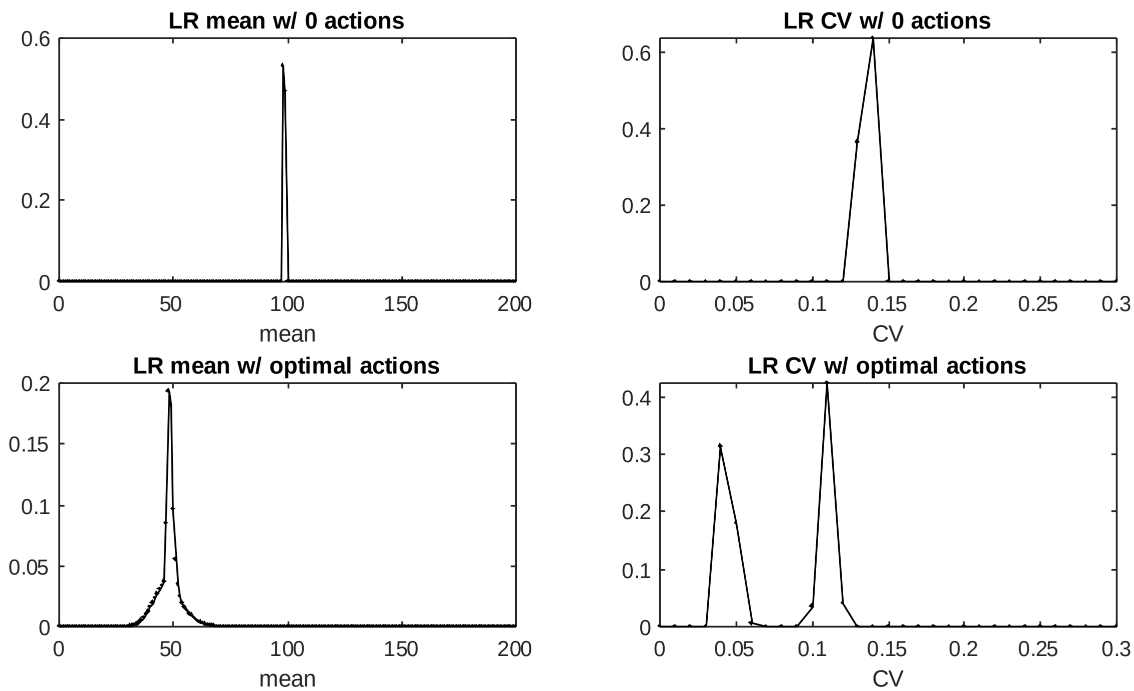

Fig 6. Long-run distributions of the mean and CV of beliefs about the population. The top panels display results with no harvesting occurs. The bottom panel arises when optimal harvesting effort and monitoring occurs.

4. Conclusion

This analysis aims to put the problem of investing in measuring a managed population on equal terms with interventions like harvesting within a bioeconomic optimization model. When population measurement may be done using a combination of costly controls, modest but conceptually important differences with previous solution concepts based on free monitoring emerge. When monitoring the population is not free, it is not in general optimal to do every period. This suggests bioeconomic models that assume population monitoring is done each period entail population measurement when it is not cost-effective, and other actions like harvesting may be wrongly specified to depend on over-investment in one form of learning. These same models likewise typically consider only one channel of information about a population. We show that it is possible to optimally invest in multiple forms of population measurement. In our application, learning from CPUE and monitoring effort appear to be substitutes, with the pattern of substitution strongly influenced by cost differences. This finding needs further investigation with a more refined solution to our full model, and sensitivity analysis over the process and observation shock densities involved in our bioeconomic model.

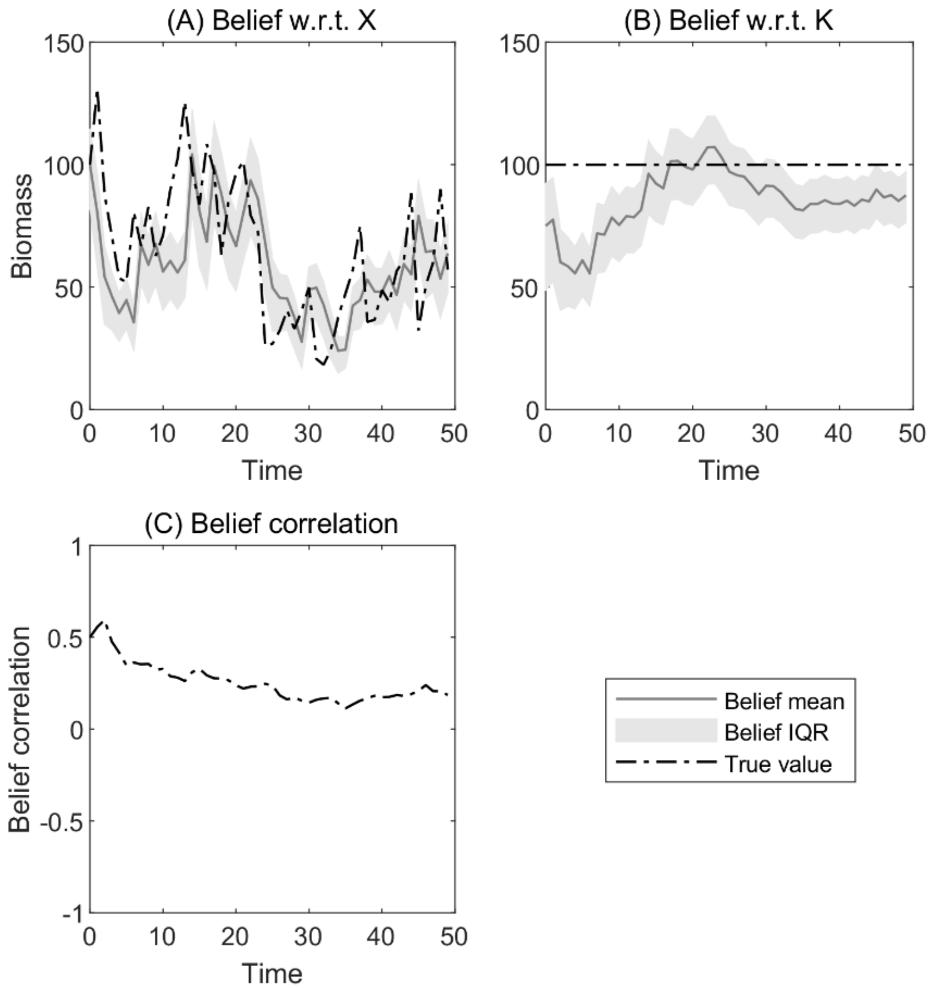

Fig. 7: Belief dynamics under state and parameter uncertainty

Results from this analysis are preliminary, but they so far confirm our sense that the problem of optimally investing in information gathering in bioeconomic models is best characterized as one involving state and parameter uncertainty simultaneously (Fackler and Pacifici 2014). Similar to related prior work on state uncertainty, we assume that parameters of the process and observation models are known with certainty, which likely removes much of the potential value of learning from the model.

Fig. 7 above provides a sample simulation of belief dynamics from a generalization of the model developed in this paper. In that model, the process and observation components are the same as in this paper, however the long-run carrying capacity parameter K is no longer known by the resource manager with certainty. Instead, the belief state is developed over both the population and the parameter. We approximate the belief state in this example with a bivariate lognormal density, and compute belief updates using a procedure closely related to the one described for this model in Section 2.4.

Our conjecture is that investment in learning in this more realistic model will differ substantially from results of this analysis. In particular, more substantial differences between this model and the alternative solution concepts considered in Section 3 seem likely to emerge. We are in the process of working toward an approximately-optimal solution for this model, however due to the greater dimensionality (5 vs 2 belief state parameters that are continuous), the problem requires approximate dynamic programming rather than conventional modified policy iteration that we are able to apply to this model. Recent advances in ADP methods suitable for bioeconomic analysis make us optimistic that a practical solution can be found (Springborn and Faig 2019).

References

Beverton, R.J.H. & S.J. Holt. 1957. On the dynamics of Exploited Fish Populations. Fish. Invest. London. Ser. II. 19: 1–533.

Chase, Michael J., Scott Schlossberg, Curtice R. Griffin, Philippe JC Bouché, Sintayehu W. Djene, Paul W. Elkan, Sam Ferreira et al. 2016. "Continent-wide survey reveals massive decline in African savannah elephants." *PeerJ* 4: e2354. https://peerj.com/articles/2354/

D'evelyn, Sean T., Nori Tarui, Kimberly Burnett, and James A. Roumasset. 2008. "Learning-by-catching: uncertain invasive-species populations and the value of information." *Journal of Environmental Management* 89(4): 284-292.

Fackler, Paul L., and Robert G. Haight. 2014. "Monitoring as a partially observable decision problem." *Resource and Energy Economics* 37: 226-241.

Fackler, Paul L., and Krishna Pacifici. 2014. "Addressing structural and observational uncertainty in resource management." *Journal of Environmental management* 133: 27-36.

Gelman, Andrew, Hal S. Stern, John B. Carlin, David B. Dunson, Aki Vehtari, and Donald B. Rubin. 2013. *Bayesian Data Analysis*. 3rd ed. New York, NY: CRC Press.

Gordon, H.S. 1953. An economic approach to the optimum utilization of fishery resources. *J. Fish. Res. Bd. Can.* 10: 442–447.

Gordon, H.S. 1954. The economics of a common property resource: the fishery. J. Polit. Econ. 62: 124–142.

Kling, David M., Sanchirico, James N., & Fackler, Paul L. 2017. Optimal monitoring and control under state uncertainty: Application to lionfish management. Journal of Environmental Economics and Management: 223-245.

Kotani, Koji, Makoto Kakinaka, and Hiroyuki Matsuda. 2011. "Optimal invasive species management under multiple uncertainties." *Mathematical biosciences* 233(1): 32-46.

LaRiviere, Jacob, David M. Kling, James N. Sanchirico, Charles Sims, and Michael Springborn. 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.

MacLachlan, Matthew J., Michael R. Springborn, and Paul L. Fackler. 2016. "Learning about a moving target in resource management: optimal Bayesian disease control." *American Journal of Agricultural Economics* 99(1): 140-162.

Memarzadeh, M., Britten, G. L., Worm, B., & Boettiger, C. 2019. Rebuilding global fisheries under uncertainty. *Proceedings of the National Academy of Sciences*, 116: 15985-15990.

Nichols, James D., and Byron K. Williams. 2006. "Monitoring for conservation." *Trends in Ecology & Evolution* 21(12): 668-673.

Ralston, Stephen, André E. Punt, Owen S. Hamel, John D. DeVore, and Ramon J. Conser. 2011. "A meta-analytic approach to quantifying scientific uncertainty in stock assessments." *Fishery Bulletin* 109(2): 217-231.

Sethi, Gautam, Christopher Costello, Anthony Fisher, Michael Hanemann, and Larry Karp. 2005. "Fishery management under multiple uncertainty." *Journal of environmental economics and management* 50(2): 300-318.

Sloggy, Matthew R., David M. Kling, and Andrew J. Plantinga. 2019. Measure twice, cut once: optimal inventory and harvest under volume uncertainty and stochastic price dynamics. Working paper.

Springborn, Michael R., and Amanda Faig. 2019. "Moving forward: a simulation-based approach for solving dynamic resource management problems." *Marine Resource Economics* 34(3): 199-224.