

# Learning in a noisy environment: Adaptive management for inconvenient models

## Abstract

In this paper we extend the deterministic mangrove fishery, ecosystem service model of Sanchirico and Springborn (*In press*) to incorporate both risk (irreducible uncertainty) and uncertainty that is reduced over time through learning. We demonstrate how to handle learning models that are more realistic but “inconvenient” in the sense that the function describing a decision-maker’s updated beliefs about the nature of the system does not follow a convenient closed (conjugate) form. To facilitate this we develop a method for approximating the decision-maker’s posterior beliefs using a Kullback–Leibler divergence approach. The full management model describes optimal management of both a harvestable resource (fish) and the ecosystem on which it may depend (mangroves) as a function of the state of the system, which includes both fish and mangrove stocks as well as current information or beliefs about mechanics. Since describing a non-trivial information state space will typically involve at least two state variables, overall we solve a four-state, two-control dynamic programming problem. Since the rate of learning over time is influenced endogenously by both control variables our setting is one of active adaptive management where current controls are selected to maximize returns which include the expected value of information.

Keywords: dynamic programming, Bayesian learning, adaptive management, bioeconomic

JEL:

*Author 1:* Michael Springborn\*, Assistant Professor, Department of Environmental Science and Policy, University of California, Davis, [mspringborn@ucdavis.edu](mailto:mspringborn@ucdavis.edu)

*Author 2:* James N. Sanchirico, Professor, Department of Environmental Science and Policy, University of California, Davis and Nonresident fellow, Resources for the Future, Washington DC, [jsanchirico@ucdavis.edu](mailto:jsanchirico@ucdavis.edu)

**\* Contact author: Michael Springborn, [mspringborn@ucsdavis.edu](mailto:mspringborn@ucsdavis.edu), 530-752-5244.**

# Learning in a noisy environment: Adaptive management for inconvenient models

## Introduction

Understanding the dynamics of ecosystem service provision and designing management measures to maintain the flow of these services over time depends critically on the underlying ecosystem production function. These production functions are models that map non-linear biological and physical processes into the provision of services (Barbier et al. 2008). For the most part, researchers who have investigated the relationship between functions and services have assumed known associations (see, e.g., reviews by Barbier (2007) and Heal et al (2007)), even though they are not sufficiently understood (Daily and Matson 2008).

Recent work in a sea grass, mangrove, coral-reef ecosystem, however, highlights how various assumptions on the nature of these production functions can yield different policy prescriptions (Sanchirico and Mumby 2009; Sanchirico and Springborn *In press*). Misspecification of the ecological production function, therefore, has potentially important implications for setting fish catches, restoring or clearing habitat, and controlling pollution.

The traditional method of attempting to resolve the uncertainty on ecosystem dynamics is to invest in natural science research (Murkowski et al. 2010). Because management decisions are made on a repeated basis, however, there is also the possibility of using these decisions to learn about the ecosystem. Ideally, learning would be driven not only by the exogenous arrival of information (e.g., research outputs), but by an endogenous investment process where the optimal course of action depends in part on the value of information expected to be generated (see., e.g. Hartmann et al 2007, Springborn et al. 2010). All else equal, the faster a manager resolves uncertainty, the sooner the manager is able to identify the ideal policy under that expanded information set, which in this case is over the reversibility of system. This accelerated learning comes at an opportunity cost, which is forgoing whatever action would be ideal in the short run, ignoring learning. An adaptive (endogenous) learning approach seeks to balance these tradeoffs, pursuing informative actions, but only when they are worth the cost.

In this paper, we extend the economic-ecological coral-reef ecosystem model of Sanchirico and Springborn (*In press*) by incorporating both reducible and irreducible uncertainty. The reducible uncertainty is present in the nature of the recruitment function, which captures how the availability of different habitat along with ontogenic migrations of the species affects both the carrying capacity and the growth rate of the population on the reef. Learning is accomplished by observing fish population dynamics, which a manager can influence by changing fish catches on the coral reefs and restoration or clearing of the mangrove habitat.

We demonstrate how to handle learning models that are more realistic but “inconvenient” in the sense that the function describing a decision-maker’s updated beliefs about the nature of the system does not follow a convenient closed (conjugate) form. To facilitate this we develop a method for approximating the

decision-maker's posterior beliefs using a Kullback–Leibler divergence approach. The full management model describes optimal management of both a harvestable resource (fish) and the ecosystem on which it may depend (mangroves) as a function of the state of the system. This state space includes both fish and mangrove stocks as well as current information, or beliefs, about the process governing population dynamics.

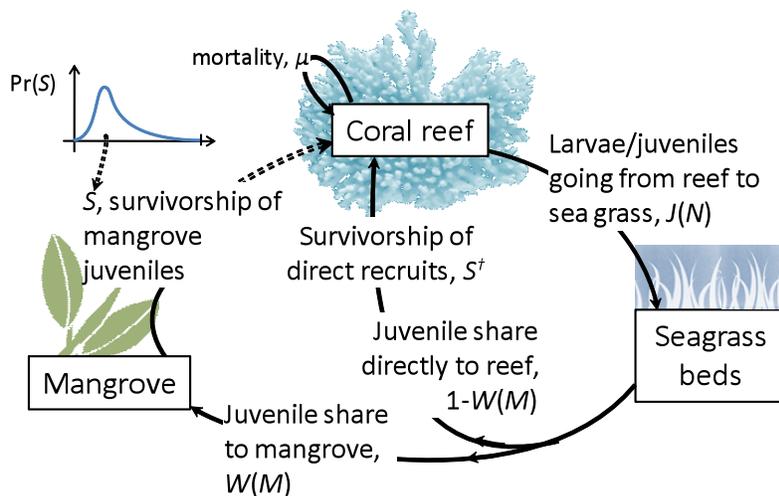
Since describing a non-trivial information state space will typically involve at least two state variables, overall we solve a four-state, two-control dynamic programming problem. After describing our model for the ecology, learning process and the economics we present initial results for a deterministic base case as and the stochastic case with learning. We find that differences in the stock of information held by the decision maker can drive strong differences in resource use policy.

## Model

We present the model in three parts: (1) the ecological model, (2) uncertainty and learning in the ecological model, and (3) the economics and integrated management problem. We take as our starting point the ecological and economic model described in Sanchirico and Springborn (*In press*). While necessary components are described here for coherence, further details and in depth analysis of the deterministic model can be found in the aforementioned article.

### *Ecological Model*

The ecological model, depicted in Fig. 1, describes the dynamics of a biological (fish) species whose life history spans coral reef, seagrass bed and mangrove habitats. Adult fish are limited to the coral reef where they are subject to a fixed rate of natural mortality ( $\mu$ ) and a time-varying rate of fishing pressure on a coral reef ( $h_t$ ). Juveniles recruit to the adult habitat either directly from seagrass beds or subsequent to an additional nursery stage within the mangroves.



**Fig. 1:** Life-cycle schematic for the mangrove, sea grass, and coral-reef fish population model. Dashed lines indicate stochastic processes. Time subscripts are suppressed; all variables are dynamic except for parameters  $\mu$  and  $S^\dagger$ .

In each period, the adult population given by  $N_t$  produces a number of juveniles according to the function

$J(N_t) = \theta N_t^\gamma$ , where  $\gamma$  and  $\theta$  are both nonnegative. If, as commonly thought, production of young demonstrates increasing returns to scale, this can be captured by setting  $\gamma > 1$ . The favorability of conditions in the seagrass beds can be described by  $\theta$ , which we assume to be constant for simplicity.

The life-cycle of juveniles is determined in part by the availability of habitat, specifically the extent of mangrove habitat,  $M_t$ . For simplicity, we measure mangrove extent as a proportion of the mangrove coverage in a pristine and undisturbed setting, thus  $M_t \in [0, 1]$ . From the seagrass beds, juveniles recruit either directly to the reef or via mangroves according to shares determined by  $M_t$ . A fraction given by  $(1 - W(M_t))$  heads directly to the reef, where  $W(M_t) \in [0, 1]$  is assumed to be a continuous function of  $M_t$ . Following the reasoning in Sanchirico and Mumby (2009) and Sanchirico and Springborn (*In press*) we assume the following conditions hold for  $W(M_t)$ : (1) if there are no mangroves, no juveniles can utilize them ( $W(0) = 0$ ); (2) even when the mangroves are at their maximum extent, some of the juveniles might recruit directly from seagrass to reef ( $W(1) \leq 1$ ); and (3) the fraction utilizing the mangroves increases as the coverage of mangroves increases, everything else being equal ( $W'(M_t) > 0$ ).

The total number of direct recruits forgoing the mangroves is given by the deterministic expression

$$K^\dagger(N_t, M_t) = J(N_t)(1 - W[M_t])S^\dagger, \quad (1.1)$$

where  $S^\dagger \in [0, 1]$  is a direct-recruit survivorship parameter.

Uncertainty enters our model in the survivorship of juveniles which utilize the mangroves. The share of juveniles whose life-cycle includes an intermediate nursery stage in the mangroves is given by  $W(M_t)$ . Let  $Z(N_t, M_t) \equiv J(N_t)W(M_t)$  represent the number of inbound juveniles to the mangroves. Survivorship of these juveniles is assumed to be a binomial random variable with a survivorship probability of  $S_t$ :

$$K(N_t, M_t | S_t) \sim \text{Binomial}(Z(N_t, M_t), S_t). \quad (1.2)$$

A description of the full uncertainty and learning model, including Equation (1.2) and the nature of  $S_t$  is provided in the next section. However, first we complete the description of the ecological model. Combining direct and mangrove recruits, recruitment to the reef at time  $t$  is equal to:

$$R(N_t, M_t) = K(N_t, M_t) + K^\dagger(N_t, M_t) \quad (1.3)$$

Before the adult population is augmented we assume that density-dependent mortality occurs among the new recruits. Following Armsworth (2002), the density-dependent process is captured by recruits competing with other recruits for space and resources during settlement. In particular, we assume that recruits enter the reef according to a Beverton-Holt recruitment function,  $G(R_t) = b_1 R_t / (1 + b_2 R_t)$  where  $b_1$  describes the survival rate at low densities, and  $b_1/b_2$  is the saturation limit with respect to the recruitment. Combining recruitment, fishing and natural mortality, the change in the fish stock on the reef is:

$$N_{t+1} - N_t = \frac{b_1 R_t(N_t, M_t)}{1 + b_2 R_t(N_t, M_t)} - \mu N_t - h_t. \quad (1.4)$$

where  $h_t$  specifies harvest, one of the two control variables.

The second control variable is effort devoted to mangrove conversion in period  $t$  which is represented by  $D_t$ . The extent of mangroves connected (within a certain distance) to the reef depends on whether the planner engages in restoration ( $D_t < 0$ ) or clearing ( $D_t > 0$ ). The mechanism by which these activities translate into changes in mangrove coverage is described by a conversion production function,  $F(D_t)$ . The mangrove dynamics are

$$M_{t+1} - M_t = F(D_t) = 1 - e^{D_t} \quad (1.5)$$

which can be positive (clearing for development) or negative (restoration) and  $F(D_t)$  is the change in mangroves.<sup>1</sup> Equation (1.5) models a process where mangrove conversion is reversible (though conversion is costly). Of course, since we include restoration or clearing as a control variable in our economic model, the planner can decide whether reversing development is optimal. Reversible development is more likely, for example, when the mangroves are cleared for aquaculture, such as shrimp farms.

We account for asymmetry in the ability to restore mangroves and clearing mangroves within  $F(D_t)$  by assuming that the marginal change  $F'(D_t)$  depends on whether  $D_t$  is positive or negative. In particular, we assume that  $F(D_t)$  has the following properties:  $F(0) = 0$ ,  $F_D < 0$ ,  $F_{DD} \leq 0$ . This captures the notion that restoring mangroves may be more difficult than clearing mangroves. Since developed areas would likely be protected from mangrove encroachment, we do not include a natural growth process for mangroves that could change the extent of coverage over time.

### ***The Learning Model and Approximate Posterior***

We now return to our model of the uncertainty embedded in equation (1.2), that is, in the process of juvenile recruitment from the mangrove habitat. A simple approach to modeling *reducible* uncertainty in this framework would be to assume that the decision maker does not know  $S_t$  with certainty but rather has beliefs about the probability that  $S_t$  takes on any particular level. With observations of “trials”  $Z_t$  and successes  $K_t$  (where functional arguments have been suppressed) each period the decision-maker could update beliefs over what the true value of  $S_t$  is. This framework would be computationally convenient but results in an overly-trivial learning process.

The framework is convenient in the following sense. A reasonable model for a random variable like  $S_t$  is a beta distribution, a flexible functional form that is restricted to the unit interval. If prior beliefs on the distribution of  $S_t$  at the beginning of a period are described by a beta distribution, then given observations on  $Z_t$  and  $K_t$  and the application of Bayes’ theorem it turns out that the posterior distribution describing updated beliefs also follows a beta distribution (see Gelman 2004, p. 34). This convenient property is known as conjugacy. In the beta-binomial case the posterior parameters are simple linear combinations of the prior parameters and  $Z_t$  and  $K_t$ .

While computationally simple, this approach results in an overly-trivial learning process. This is because it takes relatively few observations to narrow the distribution, even with a relatively diffuse prior. Learning happens unrealistically quickly—any meaningful irreducible uncertainty is wrung from the system in a handful of periods. Next we describe a hierarchical approach consistent with the notion that

---

<sup>1</sup> Given that we rescaled  $M_t$  to be a proportion of the maximum extent (pristine area), the rate of mangrove conversion,  $D_t$ , is correspondingly scaled to be in the same units.

learning about natural resource dynamics occurs in a noisy environment and takes a non-trivial amount of time.

The first level of the hierarchical model is given by Equation (1.2)—successes  $K_t$  are given by a binomial process with parameter  $S_t$ . The next step in the hierarchy is modeling the parameter  $S_t$  itself as a random variable. For simplicity, it is common in hierarchical models to assume that the distribution of an unknown parameter is given by a distribution with a known scale but an unknown location parameter (see Gelman et al. 2004, pg. 46). Consistent with this approach, we assume  $S_t$  is drawn each period from a distribution with known dispersion but an unknown mean,  $\bar{S}$ . Specifically, we model  $S_t$  as a beta random variable

$$S_t \sim \text{Beta}(\alpha, \beta), \quad (1.6)$$

where parameters  $\alpha$  and  $\beta$  are unknown.<sup>2</sup> The motivation for the “unknown location, known scale” approach is to simplify from two unknown parameters to one. For example, in a normal model, this means dealing with an unknown mean and setting aside the variance. Analogously, for the beta distribution this means simplifying from two unknown parameters— $\alpha$  and  $\beta$ —to one. However, for the beta distribution, both parameters are required for determining location and scale. Thus we now describe how, given the assumption of known scale, we simplify from two unknown parameters to one:  $\bar{S}$ .

Given the expression for the mean of a beta random variable we have

$$\bar{S} = \alpha / (\alpha + \beta). \quad (1.7)$$

For a condition describing known dispersion it would be natural to specify directly the known variance of  $S_t$ . Instead we take a computationally simpler but essentially equivalent approach of characterizing the known dispersion for  $S_t$  via the so-called “concentration parameter”

$$\rho = \alpha + \beta, \quad (1.8)$$

where  $c$  is assumed given.<sup>3</sup> Using expressions (1.7) and (1.8) the parameters of equation (1.6) are given by  $\alpha(\bar{S}) = \bar{S}\rho$  and  $\beta(\bar{S}) = (1 - \bar{S})\rho$ .

The final level in the hierarchical model is the distribution describing beliefs over  $\bar{S}$ , which is the focal variable for learning or reducible uncertainty. Given an infinite number of observations, we could expect to learn  $\bar{S}$  with certainty. For the moment, we will generically specify the prior density of beliefs over

---

<sup>2</sup> While the deterministic model of Sanchirico and Springborn (*In press*) encodes the argument that juvenile survivorship of mangrove users is greater than that of direct recruits (Chittaro et al. 2005) because they are less prone to predation (Aburto-Oropeza et al. 2008), this condition will no longer hold in every period since  $S_t$  will reflect a stochastic draw.

<sup>3</sup> The variance of  $S_t$  is given by  $\text{Var}(S_t) = \alpha\beta / ((\alpha + \beta)^2(\alpha + \beta + 1))$ , which results in a much more complicated expression for  $\alpha(\bar{S})$  and  $\beta(\bar{S})$ . The concentration parameter approach encodes information about the scale of the distribution while resulting in much simpler expressions.

$\bar{S}$ —beliefs at the beginning of any particular period—as  $f_{\bar{S}}(\bar{S})$  and return to this distribution once the updating process is clear. The challenge in the learning model is to take  $f_{\bar{S}}(\bar{S})$  and observations  $(Z_t, K_t)$  as inputs and succinctly describe posterior beliefs over  $\bar{S}$ . A key component of the posterior is the likelihood of observing  $K_t$  successes given  $Z_t$  trials and  $\bar{S}$ . This likelihood function is given by

$$L(K_t | Z_t, \bar{S}) = \int_0^1 f_{\text{Binomial}}(K_t | Z_t, S_t) f_{\text{Beta}}(S_t | \bar{S}) dS_t. \quad (1.9)$$

Using Bayes' theorem, the posterior distribution for  $\bar{S}$  is given by

$$f_{\bar{S}}(\bar{S} | Z_t, K_t) = \frac{f_{\bar{S}}(\bar{S}) \cdot L(K_t | Z_t, \bar{S})}{\int_0^1 f_{\bar{S}}(\bar{S}) \cdot L(K_t | Z_t, \bar{S}) d\bar{S}}. \quad (1.10)$$

To implement the learning model in a dynamic programming framework we must be able to succinctly summarize a particular outcome of posterior distribution in Equation (1.10) with a small number of parameters. A technical challenge is posed by the fact that, given the hierarchical model described above, we do not have the convenience of conjugacy—the structure of Equation (1.10) does not conform to a known distribution. For example, if we suppose that initially prior beliefs over  $\bar{S} \in [0, 1]$  are described with another beta distribution, the posterior is not a beta distribution as in the simple, non-hierarchical beta-binomial model.

To address this problem we approximate the posterior to support a succinct parameterized characterization to enable solution of the dynamic program. The approximation is carried out by specifying a candidate approximate function  $\hat{f}_{\bar{S}}(\bar{S}; g, h)$  for the true function  $f_{\bar{S}}(\bar{S} | Z_t, K_t)$ . The parameters  $g$  and  $h$  are fitted to minimize the divergence between the true and approximating function. Specifically we fit these parameters to minimize the so-called information divergence or Kullback–Leibler divergence between the two. While this approach has been used in the Bayesian empirical literature (Chen and Shao 1997) to our knowledge it is a novel approach in a decision-theoretic setting.

The Kullback–Leibler divergence,  $D_{KL}$ , is given by the expectation of the log difference between the two densities over the domain of  $\bar{S}$ :

$$D_{KL} \left( f_{\bar{S}}(\bar{S} | Z_t, K_t) ; \hat{f}_{\bar{S}}(\bar{S}; g, h) \right) = \int_0^1 f_{\bar{S}}(\bar{S} | Z_t, K_t) \log \frac{f_{\bar{S}}(\bar{S} | Z_t, K_t)}{\hat{f}_{\bar{S}}(\bar{S}; g, h)} d\bar{S}. \quad (1.11)$$

The final requirement is to select a suitable functional form for the candidate posterior  $\hat{f}_{\bar{S}}(\bar{S}; g, h)$  that will allow for a satisfactory fit vis-à-vis the true posterior. If the prior  $f_{\bar{S}}(\bar{S})$  is specified by a beta distribution, then it can be shown that the posterior described by Equation (1.10) is equal to a beta *density* multiplied by a beta *function*. It is also restricted to the unit interval. This suggests that a beta density is a reasonable option for  $\hat{f}_{\bar{S}}(\bar{S}; g, h)$ . While we intend to conduct a formal analysis of the quality of fit,

initial ad hoc evidence suggests that fit is good—when plotted, the true and fitted distributions in several randomly selected cases were visually indistinguishable.

Next we turn to the economic model and specify the Bellman equation. We will use  $\mathbf{I}_{t+1}(Z_t, K_t)$  as shorthand to refer to the dynamics of the information state space as described above.

### ***Economic Model***

Similar to Swallow (1990) and following the long tradition in bioeconomic modeling (Clark 1990), we model a benevolent social planner that can choose the level of mangrove conversion and fish catch in each period. In our most general formulation, controls are chosen to maximize the net present value from fishing, development, and mangrove protection.

Let  $v(h_t, D_t | N_t, M_t)$  represent the immediate benefits of harvest and development conditional on fish and mangrove stocks. These immediate benefits are comprised of fisher profit  $\pi(h_t, N_t)$ , benefits from the current extent of development  $B(1-M_t)$ , the cost of any current mangrove conversion  $C(D_t)$  and any *in situ* benefit of the mangroves  $P(M_t)$  that could be due to providing coastal protection (Barbier et al. 2008) or from intrinsic value associated with the habitat. The discount factor is described by  $\delta$ .

The Bellman equation specifying the control problem is given by

$$\begin{aligned}
 V(N_t, M_t, I_t) &= \max_{h_t, D_t} v(h_t, D_t | N_t, M_t) + \delta E_{\{N, I\}} V(N_{t+1}, M_{t+1}, \mathbf{I}_{t+1}) \\
 s.t. \quad v(h_t, D_t | N_t, M_t) &= \pi(h_t, N_t) + B(1 - M_t) - C(D_t) + P(M_t) \\
 N_{t+1} - N_t &= \frac{b_1 R_t(N_t, M_t)}{1 + b_2 R_t(N_t, M_t)} - \mu N_t - h_t \\
 M_{t+1} - M_t &= F(D_t) \\
 \mathbf{I}_{t+1}(Z_t, K_t) \\
 0 &\leq M_t \leq 1 \\
 0 &\leq N_t, \quad 0 \leq h_t \\
 N_{t=0}, M_{t=0}
 \end{aligned} \tag{12}$$

For simplicity, we will refer to  $P(M_t)$  as storm protection for the remainder of the paper. Mangroves, therefore, contribute to the value of the system indirectly through the production of fish and directly in their protection of the coastal area. Fishing profit is assumed to be increasing at a decreasing rate in harvest and fish population on the reef ( $\pi_h > 0$ ,  $\pi_{hh} \leq 0$ ,  $\pi_N > 0$ ,  $\pi_{NN} \leq 0$ ).

We model the benefits of development,  $B(1-M_t)$ , as a function of the amount of mangroves cleared (e.g. extent of total development which is  $1-M_t$  in any  $t$ ) rather than from the flow of conversion (Swallow 1990). Our approach is consistent with the idea that developed areas will return a flow of rents from some alternative use. We model the total cost of conversion by a quadratic function, which is symmetric with respect to zero and has the following properties:  $C(0) = 0$ ,  $C_{DD} > 0 \forall D$ ;  $C_D > 0$  if  $D > 0$ ; and  $C_D < 0$  if  $D < 0$ . Because in our set-up restoration is simply the negative of development, the appropriate

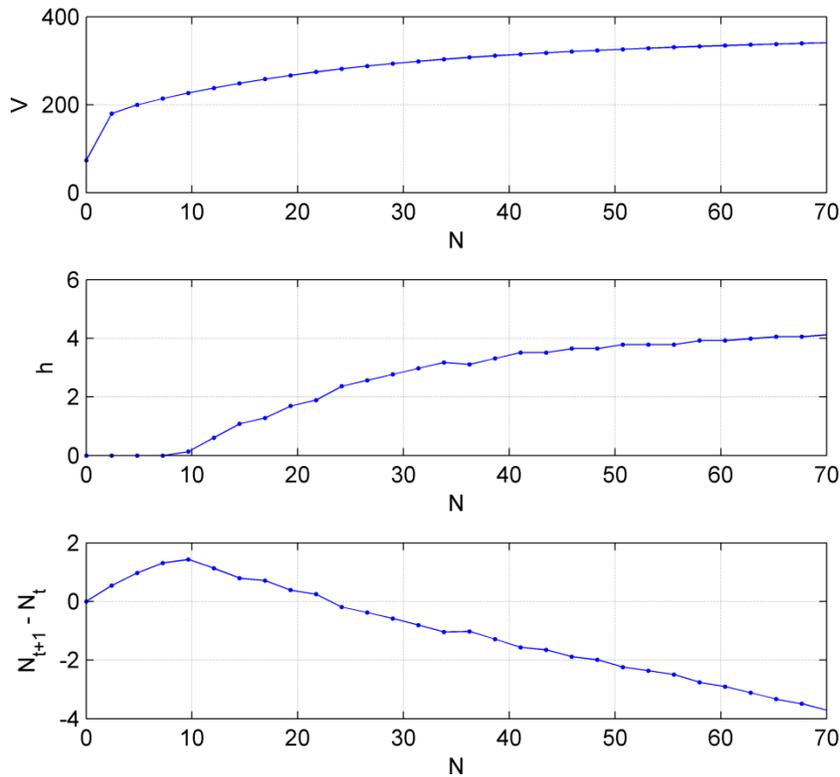
interpretation of the marginal cost of restoration is  $-C_D$  and for the marginal cost of development is  $C_D$ . The increasing cost of conversion takes into account adjustment costs that penalize the planner for either trying to ramp up restoration or development too quickly.

We also include the non-negativity restrictions on the states and control (fishing catch) along with the restriction that  $M_t$  is bounded from above by one (by assumption).

Next we present initial results from numerical solutions to the model. Parameter values for the model are listed in Table 1 in the Appendix and follow the levels used in Sanchirico and Springborn (*In press*). The dynamic programming problem is solved using value function iteration (Judd, 1998).

## Results

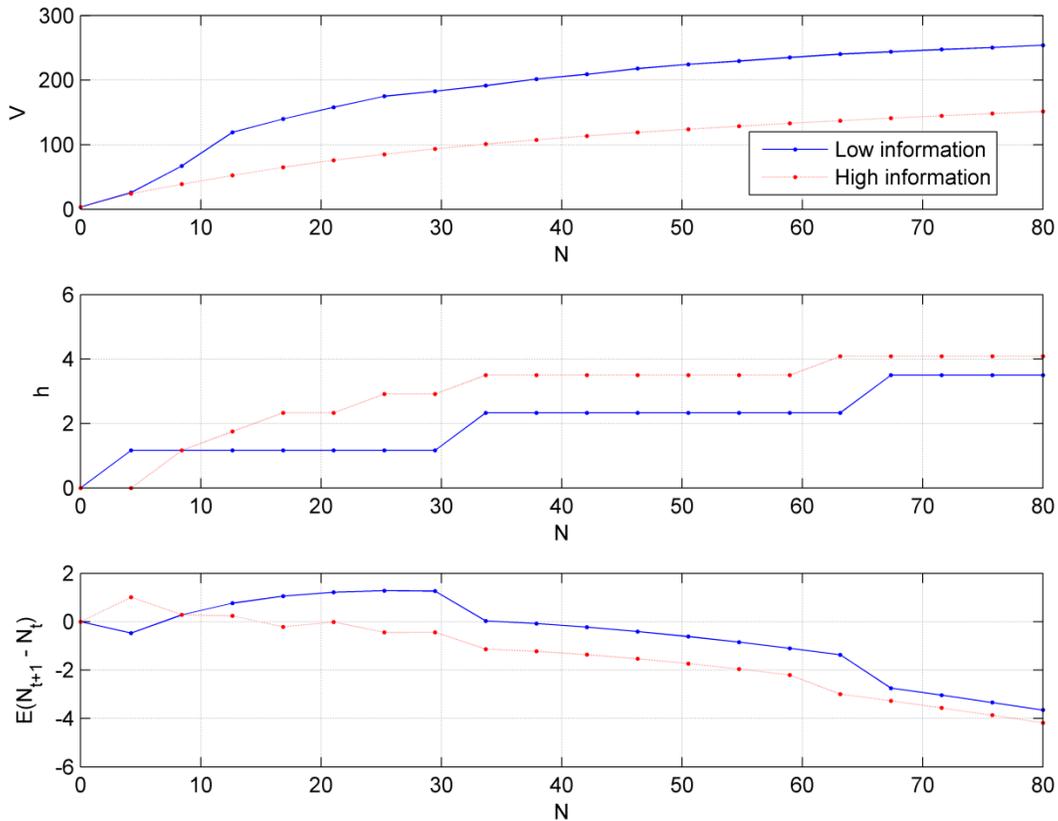
For comparison and to introduce model outcomes in the simplest fashion, we first present results for a deterministic version of the model where recruits from the mangroves stem from a survivorship parameter of  $S = 0.5$  in the same non-stochastic fashion as direct recruits. In Figure 2 we present the value function (top), policy function (middle), and change in fishery stock (bottom) are plotted as a function of the fishery stock,  $N$ . The value function does not go to zero as  $N$  goes to zero because the model includes benefits from coastal areas converted from mangroves to development,  $1-M$ . The optimal policy function for harvest in this case includes a moratorium on fishing ( $h = 0$ ) when the stock level is much lower than 10. The final plot at bottom shows the change in stock over one period given harvest, recruitment and adult mortality. The steady state population is given by the intersection with the horizontal line at zero, i.e. at  $N = 22$ .



**Fig. 2:** Results from the deterministic dynamic programming problem. The value function (top), policy function (middle), and change in fishery stock (bottom) are plotted as a function of the fishery stock,  $N$ , for the case of  $M = 0.5$  and  $S_t = S = 0.5 \forall t$ .

Next we present results from the dynamic programming problem with learning. In Figure 3 we present outcomes from one example state in which  $M = 0.5$  and  $E(\bar{S}) = 0.56$ . As in the deterministic case we depict the value function (top) and policy function (middle). The *expected* rather than actual change in fishery stock is depicted in the bottom panel since in this setting recruitment is a random variable. The solid line reflects a “low information” case where current beliefs about  $\bar{S}$  are relatively diffuse. The dotted line reflects a “high information” case where current beliefs about  $\bar{S}$  are relatively concentrated or well-informed.

In the first panel we see that the value function for the high information case lies *below* that of the low information case. While this may be surprising given that information is valuable, in the high information case beliefs in the potential for a system with high expected survivorship from the mangroves have been largely eliminated.



**Fig. 3:** Results from the dynamic programming problem with learning. The value function (top), policy function (middle), and expected change in fishery stock (bottom) are plotted as a function of the fishery stock,  $N$ . The solid line reflects a “low information” case where current beliefs about  $\bar{S}$  are relatively diffuse. The dotted line reflects a “high information” case where current beliefs about  $\bar{S}$  are relatively concentrated. In both cases,  $M = 0.5$  and  $E(\bar{S}) = 0.56$ .

In the second panel we see how the optimal policy function for harvest depends on the current state of information. We see that harvest in the low information case lies largely below that of the high information case. This is consistent with the notion that in a low information case a manager might adjust harvest to leave more fish in the system to generate a higher observation rate and increase the rate of learning.

In further work we will also examine the possibility that low information state also induces lower harvest to provide more of a buffer against the possibility of extinction. In the final panel the expected change in fishery stock adds weight to this proposition since the anticipated stock in the high information case is much lower than in the low information case. The terminology “anticipated stock” is used here since in the stochastic setting the system will not settle into a true steady state.

## Discussion

In this paper we have extended the deterministic mangrove fishery, ecosystem service model of Sanchirico and Springborn (*In press*) to incorporate irreducible uncertainty as well as uncertainty that can be reduced over time through observational learning. We have described an approach for handling Bayesian learning processes that do not conform to convenient conjugate models. We expect that this approach will expand the range of problems for which credible policy analysis with learning can be conducted. Current evidence in support of our novel approach in using a Kullback-Leibler divergence-motivated approximate posterior is only informal. In further work we will conduct a detailed analysis to explore conditions in which the method performs better or worse.

In our initial results we have shown that differences in the stock of information held by the decision maker can drive strong differences in resource use. The results presented here include only a small slice of the four-dimensional state space. In further analysis we will explore interactions between states, for example, considering the effect that the extent of the habitat ( $M$ ) has on the influence of information and stock size in setting optimal harvest. In addition we will more fully characterize how optimal policy shifts along a progression from (1) a deterministic setting, to (2) irreducible uncertainty without learning, to (3) irreducible uncertainty combined with reducible uncertainty which declines with learning.

## Appendix

**Table 1: Ecological and economic parameters**

	<i>Parameter</i>	<i>Level</i>	<i>Notes</i>
<b>Ecology</b>	$b_1$	1	Survival rate of juvenile recruits at low density
	$b_1/b_2$	10	Saturation rate of recruitment in each t
	Natural mortality rate, $\mu$	.1	10% mortality of the adult standing stock in each t, $\mu \in [0,1]$
	Seagrass survivorship rate, $\theta$	1	Survivorship of larval and juveniles in the seagrass beds, $\theta \in [0,1]$
	Larval production per adult, $\gamma$	1	If $\gamma$ is greater than one, then larval production is increasing in the adult standing stock
	Mangrove utilization, $\omega$	.5	Share of juveniles going to the mangroves is $W(M) = M^5$
<b>Economics</b>	Choke price, $\kappa_1$	7	Vertical intercept of the demand curve
	Slope of demand curve, $\kappa_2$	.75	Slope of the demand curve, when harvest equals to $\kappa_1/\kappa_2$ the price is zero
	Harvesting costs, c	20	Cost per unit of harvesting, when holding the stock size constant
	Discount factor, $\delta$	1/(1.05)	
	Benefit of development,	$v_1 = 7, v_2 = 1$	Describes the magnitude and curvature of the benefits of development
	Conversion cost, $c_d$	15	Costs per unit of conversion
	Benefit of storm protection	$\rho_1 = 7.7$ $\rho_2 = .5$	Describes the magnitude and curvature of the benefits of storm protection

Table 1: Ecological and economic parameters used in the numerical dynamic programming solution.

## References

- Aburto-Oropeza, O., Dominguez-Guerrero I, Cota-Nieto J, Plomozo-Lugo T. (2009). Recruitment and ontogenetic habitat shifts of the yellow snapper (*lutjanus argentiventris*) in the gulf of california. *Marine Biology*, 156, 2461-72.
- Aburto-Oropeza, O. , Ezcurra E, Danemann G, Valdez V, Murray J, Sala E. (2008). Mangroves in the gulf of california increase fishery yields. *Proceedings of the National Academy of Sciences of the United States of America*, 105, 10456-59.
- Armsworth, P. R. (2002). Recruitment limitation, population regulation, and larval connectivity in reef fish metapopulations *Ecology*, 83, 1092-104.
- Arvedlund, M. and Takemura, A. (2006). The importance of chemical environmental cues for juvenile *lethrinus nebulosus* forsskål (lethrinidae, teleostei) when settling into their first benthic habitat. *Journal of Experimental Marine Biology and Ecology*, 338, 112-22.
- Ascher, U. M. and Petzold, L. R. (1998). *Computer methods for ordinary differential equations and differential-algebraic equations*. Philadelphia: Society for Industrial and Applied Mathematics (SIAM).
- Barbier, E. B. (1993). Sustainable use of wetlands-valuing tropical wetland benefits-economic methodologies and applications. *Geographical Journal*, 159, 22-32.
- Barbier, E. B. (2000). Valuing the environment as input: Applications to mangrove-fishery linkages. *Ecological Economics*, 35, 47-61.
- Barbier, E. B. (2007). Valuing ecosystem services as productive inputs. *Economic Policy*, 22, 177-229.
- Barbier, E. B., Koch EW, Silliman BR, Hacker SD, Wolanski E, Primavera J, Granek EF, Polasky S, Aswani S, Cramer LA, Stoms DM, Kennedy CJ, Bael D, Kappel CV, Perillo GME, Reed DJ. (2008). Coastal ecosystem-based management with nonlinear ecological functions and values. *Science*, 319, 321-23.
- Bockstael, N. E., Freeman AM, Kopp RJ, Portney P, Smith VK. (2000). On measuring economic values for nature. *Environ. Sci. Technol.*, 34, 1384-89.
- Bryson, A. E., Jr. (1999). *Dynamic optimization*. Reading, MA: Addison Wesley Longman, Inc. p 434.
- Chittaro, P., Usseglio P, Sale. P. (2005). Variation in fish density, assemblage composition and relative rates of predation among mangrove, seagrass and coral reef habitats. *Environmental Biology of Fishes*, 72, 175-87.
- Chang, A.C. (1992) *Elements of Dynamic Optimization*. Long Grove, IL.: Waveland Press, Inc., p. 327.
- Chen, M. H. and Q. M. Shao (1997). Performance study of marginal posterior density estimation via Kullback-Leibler divergence. *Test* 6(2), 321-350.
- Clark, C. W. (1990). *Mathematical bioeconomics: The optimal management of renewable resources*. Second edition, Pure and Applied Mathematics series. New York, Chichester, p. 386.

- Daily, G. C. and Matson, P. A. (2008). Ecosystem services: From theory to implementation. *Proceedings of the National Academy of Sciences of the United States of America*, 105, 9455-56.
- Duke, N.C., J.O. Meynecke, S. Dittmann, A.M. Ellison, K. Anger, U. Berger, S. Cannicci, K. Diele, K.C. Ewel, C.D. Field, N. Koedam, S.Y. Lee, C. Marchand, I. Nordhaus, F. Dahdouh-Guebas(2007). A world with mangroves? *Science (in Letters)*, 317, 41-42.
- Freeman, A. M. (1993). *The measurement of environmental and resource values : Theory and methods*. Washington, D.C.: Resources for the Future.
- Gelman, A., J. Carlin, H. Stern, and D. B. Rubin (2004). *Bayesian Data Analysis* (2 ed.). Washington, District of Columbia: Chapman and Hall/CRC.
- Goto, N. and Kawable, H. (2000). Direct optimization methods applied to a nonlinear optimal control problem. *Mathematics and Computers in Simulation*, 51, 557-77.
- Hartmann, K, L. Bode, and P. Armsworth. (2007). The economic optimality of learning from marine protected areas. *ANZIAM J* 48(CTAC2006) pp. C307-C329.
- Heal, G. M., , Barbier E.B., Boyle K.J., Covich A.P., Gloss S.P., Hershner C.H., Hoehn J.P., Pringle C.M., Polasky S., Segerson K. & Shrader-Frechette K.. (2005). *Valuing ecosystem services: Toward better environmental decision-making*. Washington, D.C.: The National Academies Press
- Judd, K. L. (1998). *Numerical methods in economics*. Cambridge, Mass.: MIT Press.
- Kamien, M. I. and Schwartz, N. L. (1991). *Dynamic optimization: The calculus of variations and optimal control in economics and management*, 2nd edn. Amsterdam ; New York, N.Y.: North-Holland ; Elsevier Science.
- Lugo, A. E. (2002). Can we manage tropical landscapes? An answer from the caribbean perspective. *Landscape Ecol.*, 17, 601–15.
- Mumby, P. J., Edwards AJ, Ernesto Arias-Gonzalez J, Lindeman KC, Blackwell PG, Gall A, Gorczyńska MI, Harborne AR, Pescod CL, Renken H, C. C. Wabnitz C, Llewellyn G. (2004). Mangroves enhance the biomass of coral reef fish communities in the Caribbean. *Nature*, 427, 533-36.
- Murawski, S. A., J.H. Steele, P. Taylor, M.J. Fogarty, M.P. Sissenwine, *et al.*, (2010). Why compare marine ecosystems? *ICES Journal of Marine Science: Journal du Conseil* 67 1-9.
- Myers, R. A. and Worm, B. (2003). Rapid worldwide depletion of predatory fish communities. *Nature*, 423, 280-83.
- Nagelkerken, I., C. M. Roberts, G. van der Velde, M. Dorenbosch, M. C. van Riel, E. Cocheret de la Morinière, P. H. Nienhuis (2002). How important are mangroves and seagrass beds for coral-reef fish? The nursery hypothesis tested on an island scale. *Mar. Ecol. Prog. Ser.*, 244, 299-305.
- Nalle, D. J., Montgomery CA, Arthur JL, Polasky S, Schumaker NH. (2004). Modeling joint production of wildlife and timber. *Journal of Environmental Economics and Management*, 48, 997-1017.
- Polasky, S., Nelson E, Camm J, Csuti B, Fackler P, Lonsdorf E, Montgomery C, White D, Arthur J, Garber-Yonts B, Haight R, Kagan J, Starfield A, Tobalske C. (2008). Where to put things? Spatial land management to sustain biodiversity and economic returns. *Biological Conservation*, 141, 1505-24.

- Rodwell, L. D., Barbier EB, Roberts CM, McClanahan TR. (2003). The importance of habitat quality for marine reserve – fishery linkages. *Can. J. Fish. Aquat. Sci.*, 60, 171–81
- Rönnbäck, P. (1999). The ecological basis for economic value of seafood production supported by mangrove ecosystems. *Ecological Economics*, 29, 235-52.
- Sanchirico, J. N. and Mumby, P. (2009). Mapping ecosystem functions to the valuation of ecosystem services: Implications of species-habitat associations for coastal land-use decisions. *Theoretical Ecology*, 2, 67-77.
- Sanchirico, J.N. and M. Springborn. (*In press*). How to get there from here: Ecological and economic dynamics of ecosystem service provision. *Environmental and Resource Economics*.
- Simpson, S. D., Meekan M, Montgomery J, McCauley R, Jeffs A. (2005). Homeward sound. *Science*, 308, 221.
- Springborn, M., Costello, C. & Ferrier, P. (2010) Optimal random exploration for trade-related non-indigenous. *Bioinvasions and Globalization: Ecology, Economics, Management, and Policy* (ed. by C. Perrings & H. Mooney & M. Williamson), 127-144. Oxford University Press, Oxford.
- Swallow, S. K. (1990). Depletion of the environmental basis for renewable resources: The economics of interdependent renewable and nonrenewable resources. *Journal of Environmental Economics and Management*, 19, 281-96.
- Valiela, I., Bowen JL, York JK. (2001). Mangrove forests: One of the world's threatened major tropical environments *Bioscience* 51, 807-15.
- Vlassenbroeck, J. and Vandooren, R. (1988). A chebyshev technique for solving nonlinear optimal-control problems. *Ieee Transactions on Automatic Control*, 33, 333-40.