Alternative Information Regimes in Renewable Resource Management with Irreversible Regime Shifts: A State Space Approach

Craig A. Bond, Colorado State University

Working Paper, Colorado State University, November, 2011

Abstract:

Especially in systems characterized by thresholds and irreversibilities, the ex ante information available to a decision-maker has the potential to significantly affect optimal management. Most past explorations of regime shifts have assumed that the optimizing agent can either directly or indirectly observe or infer the past regime with certainty, leading to specific characterizations of cases when management is precautionary or more exploitive relative to the no-threshold cases. In this paper, we relax the assumption of resolution of uncertainty, and show that the information effects result in qualitatively different prescriptions for optimal management. In particular, even in the case of an exogenous probability of a regime shift, the strategy is no longer to manage according to the regime certainty solution, but rather to be non-precautionary. As such, persistent regime risk has a similar effect to an increase in the discount rate, but the effects are endogenous and may be complex.

Keywords: Renewable resource management, thresholds, learning-by-doing, dynamic programming, uncertainty

JEL Codes: Q2, C61, C63, D81, D83

Introduction

Management of complex ecosystems is often fundamentally a problem of choosing actions in the face of large degrees of uncertainty distinct from mere stochasticity of the underlying physical processes that generate environmental outcomes. For example, managers may have competing theories regarding the underlying structure of the system, such as the shallow lakes model in Peterson, et al. (2003), may be uncertain about a key structural parameter or parameters, as in the climate change model considered by Kelly and Kolstad (1999), or may be subject to a partially-observable process in which an underlying state cannot be observed, as in the invasive species scenario considered by Haight and Polasky (2010). Recognizing this complexity, ecologists have proposed the paradigm of "adaptive management", under which the uncertainty of the underlying dynamic process is recognized explicitly and management is prescribed as an "experimental approach that allows policy makers to learn from their mistakes and apply those lessons..." in the future (Thrower and Martinez, 2000, p. 88).

The adaptive management approach thus suggests the use of observable data to update expectations about potentially uncertain parameters or states in the dynamic system as information is generated along a management path (Ludwig and Walters, 1982). In essence, problems of this sort essentially introduce a new margin over which managers must trade off; namely, the potential for endogenous learning about the future through deviations in the management strategies that would have otherwise been "optimal" if current beliefs were to persist. In this manner, the manager can experiment by taking actions that generate a beneficial data series in terms of information content, and process this information to reduce uncertainty about the state of the system in the future, thus providing a tangible economic benefit (Bond and Loomis, 2009; Bond, 2010). Examples and the history of reasonably simple models of parametric uncertainty in primarily a regression context can be found in Kendrick (2005), with Wieland (2000) providing a dynamic programming application.

In systems characterized by thresholds and irreversibilities, however, the uncertainty may be even deeper, in that the optimizing agent(s) may not know for certain which regime is generating the ultimate outcomes, and the potential for experimentation and data generation may be limited by the fact that there is an "absorbing process" that cannot change once the threshold is crossed (or that can switch back to an original regime according to some other process). Nevertheless, endogenous learning about the potential and existence of a regime shift is possible, and management that incorporates this learning will most likely result in outcomes that dominate strategies that assume beliefs about the threshold are fixed (at least ex ante). Recent work has shed insight into optimal management strategies when the information regime is constant for these types of problems, but the implications of endogenous learning about the unknowns has not been explored (Brozovic and Schlenker, 2011; Polasky, et al. 2011). Of particular interest are the incentives related to precautionary behavior when parameter values are uncertain or state variable values are unknown.

This paper presents an analysis of endogenous learning in a resource management models that under some parameterizations admits thresholds and irreversibilities with an exogenous probability of regime shift. The model is stochastic in the biological processes, suggesting that regime uncertainty is persistent along the planning horizon, but a Bayesian learning process is modeled to allow for updating of beliefs. The framework is a discrete-time probabilistic state space model, perhaps familiar to most readers in the particular case of linear equations and independent Gaussian noise as the Kalman filter (Kalman, 1960). The state-space model is used to model the evolution of beliefs related to the unobservable components (either parameters or physical states) of the optimization problem, and is incorporated into the dynamic program to allow for active learning about the system. As seen below, several papers in the literature incorporate a state-space approach to update beliefs over either unknown parameters or unknown states; here, we allow for both.

There is a fairly extensive literature dealing with the effects of regime shifts on optimal management in a variety of ecosystems, though most include discontinuities in the expected value of the state variable once a threshold is crossed or are deterministic with respect to the regime shift (Reed, 1988; Clarke and Reed, 1994; Tsur and Zemel, 1996; Nævdal, 2003; Brozovic and Schlenker, 2011; Polasky, et al., 2011).[1] In this paper, we model a change in underlying biophysical process once a threshold is reached, rather than a discrete shock to the stock, as well as a stochastic, exogenous regime shift.[2] As such, we build on the research of Polasky, et al. (2011), who use a hazard rate in continuous time to analyze a model in which the shift in biological regime is a random variable, but the process over time is otherwise deterministic (i.e., there is no error term on the equation of motion on the resource stock). In a complementary paper, Brozovic and Schlenker (2011) examine the effects of uncertainty of threshold location on the optimal management of a stochastic shallow lake/pollution loading model. Both note that that while useful, endogenous learning is precluded from their respective specifications. The method and examples presented here explicitly address the learning framework while nesting the more restrictive information assumptions in these previous papers.

This paper contributes to the literature as follows. First, we contribute to the literature on thresholds and regime shifts by nesting a few of the basic models of system dynamic shifts into a broader framework that accounts for more realistic uncertainty regimes. In particular, we use an unobservable state variable in conjunction with a hazard parameter (rather than function) to model the probability of an exogenous regime shift, and show that the results in the previous literature are critically dependent on the resolution of one-period behind uncertainty in each subsequent time step. Second, we show how the state-space methodology can be used to model parametric uncertainty, partially-observable Markov decision processes, or both, depending on the problem at hand. In so doing, we highlight the fact that assumptions regarding the information regime (or what the

---

[1] A few examples of such ecosystems include freshwater lakes, coral reefs, grasslands and forests (both within a system and between systems), El Nino/La Nina-type ocean events, and climate change, just to name a few. See Scheffer, et al. (2001) for an excellent review.

[2] The assumption of an exogenous regime shift allows for isolation of the incentive effects associated with the information regime of the problem without confounding the incentives of an endogenous hazard.

optimizer knows, when, and how that information is used) can influence optimal management strategies, and that endogenous adaptive management strategies can be identified ex ante. This approach is fairly general, though not novel, and should be useful in a number of exercises that admit complicated uncertainty structures. Finally, our numeric results illustrate the feasibility of the approach and some key results of differing assumptions on a particular specification of a renewable resource management problem with exogenous regime shift probabilities.

The next section discusses the parameterizations that can represent alternative information regimes in dynamic problems, with a particular emphasis on a renewable resource model with a potential “systems dynamic effect” , defined as a regime shift characterized by continuity of stock levels at the time of shift as termed by Polasky, et al. (2011). We then briefly review the state space methodology that can be used to model Bayesian updating of beliefs over the uncertain parameters or states of the problem, and show how this applies to the renewable resource model. Next, three numerical simulations that vary parameter values and information regimes are presented: a competing ecosystems model, a model of regime shift with known exogenous hazard, and a regime shift model with unknown exogenous hazard. We discuss the incentives for precautionary behavior embedded in each. The final section provides some context for the results in terms of prior literature, and suggests some avenues for future research.

Treatment of Information in Dynamic Resource Management Models

A key component in models of complex systems is the treatment of the information regime under which the manager is operating; in other words, what does the manager know, and when does s/he know it? Recently, the natural resource economics literature has begun to incorporate these ideas by augmenting the state space of dynamic programs to include transition equations for both the resource itself and the beliefs about the state of the system (Walters and Holling, 1990). Broadly termed “adaptive control”-type models, applications have included fisheries (Ludwig and Walters, 1982), non-point source pollution (Kaplan et al., 2003), climate change (Kelly and Kolstad, 1999), air pollution (Cunha-e-Sa and Santos, 2008), invasive species (Springborn, 2008), shallow lakes (Peterson et al., 2003; Bond and Loomis, 2009), and general environmental policy (Brock and Carpenter, 2007; Bond, 2010). Most of these cases could be classified as problems of parametric uncertainty, in which one or more parameters of a stochastic state transition equation is not known to the manager, but rather characterized by a distribution which is updated as more data about the system becomes available.

In the case of thresholds and irreversibilities, however, the problem is potentially even more complicated, in that the unknown quantities are not parameters, but rather the values of the state of the dynamic system itself. Such partially-observable Markov systems are more rare in the literature (Haight and Polasky 2010 is a notable exception), but represent a more realistic setting faced by real-world managers, in which the evolution of the system is stochastic and there is uncertainty about the state-transition mechanisms that govern it, including both key states and parameters of the system.

To illustrate the potential information regimes, consider the discrete-time analog to the model in Polasky, et al. (2011), in which an optimizing agent maximizes the net present value of (constant) net price times harvest in each period *t*, ph_t, subject to the evolution of a stochastic stock evolution process with *i.i.d.* zero mean and constant variance errors:

$$s_{t+1} = s_t + G_i(s_t) - h_t + \varepsilon_{t+1},\ i = 1,2, \tag{1}$$

and the standard initial conditions on the stock level. Assume that before the regime shift, $G_i(s_t) = G_1(s_t)$, while $G_i(s_t) = G_2(s_t)$ thereafter, s_{t+1} is observable in period *t*+1, and these processes are known with certainty.[3] We assume $G_1(s_t) > G_2(s_t)$ and $G_1'(s_t) > G_2'(s_t)\ \forall s_t > 0$, with $G_i(\cdot)$ taking a logistic form and carrying capacity for regime 1 greater than that for regime 2. In a departure from the original model, assume an additional indicator state variable $I_t \in \{0,1\}$ that denotes if a regime change has occurred, governed by the process

$$I_{t+1} = \begin{cases} 0\ if\ I_t = 0 \text{ and } f\left(s_t, \zeta_{t+1}; \boldsymbol{\beta}\right) < \alpha \\ 1\ otherwise \end{cases}, \tag{2}$$

with ζ_t ~ iid uniform(0,1), α and $\boldsymbol{\beta}$ are parameters, and $f(\cdot)$ is a function specific to the problem under consideration.[4]

There are a variety of structural and information regimes that might be assumed for this problem. For example, assume that $f\left(s_t, \zeta_{t+1}; \boldsymbol{\beta}\right) = \zeta_{t+1}$, so that the hazard rate is defined as $(1-\alpha)$, and is interpreted as the exogenous probability of a shift in regime from regime 1 ($G_1(\cdot)$) to regime 2 ($G_2(\cdot)$) when the true data generation process is regime 1. At least four possible information regimes can be assumed, depending on the assumptions regarding the ability to observe ζ_{t+1} and whether or not α is known.

In the case of known hazard rate with observability of ζ_{t+1}, the problem is stochastic in the regime switch but fully observable otherwise, and is essentially an identical information treatment as the case of exogenous probability of regime shift with system dynamics effect in Polasky, et al. (2011). In that model, the state equation on the resource stock is assumed deterministic, but note that so long as s_{t+1} is observable, the lack of a stochastic process corresponds to the assumption of (perfect) observability of the regime shift. As such, in this treatment, the model is a straightforward stochastic dynamic programming problem with 2 states, and there is no need to include beliefs in the state space (though it can be modeled as such, as seen in the next section).

---

[3] The assumption that the net growth functions are known with certainly can be relaxed given the framework presented here. Also note that we assume no discontinuity in the state in order to focus on the treatment of information in the problem. Such discontinuities could be incorporated through, say, including an additional value of the discrete state I_t that is associated with the discontinuity, and a physical state transition that accounts for the stock change at that level.

[4] In other words, $G_i(s_t) = G_1(s_t)$ if $I_{t+1} = 0$, $G_i(s_t) = G_2(s_t)$ otherwise.

Retaining the assumption of known α but relaxing the observability assumption on the regime switch results in a model with no parametric uncertainty in the adaptive control sense, but with an unobservable state variable I_t, since the resource stock evolves stochastically and is the only observed state of the system from period to period. We thus have a partially-observable Markov sequence with constant, known parameters, and assume that a manager's optimal actions will depend on her beliefs about the current (and future) growth regimes. As such, we augment the state space to include the probability that the regime switched, and use an updating rule to describe the evolution of these beliefs.

If, on the other hand, α is unknown, regardless of the observability of ζ_{t+1}, we have a case of both parametric uncertainty and a partially-observable Markov sequence, as the error term on the resource stock transition equation precludes the manager from knowing for sure which regime is operative. This is not to say, however, that the problems are equivalent, since the ability to observe the random draw governing the stochastics of the regime shift provides additional information over which to update the belief space, which now consists of the joint probability that the regime shift has occurred and the value of the unknown parameter.

These same information regimes are possible if one allows for various endogenous specifications of the probability of the regime shift (e.g., $s_t < \alpha$, $\zeta_{t+1} < f(s_t;\boldsymbol{\beta}),\ 0 \le f(s_t;\boldsymbol{\beta}) \le 1,$ etc...), degrees of irreversibility within the problem, and/or other extensions depending on the problem. For example, in the case of a purely stochastic process governing a switch "back", one may rewrite the transition equation related to the indicator as

$$I_{t+1} = \begin{cases} 0\ if\ \ \left(I_t = 0 \text{ and } f\left(s_t, \zeta_{t+1}; \boldsymbol{\beta}\right) < \alpha\right) \text{ or } \left(I_t = 1 \text{ and } \xi_{t+1} < \gamma\right) \\ 1\ otherwise \end{cases}, \tag{3}$$

where $\xi_t \sim uniform(0,1)$ and the probability of a reverse regime shift is given by $0 \le \gamma \le 1.$ The information regime in this case is defined by the assumptions regarding the nature of the parameter vector $(\boldsymbol{\beta}, \alpha, \gamma)$ (with elements either known and constant or unknown) and the observability of the stochastic elements of the problem (ζ_{t+1}, ξ_{t+1}). As shown below, the information regime could have considerable effects on the optimal management plan, just as the endogeneity of the regime switch did in Polasky, et al. (2010) or the variance associated with the unknown threshold did in Brozovic and Schlenker (2011). However, to fix ideas in what follows, we assume $f(\cdot) = \zeta_{t+1}$ and $\gamma = 0$, rendering regime 2 the absorbing state and the regime shift irreversible.

State Space Methodology

This section demonstrates a unifying state-space filtering framework for treatment of alternative information regimes in resource management problems, and provides examples within the context of exogenous hazards in the regime-switching renewable resource problem above.

A Primer on State Space Modeling[5]

A discrete-time state space model is a representation of a Bayesian probabilistic process in which inferences about unobserved "states" of the filtering system, $\mathbf{z_t} \in \mathbb{R}^n$, are made through processing of observable information $\mathbf{y_t} \in \mathbb{R}^m$.[6] The model is defined by

$$\begin{aligned} \mathbf{z_{t+1}} &\sim g\left(\mathbf{z_{t+1}} \,|\, \mathbf{z}_t\right) \\ \mathbf{y_{t+1}} &\sim g\left(\mathbf{y_{t+1}} \,|\, \mathbf{z_{t+1}}\right), \end{aligned} \tag{4}$$

where $g\left(\mathbf{z_{t+1}} \,|\, \mathbf{z_t}\right)$ describes the dynamic process of the unobservable quantities and $g\left(\mathbf{y_{t+1}} \,|\, \mathbf{z_{t+1}}\right)$ describes the distribution of the observables conditional on the contemporaneous values of the unobservables. It is assumed that (4) is a Markov process, and thus the history of the system through time t is completely characterized by $\mathbf{z_t}$ and $\mathbf{y_t}$.

This type of model can be used in the context of dynamic programming models with partially-observable Markov processes, parametric uncertainty, or both, to model the evolution of beliefs about the unknown quantities as a state transition equation. To do so, define the predictive distribution of $\mathbf{z_{t+1}}$ as

$$g\left(\mathbf{z_{t+1}} \,|\, \mathbf{y_t}\right) = \int g\left(\mathbf{z_{t+1}} \,|\, \mathbf{z_t}\right) g\left(\mathbf{z_t} \,|\, \mathbf{y_t}\right) d\mathbf{z_t}, \tag{5}$$

where $g\left(\mathbf{z_t} \,|\, \mathbf{y_t}\right)$ is the prior over $\mathbf{z_t}$ at time t (perhaps from a previous update of the system). Equation (5) is known as the Chapman-Kolmogorov equation, which describes the dynamics of the state of the system conditional on what can be observed at time t, and is essentially a marginal distribution obtained by the integration of the conditional distribution $g\left(\mathbf{z_{t+1}} \,|\, \mathbf{z_t}\right)$ times the prior.

Once the predictive distribution is obtained, Bayes' rule can be used to update the prior and obtain the new distribution of the unobservables $\mathbf{z_{t+1}}$ conditional on the observed data $\mathbf{y_{t+1}}$:

$$g\left(\mathbf{z_{t+1}} \,|\, \mathbf{y_{t+1}}\right) = \frac{g(\mathbf{y_{t+1}} \,|\, \mathbf{z_{t+1}}) g\left(\mathbf{z_{t+1}} \,|\, \mathbf{y_t}\right)}{\int g(\mathbf{y_{t+1}} \,|\, \mathbf{z_{t+1}}) g\left(\mathbf{z_{t+1}} \,|\, \mathbf{y_t}\right) d\mathbf{z_{t+1}}}. \tag{6}$$

Equation (6), also known as the filtering distribution, thus provides an implicit equation that can be used to update distributional information about the unobservable parameters or dynamic programming state variables between time periods through modeling the

[5] The general framework described here relies heavily on Särkkä (2006). The reader is referred to this resource for additional details.

[6] The terminology "state of the filtering system" is used as in the state-space literature, where the "state" is a (possibly unobservable) parameter value and the measurements are the observable features of the system. From a dynamic programming standpoint, both s_t and I_t are states of the program.

evolution of the sufficient statistics $\boldsymbol{\theta}_{\mathbf{t}+1}$ of the distribution $g\left(\mathbf{z}_{\mathbf{t}+1} \,|\, \mathbf{y}_{\mathbf{t}+1}\right)$. We illustrate the use of the general model in the context of thresholds and irreversibilities below.

Modeling an Exogenous, Known Hazard

Consider the case where $g\left(\mathbf{z}_{\mathbf{t}} \,|\, \mathbf{y}_{\mathbf{t}}\right)$ represents the probability that the regime has switched given the contemporaneous data on the renewable stock, and define the prior at time *t* as $\Pr(I_t = 1 \,|\, s_t) = \pi_t$. As will be shown, we can model both ζ_{t+1} observable or unobservable and α known or unknown using the state space specification and filtering methodology.

Begin with the case of α known, and let the state of the filtering system in period *t*+1 be I_{t+1}, with observable measurements $s_1 \ldots s_{t+1}$ in each time period.[7] Our objective from filtering is to find $\Pr(I_{t+1} = 1 \,|\, s_0 \ldots s_{t+1})$, which is trivial in the case of perfect observation. Nevertheless, it is instructive to set up the filter to trace through the predictive and filtering steps.

The predictive distribution is given by $\Pr(I_{t+1} = 1 \,|\, s_0 \ldots s_t),$ and the filtering distribution as $\Pr(I_{t+1} = 1 \,|\, s_0 \ldots s_{t+1}).$ The former is defined by

$$\begin{aligned} &\Pr(I_{t+1} = 1 \,|\, s_0 \ldots s_t) \\ &\quad = \Pr(I_t = 0)\Pr(I_{t+1} = 1 \,|\, I_t = 0) + \Pr(I_t = 1)\Pr(I_{t+1} = 1 \,|\, I_t = 1) \\ &\quad = (1 - \pi_t)\Pr(I_{t+1} = 1 \,|\, I_t = 0) + \pi_t \Pr(I_{t+1} = 1 \,|\, I_t = 1). \end{aligned} \tag{7}$$

$\Pr(I_{t+1} = 1 \,|\, I_t = 1) = 1$ for the purely irreversible case, regardless of observability of the process governing the switch. In the case of purely exogenous known hazard with observable ζ_{t+1}, $\Pr(I_{t+1} = 1 \,|\, I_t = 0, s_0 \ldots s_t) = \Pr(I_{t+1} = 1 \,|\, I_t = 0, \zeta_{t+1}),$ and thus $\Pr(I_{t+1} = 1 \,|\, I_t = 0, \zeta_{t+1}) \in \{0,1\}$. As such, $\Pr(I_{t+1} = 1 \,|\, s_0 \ldots s_t) \in \{\pi_t, 1\}$ depending on the realization of the stochastic term. Note that uncertainty over the unknown state is perfectly resolved if $\zeta_{t+1} \geq \alpha,$ in which case a switch has certainly happened; otherwise, the predictive distribution is equal to the prior. However, if ζ_{t+1} cannot be observed directly, $0 \leq \Pr(I_{t=1} = 1 \,|\, s_0 \ldots s_t) = (1 - \pi_t)(1 - \alpha) + \pi_t \leq 1,$ with strict inequalities so long as $0 < \alpha < 1$. In the presence of confounding errors on the resource transition equation, then, it is possible that the maximizing agent does not know if the regime has shifted or not.

Regardless of the assumption on observability of the process governing the shift, the filtering distribution can be used to obtain the posterior of the distribution once the new data arrives in the form of s_{t+1}. Through straightforward application of Bayes' rule,

[7] It is not necessary to assume measurement in each time period, though this assumption is maintained throughout the paper.

$$\Pr(I_{t+1}=1\,|\,s_0\ldots s_{t+1})=\frac{g(s_{t+1}\,|\,I_{t+1}=1)\Pr(I_{t+1}=1\,|\,s_0\ldots s_t)}{\sum_{a\in\{0,1\}} g(s_{t+1}\,|\,I_{t+1}=a)\Pr(I_{t+1}=a\,|\,s_0\ldots s_t)}, \tag{8}$$

where $g(\cdot)$ is the distribution (or kernel of the distribution) of s_{t+1} and $g(s_{t+1}\,|\,I_{t+1}=1)$ is the likelihood of observing s_{t+1} conditional on $I_{t+1}=1.$ With observability, (8) implies that $\Pr(I_{t+1}=1\,|\,s_0\ldots s_{t+1})=1$ if $\zeta_{t+1}\geq\alpha,$ and the data adds no additional information. If $\zeta_{t+1}<\alpha$ or with unobservable ζ_{t+1}, however, this is no longer strictly the case, and Bayes' rule gives an efficient information processing rule for modeling the evolution of beliefs about the state of the system, dependent on all observations up to time t+1 (perhaps embodied by the priors in the case of t=1), the (assumed known) hazard rate, and the (assumed known) distribution of ε_{t+1}.[8]

Equations (7) and (8) thus provide the necessary information to create a difference equation that describes the evolution of π_t over the course of the optimization problem by assuming that the posterior distribution in (8) becomes the new prior before each new decision is made.

Modeling an Exogenous, Unknown Hazard

The next layer of complexity to consider is an unknown hazard rate; that is, α in (2) is not known with certainty, but rather can be characterized by a (discrete or continuous) probability distribution function $g(\alpha).$ As in the last section, the objective is to find an expression for the (now bivariate) filtering distribution $g(I_{t+1}=j,\alpha_{t+1}\,|\,s_0\ldots s_{t+1}),\ j\in\{0,1\}.$ The notation makes clear that in this case, the model admits both a partially-observable dynamic process in physical (s_t) space and parametric uncertainty through the unknown $\alpha.$

Begin with a joint prior distribution over both parameters in time t; say, $\pi_{I=j,\alpha,t}=g\left(I_t=j,\alpha\right),\ j\in\{0,1\},$ and denote the sufficient statistics of this distribution $\boldsymbol{\theta}_t.$ As before, we use the filtering mechanism to create the probability updating equation $g(I_{t+1}=j,\alpha_{t+1}\,|\,s_0\ldots s_{t+1})$ using the predictive and filtering distributions. The former now takes the form

$$g(I_{t+1}=j,\alpha_{t+1}\,|\,s_0\ldots s_t)=\sum_{k\in\{0,1\}}\int \pi_{I=j,\alpha,t}\,g(I_{t+1}=j,\alpha_{t+1}\,|\,I_t=k,\alpha_t,)da_t, \tag{9}$$

while the latter generalizes to

$$g(I_{t+1}=j,\alpha_{t+1}\,|\,s_0\ldots s_{t+1})=\frac{g(s_0\ldots s_{t+1}\,|\,I_{t+1}=j,\alpha_{t+1})g(I_{t+1}=j,\alpha_{t+1}\,|\,s_0\ldots s_t)}{\sum_{k\in\{0,1\}}\int g(s_0\ldots s_{t+1}\,|\,I_{t+1}=k,\alpha_{t+1})g(I_{t+1}=k,\alpha_{t+1}\,|\,s_0\ldots s_t)d\alpha_{t+1}}. \tag{10}$$

---

[8] Note that the case of unobserved ζ_t and deterministic evolution of the resource is also nested within this framework, as in this case the likelihood $g(s_{t+1}\,|\,I_{t+1}=1)\in\{0,1\}.$

In the case of a discrete number of support points on α_{t+1}, the integration becomes summation over those points.

There are basically no fundamental conceptual differences between (10) and (8), save for the addition of an unknown (possibly continuous) random variable α which adds some complexity to the relationships, and an application to an endogenous hazard rate is straightforward.[9] In some cases, there may be closed form solutions to (10), in that the evolution of the sufficient statistics of $g(I_{t+1} = j, \alpha_{t+1} \mid s_0 \ldots s_{t+1})$ can be modeled explicitly. Examples include cases where the distributions of α and ε_{t+1} are conjugate, and thus the resultant posterior is of identifiable form; for example, Gaussian priors would enable the use of a slightly modified Kalman filter that admits closed-form solutions, as would the assumption of discrete distribution on the hazard rate. In the case of lack of closed-form solutions, numerical techniques may be used to model the evolution of belief space, so long as sufficient statistics could be identified. There is little doubt that such techniques are generally complicated (and suffer from the curse of dimensionality when used in a dynamic programming context), but the advance of computational methods and associated computing speeds have increased the feasible set of problems that can be solved this way (e.g., Rust, 1997; Han et al., 2006). In addition, significant insights might be gained from simple problems that focus on a discrete belief space. It is to such examples that we now turn, with a focus on the implications of these assumptions on precautionary behavior within the renewable resource model.

Numerical Examples

In most cases, analytical insight into the more complicated of these problems is not likely, especially given the large expansion of the state space associated with the introduction of the joint priors in belief space. Even using modern numerical approximation methods such as flexible polynomial forms for the unknown value function, may prove taxing for realistic problems (Judd 1999, Miranda and Fackler, 2002), especially given the natural probabilistic restrictions on the state space. Here, however, we restrict attention to relatively simple problems where basic methods are feasible.[10] In addition, because analytical solutions to a subset of these problems are possible, the numerical results of the restricted models were externally validated.

In order to explore the effects of non-observability and parametric uncertainty on the optimal management of the renewable resource three stochastic models are presented under alternative parameterizations and information assumptions. All take the form:

[9] Recall that in a Bayesian rather than frequentist framework, the underlying true parameter is interpreted as random.

[10] In particular, we use a linear splines and discretized controls over the state space in each case, resulting in rather crude but comparable approximations to the true value function. The numerical models were coded using the COMPECON toolbox of Miranda and Fackler (2002). The author is grateful to Paul Fackler for sharing his code for building the simplex in the model with four state variables.

$$
\begin{aligned}
&\max_{h_1 \ldots h_\infty} \sum_{t=0}^{\infty} \delta^{t} p h_t \\
&s.t.\ s_{t+1} = \begin{cases} s_t + G_1(s_t) - h_t + \varepsilon_{t+1} \text{ if } I_{t+1} = 0 \\ s_t + G_2(s_t) - h_t + \varepsilon_{t+1} \text{ if } I_{t+1} = 1 \end{cases} \\
&\quad I_{t+1} = \begin{cases} 0 \text{ if } I_t = 0 \ \& \ \zeta_{t+1} < \alpha \\ 1 \ otherwise \end{cases}, \\
&\quad s_0 > 0,\ 0 \le \Pr(I_0 = 1) \le 1
\end{aligned}
$$

In each case, we assume parameter values as in Table I, with growth in each regime assumed to be logistic and occurring after harvest is made: $G_i = \left(1 - (s_t - h_t)\frac{1}{k_i}\right)(s_t - h_t)$.[11] The assumption of an unknown I_t is maintained, with differences in each model defined by alternative assumptions over α, the parameter which determines the probability of a regime shift/hazard rate. Additionally, we assume that the decision-maker in the problem updates over the unknown distributions using the state-space framework presented in the previous section.[12]

Competing Ecosystems Model

The first specification assumes a known $\alpha = 1$, which implies stationary of regime over the planning horizon. As such, it is essentially a competing ecosystems model, in which the decision maker may be unsure of which state transition regime is operative. This case is presented to illustrate the incentives related to endogenous learning without confounding the incentives related to the regime shift.

The Bellman equation becomes

$$V(s_t, \pi_t) = \max_{h_t} ph_t + \delta\left\{(1-\pi_t)E\left[V\left(s_{t+1}(G_1), \pi_{t+1}(G_1)\right)\right] + \pi_t E\left[V\left(s_{t+1}(G_2), \pi_{t+1}(G_2)\right)\right]\right\}, \quad (11)$$

where the notation indicates the dependence of the next-period state variables on the assumed growth regime, which in turn depends on stock levels and harvest at time *t*. In addition, beliefs depend on the prior belief state in *t* as well.

Figure 1 presents the approximate value function for this problem. Given the nature of the information updating process, beliefs at the limits (i.e., $\pi_t = 0$ and $\pi_t = 1$) do not update under any realizations of the future stock, and thus $V(s_t, 0)$ and $V(s_t, 1)$ are equivalent to the value functions under regime certainty, with corresponding steady-state stock and control values of approximately (0.72, 0.25) and (0.52, 0.175), respectively. For stock levels greater than the steady-state in each case, the value functions are linearly increasing in the stock with identical slopes, a result of the assumption that instantaneous

[11] Numerically, the assumption of growth after harvest simplifies coding given the natural non-negativity of stocks. This does change steady-state values relative to logistic growth in s_t alone, but does not alter the incentives in each information regime.

net benefits are linear in harvest. For stock levels less than the steady-state, the value function is concave as a result of the assumptions on the growth function and the non-negativity constraint on harvest. In these cases of known regime, the optimal solution is a most-rapid approach path (MRAP) that drives the stock to the steady-state level as soon as possible (immediately if the starting stock level is greater than the steady state).

For intermediate beliefs over the true regime, the value function is similarly shaped with respect to the stock, but is scaled up or down in accordance with the linear weights in the Bellman equation. The intuition is that in the presence of uncertainty over the true regime, there is a trade-off in potentially lost benefits if the decision-maker operates according to, say, regime 1 certainty, but the system is characterized by regime 2. As such, the ex ante expected value is decreasing in π_t, as is the point in stock space where the value function turns linear. Ceteris paribus, then, one would expect optimal control and stock values to increase as the belief that regime 1 is the true regime is strengthened. These beliefs are updated based on the observed stock in each period, thus endogenizing the probabilities on the right-hand side of (11) and changing the nature of the solution from a MRAP to a more gradual driving of the stock towards the levels associated with the limiting beliefs.

This is confirmed in the simulations presented in Figure 2, which depict the optimal paths corresponding to $\varepsilon_{t+1} = 0 \; \forall t$ for cases in which initial beliefs are incorrect relative to the true regime, with starting stock levels equal to 0.5.[13] In panel I, the true data generation regime has growth $G_1(\cdot)$ with $\pi_0 = 0.99$, while panel II assumes growth regime $G_2(\cdot)$ with $\pi_0 = 0.01$. In each case, the optimal control in the first period is essentially identical to the regime certainty case, which (were the beliefs true) would result in driving the resource stock to the corresponding steady state level. However, since the true regime is not consistent with beliefs, this results in stock levels in period 2 to be high (relative to expectations) in panel I, and low in panel II. Learning from these observations, the decision-maker updates probabilities accordingly, placing slightly more weight on the true regime, and subsequently acts accordingly.

The end result for panel I are control and resource stock paths that are slowly increasing towards the steady state levels for regime 1 as the endogenized weights in (11) adjust towards the true value. The speed of adjustment is directly related to the size of the adjustment in belief space, as can be seen by the increased slopes on the control and stock paths as the absolute value of the slope of the belief path increases. In panel II, since the initial control based on incorrect beliefs drove the stock path higher than the steady-state level associated with $G_2(\cdot)$, the stock level is gradually declining. Note, however, that as in panel I, extraction is increasing over the planning horizon. The explanation is that the positive relationship between extraction and π_t conditional on the stock level dominates

[13] If initial beliefs were essentially correct, the paths would be numerically indistinguishable from the regime certainty case. Intermediate beliefs produce state and control paths that fall in between these polar cases.

the positive relationship between the extraction and the stock level conditional on π_t, primarily due to the rapidity of the updating of beliefs.

Known, Exogenous Hazard

The second model assumes $0 < \alpha < 1$, and that the decision-maker knows the hazard rate $(1-\alpha)$ that determines the probability of a regime shift from regime 1 to regime 2. If one assumes that either a) ζ_{t+1} is observable or b) $\varepsilon_{t+1} = 0$ and this fact is known, the model is equivalent to the exogenous probability of regime shift model in Polasky, et al. (2011) (see Appendix A). In this case, the optimal solution assuming an initial state $I_0 = 0$ is for the manager to apply the MRAP solution in the regime 1 certainty solution; in other words, a known, exogenous hazard with observability over the regime in time *t* does not affect optimal management.[14]

Our insight here is that this result is a consequence of a) the equivalence of the marginal value of a unit of stock at the steady-state of each regime due to the linearity of net benefits with respect to harvest; and b) the resolution of uncertainty over the operative regime in time period *t* from the standpoint of time period *t*+1. Starting with the latter, the assumption that the uncertainty about *t* is resolved in the subsequent period implies that the belief updating process is *independent* of the control variable, and that subsequent weights on each regime will be either $\pi_{t+1} = 0$ or $\pi_{t+1} = 1$. As such, there is no dynamic trade-off to be made in the belief dimension; that is, from the perspective of time *t*, the decision-maker need not trade off between decisions today and value differences induced by differences in beliefs tomorrow (as documented in the previous subsection). Coupled with the fact that the marginal value of increasing s_t is equivalent in each regime, it makes no difference to the decision maker at the control margin if the system is currently in regime 1 or regime 2 when $\pi_t = 0$, as there is no penalty for being wrong given the resolution of uncertainty, and thus no trade-off and no change in behavior.

If the deterministic or observability conditions are not met, however, this result no longer holds. Using the state-space specification of this now partially-observable Markov decision model, the Bellman equation for this problem becomes

$$V\left(s_t, \pi_t\right) = \max_{h_t} ph_t + \delta\Big\{\left(1-\pi_t\right)\Big(\alpha E\Big[V\left(s_{t+1}(G_1), \pi_{t+1}(G_1)\right)\Big] + (1-\alpha)E\Big[V\left(s_{t+1}(G_2), \pi_{t+1}(G_2)\right)\Big]\Big)$$
$$+\pi_t E\Big[V\left(s_{t+1}(G_2), \pi_{t+1}(G_2)\right)\Big]\Big\},$$

and the operative regime at any time period (save perhaps *t*=0) is not known. Just as in the previous section, the probability weights π_t are now endogenized, are no longer independent of h_t, and will change as decisions are made and beliefs updated along a

[14] Appendix A also documents the main results in Polasky, et al. (2011) using the indicator state-variable framework.

management path.[15] Compared to the competing ecosystems model, however, the predictive distribution associated with the unknown indicator state will cause the expectation of next-period beliefs π_{t+1} to increase given that $0 < \alpha < 1$, and thus ex ante values at a particular physical/belief state will be lower than in the previous case (see panel II in Figure 1).

Figure 3 shows the optimal path in this case, again assuming $\varepsilon_{t+1} = 0 \ \forall t$ with $\pi_0 = 0$ and $s_0 = 0.5$, with the true regime starting with $I_0 = 0$ and shifting in period *t*=26. The first item of note is that relative to the regime 1 certainty case, the optimal control in the first period increases from 0.025 to 0.045, suggesting that the optimal initial decision is to drive the resource stock level *lower* than the regime 1 certainty case. The reason is that unless there is a strong negative shock to the resource, π_{t+1} will increase, and there continues to be a positive relationship between extraction and π to account for the optimality of a lower steady state under regime 2 growth. This manifests itself in the second argument of the value function, as the agent takes into account the additional information margin. Even if there remains an equivalence of the marginal value of the state s_t, as there would be if $\pi_t = 0$, the agent must trade off instantaneous benefits today with the (likely lower) expected values of the system tomorrow given the change in expected beliefs, and thus the problem becomes truly dynamic in the belief dimension. In this sense, the optimal initial control is non-precautionary, even though the probability of a regime shift is exogenous.

Continuing along the optimal path, given the assumptions of the simulation, the state path gradually declines with the control value until such time as the regime shifts. Unlike the competing ecosystems case, then, the positive relationship between extraction and π_t conditional on the stock level is dominated by the positive relationship between the extraction and the stock level conditional on π_t. Note that the belief path is concave with respect to time, contributing to this result. Following the regime shift in *t*=26, there is a discrete jump downwards in both the state and control path as the observation of the resultant (lower) state causes π_t to increase relatively rapidly. At this point in the simulation, the belief in regime 2 is relatively high (~0.81), which results in quick (downward) convergence to the steady-state solution for regime 2, which persists for the remainder of the simulation as π_t reaches the limit belief of 1 as $t \to \infty$.[16]

[15] Note that regardless of the source of information or the updating rule used by the decision-maker, so long as beliefs over the unknown state differ from time period to time period, the control decision will change. This is consistent with the model assuming observability...if beliefs change from regime 1 to regime 2 between *t* and *t*+1, h_t and h_{t+1} differ.

[16] Additional simulations (not shown) confirm this qualitative pattern of results when the regime shift occurs at different times over the planning horizon. Differences occur in the rapidity of convergence of beliefs and the amount of time it takes the resource state/control variables to reach the regime 2 steady-state level as a result of the shift taking place when beliefs in regime 1 are relatively stronger than in the presented simulation.

In a case where the regime does not shift despite the exogenous hazard (not shown), the qualitative state, control, and belief paths are similar, but are smooth given the lack of a change in the data generation process. One key difference, however, is that the state and control variables converge to a higher level even when the belief in regime 2 is very close to one, as growth in the resource stock is still governed by (G_1), yet the optimal control rule under this belief is to extract at a higher level in an attempt to drive the state variable to the regime 2 certainty steady state. Interestingly, this is a case in which the limiting belief will (in a proababilistic sense) ultimately be "correct", i.e., $\lim_{t\to\infty} I_t = 1$, but the belief for any finite *t* is incorrect in that π_t is (infinitely) close to one yet the true data generating regime is consistent with $\pi_t = 0$.

Unknown, Exogenous Hazard

The final model illustrated here assumes that α can take on one of two distinct values, say $\alpha \in \{\alpha_1, \alpha_2\}$, but the decision-maker is uncertain which is the true value. Let $\boldsymbol{\theta}_{\mathbf{t}} = (\pi_{I=0,\alpha=\alpha_1,t}, \pi_{I=1,\alpha=\alpha_1,t}, \pi_{I=0,\alpha=\alpha_2,t})$. The Bellman equation for this problem contains four state variables, and takes the form

$$\begin{aligned}
V\left(s_t, \boldsymbol{\theta}_{\mathbf{t}}\right) &= \max_{h_t} ph_t + \delta E\left[V\left(s_{t+1}, \boldsymbol{\theta}_{\mathbf{t+1}}\right)\right] \\
&= ph_t + \delta\Big(\pi_{I=0,\alpha=\alpha_1,t} E\left[V\left(s_{t+1}(G_1), \boldsymbol{\theta}_{\mathbf{t+1}}(G_1)\right) \mid I_t = 0 \,\&\, \alpha = \alpha_1\right] \\
&+ \pi_{I=1,\alpha=\alpha_1,t} E\left[V\left(s_{t+1}(G_2), \boldsymbol{\theta}_{\mathbf{t+1}}(G_2)\right) \mid I_t = 1 \,\&\, \alpha = \alpha_1\right] \\
&+ \pi_{I=0,\alpha=\alpha_2,t} E\left[V\left(s_{t+1}(G_1), \boldsymbol{\theta}_{\mathbf{t+1}}(G_1)\right) \mid I_t = 0 \,\&\, \alpha = \alpha_2\right] \\
&+ \pi_{I=1,\alpha=\alpha_2,t} E\left[V\left(s_{t+1}(G_2), \boldsymbol{\theta}_{\mathbf{t+1}}(G_2)\right) \mid I_t = 1 \,\&\, \alpha = \alpha_2\right],
\end{aligned} \tag{12}$$

where the last prior probability is calculated as $(1 - \pi_{I=0,\alpha=\alpha_1,t} - \pi_{I=1,\alpha=\alpha_1,t} - \pi_{I=0,\alpha=\alpha_2,t})$. The specific form of (12) is very similar to the model in the previous subsection, but conditioned on the values of the unknown hazard parameter. In the case of $I_t = 1$, the expectation over the unknowns is independent of α, since it is known that only regime 2 growth is possible.

A priori, one might expect that relaxing the assumption of a known hazard rate may not affect the results given that there is no means for the decision-maker to manipulate the system and learn about this parameter before a regime shift. However, as shown below, the introduction of this parametric uncertainty has significant consequences for the evolution of beliefs about the operative regime, and as these beliefs imply differences in expected future values, they can affect the degree of (non)-precaution in behavior.

We model the case of $\alpha \in \{1, 0.9\}$, which is a mixture of the previous two cases presented. Under these values, the agent has a prior belief that the system can either switch according to an exogenous hazard (the case with $\alpha = 0.9$), or is fixed in regime and will not switch. Given the non-observability of the regime indicator, the model allows for the possibility that the agent is unsure as to the present regime, but to be consistent with earlier results, we assume $\pi_{I=0,\alpha=1,0} = \pi_{I=0,\alpha=0.9,0} = 0.5$ for the simulation presented here. As such, the agent begins with prior $\Pr(I_0 = 0) = 1$, but is uniformly unsure about the hazard parameter value.

Figure 4 presents the results, once again assuming $\varepsilon_{t+1} = 0\ \forall t$ and a regime switch in *t*=26. As in the previous case, initial extraction is greater than in the regime 1 certainty case, but is not steadily decreasing as the simulation progresses. The reason can be seen in the evolution of beliefs prior to the regime shift. First, note that $\pi_{I=1,\alpha=1,0}$ is zero for the entire planning horizon given the structure of the updating equations, and thus $\Pr(I_t = 1)$ is completely characterized by one variable, $\pi_{I=1,\alpha=0.9,t}$. As the simulation progresses, this probability first increases then decreases (after *t*=12) as no detectable regime shift is observed, while the marginal belief that $I_t = 0$ follows a symmetric, opposite pattern. However, this occurs while the probability of no possible shift is monotonically increasing and $\Pr(I_t = 0\ \&\ \alpha = 0.9)$ is monotonically decreasing. The introduction of the parametric uncertainty, then, introduces the possibility of non-monotonic beliefs even absent stochastic shocks to the physical system, as the cumulative effects of no apparent shift tend to counteract the predictive distributions tendency to increase the probability that $I_t = 1$.

As such, the resource stock path is first gradually decreasing before the maximum $\Pr(I_t = 1)$ for the reasons discussed earlier, then gradually increasing thereafter towards the certainty equivalent solution to the regime 1 certainty case. As such, the incentive is to be more exploitative in the beginning of the horizon, consistent with the known, exogenous hazard case, but this incentive is mitigated after a point where the pattern of data suggests a reasonably high probability that a regime shift cannot occur, thus creating an incentive structure consistent with the competing ecosystem model case.

Once the data pattern is disrupted by the actual regime shift in *t*=26, belief patterns rapidly reflect the change, with $\pi_{I=1,\alpha=1,t}$ rapidly declining and $\pi_{I=1,\alpha=0.9,t}$ rapidly increasing. The resultant stock path is thus gradually decreasing to the regime 2 certainty steady-state level, in accordance with the previous analysis. The one difference is that since $\Pr(I_t = 1)$ is small at the time of shift, optimal extraction *increases*, rather than decreases, as the stock approaches the steady state.

The Effects of Information Structure

Taken together, these three cases illustrate that the assumptions over the information regime are non-trivial when it comes to characterizing the nature of optimal management in the presence of irreversible threshold effects. In general, risk over the true data generation regime results in lower stock levels, ceteris paribus, than the certainty regime solutions when the belief in regime 1 is strong, but the overall path of resource exploitation depends on the evolution of the subjective beliefs of the decision maker. These effects generally take two forms: 1) the tradeoff between the risk of over- and under- exploitation given $0 < \Pr(I_t = 1) < 1$, as characterized by the fact that $\frac{\partial V(\cdot)}{\partial \Pr(I_{t+1} = 1)} < 0$; and 2) the intertemporal tradeoffs in the belief dimension when past-period uncertainty is not immediately resolved, as characterized by the fact that $\frac{\partial \Pr(I_{t+1} = 1)}{\partial h_t} \neq 0$. In the case of the competing ecosystems model, the effect in 1) is paramount and the effect of 2) is minimized; in the case of even an exogenous probability of regime shift, both effects are in play. Note that 1) would hold in the case of exogenous arrival of information as well (e.g., information regimes such as that assumed in Melkonyan, 2011).

Discussion and Conclusions

In this paper, a renewable resources model with thresholds and irreversibility was analyzed under conditions where the decision-maker is not certain of the regime state of the system, due to either a stochastic data generation process or the lack of observability of the switching event. The probability of crossing the threshold is assumed to be exogenous, and beliefs over the unknown state are modeled in a state-space Bayesian framework, rendering the problem as a partially-observable Markov decision model. The context extends the analysis of Polasky et al. (2011) to a more generalized treatment of the information available to the optimizer at the time of the extraction decision. In addition, we show that if the probability of a regime shift/hazard rate is zero, then the specification presented here can be interpreted as a competing ecosystem model as in Peterson, et al. (2003).

Polasky et al. (2011) show that when past regime uncertainty at time *t*+1 is assumed, an exogenous probability of regime shift coupled with a systems dynamic effect (a change in regime with no discontinuity in resource stocks) results in no change in optimal management from the regime certainty case, but decreased exploitation when the probability is stock-dependent. Our results showed that the former can be explained by the lack of a tradeoff about being the agent being "wrong" in his/her beliefs; once the stochastic crossing event either happens or does not, the decision-maker knows which regime is prevailing, and thus there is no trade-off to make. The latter occurs because, ceteris paribus, there is always a positive expected economic benefit to reducing the hazard rate in accordance with the "Precautionary Principle" when the second, absorbing regime has a lower steady-state stock level (and thus steady-state current value) than the regime 1 certainty steady-state level. In the case of stock effect (a discontinuity in the expected stock), there is an additional incentive to exploit before the shift, as a portion of the stock can be destroyed and thus is not recoverable, even when regime uncertainty is resolved.

Similarly, Carpenter et al. (1999) and Ludwig et al. (2003) argue that precaution is always optimal in the presence of threshold uncertainty alone in a shallow-lake-type model with reversibilities, though Brozovic and Schlenker (2011) demonstrate that this result critically depends on a deterministic data generation process, and that precaution may be dominated by a countervailing effect when the uncertainty regarding the threshold location is very high. As explained in that paper, the intuition is that as the manager has “less control” of the system from a probabilistic standpoint, the marginal net benefits of lower pollution loadings decreases.[17] In effect, this is the same incentive facing the manager with an exogenous hazard in the renewable resource model with observability; in each case, a lack of ability to meaningfully trade-off the potentially poor outcome of a regime shift in the future with current actions dissipates the precautionary incentive.

The results presented here showed that these conclusions are critically dependent on the information regime assumed; namely, that the problem is essentially static in the belief dimension. If instead beliefs evolve in accordance with a (possibly endogenous) updating rule, then optimal management in the presence of an exogenous regime shift is no longer identical to the optimal rules under certainty of regime. Rather, it was shown that stock levels are kept *lower* than the regime-certainty case, and that they tend to decrease over time until they reach the limiting (lower-stock) solution of the lower-growth regime. As such, under persistent regime risk, optimal management is *not* consistent with precautionary behavior, but rather greater levels of exploitation and lower stock levels during the period before a true (but unknown) regime shift.

The intuition is that the learning essentially endogenizes the tradeoffs between being right or wrong about the true regime, in conjunction with the dynamic incentive created as a result of the updating process. In the cases presented here, the former will always result in non-precautionary behavior given the shape of the value function, while the latter has an ambiguous effect depending on the information regime assumed and the state of the system. This ambiguity is demonstrated by the effects of the decision-maker’s beliefs about regime 2 when the hazard rate is unknown, which can be non-monotonic even in the state-transition certainty-equivalent case of $\varepsilon_{t+1} = 0 \ \forall t$. As such, the initial exploitive incentive that develops can be mitigated as the belief in no possible shift increases, similar to the result in Brozovic and Schlenker (2011), though a complete reversal does not appear possible.[18]

From a policy standpoint, this suggests a countervailing effect to the unambiguous precaution result in Polasky et al. (2011) when the probability of regime shift is endogenous and there is no expected stock collapse at the time of the shift. In other words, if as society we “know we don’t know” and operate accordingly, the relative magnitudes of the endogenous risk effect and the uncertain regime effect will determine if a reduction in emissions is justified, even if there is only a system dynamics effect. This same effect would even further induce increased exploitation in the case of an

[17] In these models, the manager is assumed to admit preferences that are strictly concave in the pollution stock, thus strengthening the risk-aversion type result. In addition, it should be noted that the control action in these models affects the probabilities of crossing/re-crossing the threshold, and thus would fall into the class of endogenous hazard models.

[18] Differences in model structure explain the difference in sign on the risk effect vs. the information effect. The main point is that the information effect is an additional margin working against the deterministic solution.

exogenous probability of regime shift with a stock effect, and for much the same reason; namely, the process governing beliefs tends to lower future ex ante expected values of the system, increasing the incentive to extract sooner. This incentive is similar to an increase in the discount rate, and seems especially applicable to the argument surrounding global climate change given the uncertainty surrounding those processes.

A few notes about assumptions and possible extensions are in order. In this paper, we assumed irreversibility over the regime switch and a risk-neutral decision-maker, as well as well-defined probability distributions over the stochastic elements of the model. Future research is needed to more fully explore the implications of these assumptions, as they will undoubtedly further muddle the policy prescriptions that follow from the analysis. In fact, in terms of climate change, well-defined probability distributions may be especially problematic, and models which incorporate ambiguity/second-order probabilities may be more appropriate (see, e.g., Kilbanoff et al., 2009 and Millner, et al., 2010).

In addition, our characterization of the information regime was itself grossly simplified, with a maximum of two possible hazard rates and two possible regimes. While this renders the problem tractable, it ignores additional information issues such as the effect of a declining variance of an unknown parameter over a continuous prior. As model complexity increases, however, both the analytical and numerical techniques to analyze them become more non-standard, and the curse of dimensionality becomes a very real limiting factor. Nevertheless, given the demonstrated importance of the structure of information for ex ante decision-making under uncertainty, risk, and learning in environmental and resource problems, this line of research should be extended.

References

C.A. Bond, On the potential use of adaptive control methods for improving adaptive natural resource management, Optimal Control Applications and Methods 31 (2010) 55-66.

C.A. Bond, J.B. Loomis, Using numerical dynamic programming to compare passive and active learning in the adaptive management of shallow lakes, Canadian Journal of Agirultural Economics 57(4) (2009) 555-573.

W.A. Brock, S.R. Carpenter, Panaceas and diversification of environmental policy, Proceedings of the National. Acadamy of Sciences 104(39) (2007) 15206-15211.

N. Brozovic, W. Schlenker, Optimal management of an ecosystem with an unknown threshold. Ecological Economics 70 (2011) 627-640.

S.R. Carpenter, D. Ludwig, W.A. Brock, Management of eutrophication for lakes subject to potentially irreversible change, Ecological Applications 9(3) (1999) 751-771.

H.R. Clarke, W.J. Reed, Consumption/pollution tradeoffs in an environment vulnerable to pollution-related catastrophic collapse, Journal of Economic Dynamics and Control 18(5) (1994) 991-1010.

M.A. Cunha-e-Sá, V. Santos, Experimentation with accumulation, Jouranl of Economic Dynamics and Control 32 (2008) 470-496.

R.G. Haight, S. Polasky, Optimal control of an invasive species with imperfect information about the level of infestation, Resource and Energy Economics 32 (2010) 519-533.

J. Han, T.L. Lai, V. Spivakovsky, Approximate policy optimization and adaptive control in regression models, Computational Economics 37 (2006) 433-452.

K.L. Judd. Numerical Methods in Economics, The MIT Press, Cambridge and London, 1998.

R.E. Kalman, A new approach to linear filtering and prediction problems, Trans. ASME, Journal of Basic Engineering 82: (1960) 34-45.

J.D. Kaplan, R.E. Howitt, Y.H. Farzin. An information-theoretical analysis of budget-constrained nonpoint source pollution control, Journal of Environmental Economics and Managemetn 46 (2003) 106-130.

D.L. Kelly, C.D. Kolstad, Bayesian learning, growth, and pollution, Journal of Economic Dynamics and Control 23 (1999) 491-518.

D.A. Kendrick, Stochastic control for economic models: Past, present, and the paths ahead, Journal of Economic Dynamics and Control 29 (2005) 3-30.

Klibanoff, P., M. Marinacci, S. Mukerji, Recursive smooth ambiguity preferences, Journal of Economic Theory 144 (2009) 930-976.

D.A. Ludwig, C.J. Walters, Optimal harvesting with imprecise parameter estimates, Ecological Modelling 14 (1982) 273-292.

D. Ludwig, S.R. Carpenter, W.A. Brock, Optimal phosphorus loading for a potentially eutrophic lake, Ecological Applications 13(4) (2003) 1135-1152.

T. Melkonyan, The effect of communicating ambiguous risk information on choice. Journal of Agricultural and Resource Economics 36(2) (2011). 292-312.

J.R. Miller, F. Lad, Flexibility, learning and irreversibility in environmental decisions: A Bayesian approach, Journal of Environmental Economics and Management 46 (1984) 161-172.

A. Millner, S. Dietz, G. Heal, Ambiguity and climate policy. Centre for Climate Change Economics and Policy Working Paper No. 28, Munich Re Programme Technical Paper No. 4, and Grantham Research Institute on Climate Change and the Environment Working Paper No. 24 (2010).

M.J. Miranda , P.L. Fackler, Applied Computational Economics and Finance, The MIT Press, Cambridge and London, 2002.

E. Nævdal, Optimal regulation of natural resources in the presence of irreversible threshold effects, Natural Resource Modeling 16(3) (2003) 305-333.

G.D. Peterson, S.R. Carpenter, W.A. Brock, Uncertainty and the management of multistate ecosystems: An apparently rational route to collapse, Ecology 84(6) (2003) 1403-1411.

S. Polasky, A. deZeeuw, F. Wagener, Optimal Management with potential regime shifts, Journal of Environmental Economics and Managmenet 62(2) (2011) 229-240.

M. Springborn, Bayesian adaptive management with optimal learning, Working paper, presented at Colorado State University, Fort Collins, CO., January 28, 2008.

W.J. Reed, Optimal harvesting of a fishery subject to random catastrophic collapse, IMA Journal of Mathematics Applied in Medicine and Biology 5(3) (1998) 215-235.

J. Rust, Using randomization to break the curse of dimensionality, Econometrica 65 (1997) 487-516.

S. Särkkä, Recursive Bayesian Inference on Stochastic Differential Equations, Ph.D. Dissertation, Helsinki University of Technology, Espoo, Finland, ISBN 951-22-8127-9, 2006, accessed April 4, 2011 at http://lib.tkk.fi/Diss/2006/isbn9512281279/.
M. Scheffer, S. Carpenter, J.A. Foley, C. Folke, and B. Walker, Catastrophic shifts in ecosystems, Nature 413 (2001) 591-596.

A. Thrower, J.M. Martinez, Reconciling anthropocentrism and biocentrism through adaptive management, J. of Environment and Development 9(10) (2000) 68-97.

Y. Tsur and A. Zemel, Accounting for global warming risks: Resource management under event uncertainty, Journal of Economic Dynamics and Control 20 (6/7) (1996) 1289-1305.

C.J. Walters and C.S. Holling, Large-scale management experiments and learning by doing, Ecology 71(6) (1990) 2060-2068.

V. Wieland, Learning by doing and the value of optimal experimentation, Journal of Economic Dynamics and Control 24 (2000) 501-534.

Appendix A

Equivalence of Specifications with Observability

To show that this model is a generalization of the models in Polasky et al. (2011), consider the (relatively) simple case of a deterministic resource state-transition equation with exogenous, known hazard rate $0<(1-\alpha)<1$. This is a two-state, one control, infinite horizon, discrete time stochastic dynamic optimization problem of the form

$$\begin{aligned}
&\max_{h_1 \dots h_\infty} \sum_{t=0}^{\infty} \delta^{t} p h_t \\
&s.t.\ s_{t+1} = \begin{cases} s_t + G_1(s_t) - h_t \text{ if } I_{t+1} = 0 \\ s_t + G_2(s_t) - h_t \text{ if } I_{t+1} = 1 \end{cases} \\
&I_{t+1} = \begin{cases} 0 \text{ if } I_t = 0 \ \& \ \zeta_{t+1} < \alpha \\ 1 \ otherwise \end{cases}, \\
&s_0 > 0,\ I_0 \in (0,1)
\end{aligned} \tag{13}$$

where δ^t is the discount factor for period *t*. In this problem, the regime switches exactly once, this switch is irreversible, and the exogenous probability of switching in each period conditional on not having switched previously is $1-\alpha$. As such, it is the discrete-time analog to the exogenous regime shift with system dynamics in Polasky, et al. (2011). The discrete-time Bellman equation takes the form

$$\begin{aligned}
V(s_t,0) &= ph_t + \delta\left[\alpha V\left(s_t + G_1(s_t) - h_t, 0\right) + (1-\alpha)V\left(s_t + G_2(s_t) - h_t, 1\right)\right] \\
V(s_t,1) &= ph_t + \delta V\left(s_t + G_2(s_t) - h_t, 1\right),
\end{aligned} \tag{14}$$

and is standard in that it does not use the state-space framework.

To show that the state-space specification nests this case, note that if $\pi_{0,I} = 0$, the Bellman equation becomes

$$\begin{aligned}
V(s_0,0) &= ph_t + \delta E\left[V\left(s_t + G_1(s_t) - h_t, \pi_{t+1,I}\right) \mid I_{t+1} = 0\right] \\
&= ph_t + \delta\left[\alpha V\left(s_t + G_1(s_t) - h_t, \pi_{t+1,I}(s_t, h_t \mid \pi_{t,I}, G_1)\right)\right. \\
&\quad \left. + (1-\alpha)V\left(s_t + G_2(s_t) - h_t, \pi_{t+1,I}(s_t, h_t \mid \pi_{t,I}, G_2)\right)\right].
\end{aligned} \tag{15}$$

Since there is no error term on the state equations, the filtering distribution (8) reduces to a spike at zero or one given a new piece of data s_{t+1}, resulting in perfect prediction of I_{t+1}. We can thus rewrite (15) as

$$V(s_t,0) = ph_t + \delta\left[\alpha V\left(s_t + G_1(s_t) - h_t, 0\right) + (1-\alpha)V\left(s_t + G_2(s_t) - h_t, 1\right)\right], \tag{16}$$

which holds not only for *t*=0 but also any *t* for which $\pi_{t,I} = 0$. Similarly,

$$V(s_t,1) = ph_t + \delta V\left(s_t + G_2(s_t) - h_t, 1\right). \tag{17}$$

Equations (16) and (17) are equivalent to (14), thus demonstrating that the state-space representation generalizes the irreversibility formulations used in Polasky et al. (2011).

Equivalence of Control Rules of Exogenous, Known Hazard Model with Observability and Regime Certainty Case assuming Systems Dynamics Effects

We now show that the optimal management of this problem is equivalent to the optimal management in each regime, assuming that regime persists at all times in the future.

The necessary conditions of the problem, assuming either an exogenous or endogenous hazard rate, can be derived through differentiation of

$$V\left(s_t,\pi_{t,I}\right)=\max_{h_t} ph_t+\delta\left\{\left[1-\pi_{t,I}\right]\left(\alpha(s)EV^1(\cdot)+(1-\alpha(s))EV^2(\cdot)\right)+\pi_{t,I}EV^2(\cdot)\right\}$$

with respect to the control and each state variable, with notation $EV^i(\cdot)=EV\left(s_t+G_i(s_t)+h_t+\varepsilon_{t+1},i-1\right)$. Assuming no information effects (i.e., $\frac{\partial\pi_{t+1,I}}{\partial h_t}=\frac{\partial\pi_{t+1,I}}{\partial s_t}=\frac{\partial\pi_{t+1,I}}{\partial\pi_{t,I}}=0$), which is the case for observable regime shifts, these conditions are

$$p+\delta\left\{\left[1-\pi_{t,I}\right]\left(\alpha(s_t)E\left(-\frac{\partial V^1(\cdot)}{\partial s_t}\right)+(1-\alpha(s_t))E\left(-\frac{\partial V^2(\cdot)}{\partial s_t}\right)\right)+\pi_{t,I}E\left(-\frac{\partial V^2(\cdot)}{\partial s_t}\right)\right\}\overset{set}{=}0$$

$$\frac{\partial V\left(s_t,\pi_{t,I}\right)}{\partial s_t}\overset{set}{=}$$

$$\delta\left\{\left[1-\pi_{t,I}\right]\left(\alpha'(s_t)EV^1(\cdot)+\alpha(s_t)E\left(\frac{\partial V^1(\cdot)}{\partial s_t}\left(1+G_1'(s_t)\right)\right)-\alpha'(s_t)EV^2(\cdot)+(1-\alpha(s_t))E\left(\frac{\partial V^2(\cdot)}{\partial s_t}\left(1+G_2'(s_t)\right)\right)\right)\right.$$

$$\left.+\pi_{t,I}E\left(\frac{\partial V^2(\cdot)}{\partial s_t}\left(1+G_2'(s_t)\right)\right)\right\}$$

$$\frac{\partial V\left(s_t,\pi_{t,I}\right)}{\partial\pi_{t,I}}\overset{set}{=}\delta\left\{-\left(\alpha(s_t)EV^1(\cdot)+(1-\alpha(s_t))EV^2(\cdot)\right)+EV^2(\cdot)\right\}$$

(18)

Note that $\frac{\partial V\left(s_t,\pi_{t,I}\right)}{\partial\pi_{t,I}}\approx V\left(s_t,0\right)-V\left(s_t,1\right)$, and the third equation in (18) implies that

$$EV^2(\cdot)=E\left\{V\left(s_{t+1},1\right)\right\}\approx\frac{V\left(s_t,0\right)-V\left(s_t,1\right)}{\delta\alpha}+EV\left(s_t+G_1(s_t)+h_t+\varepsilon_{t+1},0\right).\qquad(19)$$

Differentiation and substitution of (19) into the first two equations of (18) results in the following conditions:

$$p+\delta\left\{\left[1-\pi_{t,I}\right]\left(\alpha(s_t)E\left(-\frac{\partial V^1(\cdot)}{\partial s_t}\right)+(1-\alpha(s_t))E\left(-\frac{\partial V^1(\cdot)}{\partial s_t}\right)\right)+\pi_{t,I}E\left(-\frac{\partial V^1(\cdot)}{\partial s_t}\right)\right\}=0$$

$$\frac{\partial V\left(s_t,\pi_{t,I}\right)}{\partial s_t}=$$

$$\delta\left\{\left[1-\pi_{t,I}\right]\left(\alpha'(s_t)EV^1(\cdot)+\alpha(s_t)E\left(\frac{\partial V^1(\cdot)}{\partial s_t}\left(1+G_1'(s_t)\right)\right)\right.\right.$$

$$\left.\left.-\alpha'(s_t)EV^2(\cdot)+(1-\alpha(s_t))\left(\frac{\partial V\left(s_t,0\right)/\partial s_t-\partial V\left(s_t,1\right)/\partial s_t}{\delta\alpha}+E\left(\frac{\partial V^1(\cdot)}{\partial s_t}\left(1+G_1'(s_t)\right)\right)\right)\right)\right.$$

$$\left.+\pi_{t,I}E\left(\frac{\partial V\left(s_t,0\right)/\partial s_t-\partial V\left(s_t,1\right)/\partial s_t}{\delta\alpha}+E\left(\frac{\partial V^1(\cdot)}{\partial s_t}\left(1+G_1'(s_t)\right)\right)\right)\right\}$$

Thus, so long as the marginal value of the renewable state in each regime is identical, and $\alpha'(s)=0,$ the problem reduces to managing the system according to the most rapid approach path for each regime. This is the result shown in Polasky et al. (2011), but critically depends on an unrestricted harvest when switching from one regime to the other in the case of linear harvest.

Decreased Exploitation assuming Systems Dynamics Effects and Endogenous, Known Hazard Model with Observability

With an endogenous hazard, assuming a prior reflecting certainty over regime 1, the first-order conditions become:

$$\begin{aligned}0&=p-\beta\left\{\alpha(s)V_s(s+G_1(s)-h,0)+(1-\alpha(s))V_s(s+G_1(s)-h,0)\right\}\\&=p-\beta\left\{V_s(s+G_1(s)-h,0)\right\}\end{aligned}\tag{20}$$

$$V_s(s,0)=\beta\left[\alpha'(s)\left[V(s+G_1(s)-h,0)-V(s+G_2(s)-h,1)\right]+V_s(s+G_1(s)-h,0)(1+G_1'(s))\right]\tag{21}$$

At the new steady state for this regime, from (20), $V_s(s+G_1(s)-h,0)=p/\beta,$ if harvest is unconstrained, and substitution into (21) yields $p/\beta=\beta\left[\alpha'(s)\left[V(s+G_1(s)-h,0)-V(s+G_2(s)-h,1)\right]+\left(p/\beta\right)(1+G_1'(s))\right].$ Moving the non-bracketed term over yields $p/\beta-p(1+G_1'(s))=\beta\left[\alpha'(s)\left[V(s+G_1(s)-h,0)-V(s+G_2(s)-h,1)\right]\right].$ Solving, $G_1'(s)=\dfrac{1-\beta}{\beta}-\dfrac{\beta}{p}\left[V(s+G_1(s)-h,0)-V(s+G_2(s)-h,1)\right]$, so that marginal growth is lower than in the original case, and thus stock levels are greater.

Increased Exploitation assuming Stock Effect and Exogenous, Known Hazard Model with Observability

Similarly, for a stock regime shift where $s_{t+1}=0$ if $I_{t+1}=1,$ the Bellman equation can be written as

$$V(s,0) = ph + \beta\alpha(s)E\left[V(s+G(s)-h,0)\right],$$

with first order conditions

$$0 = p - \beta\alpha(s)E\left[V_s(s+G(s)-h,0)\right] \tag{22}$$

$$V_s(s,0) = \alpha'(s)E\left[V(s+G(s)-h,0)\right] + \alpha(s)E\left[V_s(s+G(s)-h,0)(1+G'(s))\right]. \tag{23}$$

From (23), $E\left[V_s(s+G(s)-h,0)\right] = p/\left(\beta\alpha(s)\right)$ if harvest is unconstrained, and substitution into (23) yields
$V_s(s,0) = \beta\left[\alpha'(s)E\left[V(s+G(s)-h,0)\right] + \alpha(s)\left(p/\beta\alpha(s)\right)(1+G'(s))\right]$.

At the steady state,
$p/\beta\alpha(s) = \beta\left[\alpha'(s)E\left[V(s+G(s)-h,0)\right] + \alpha(s)\left(p/\beta\alpha(s)\right)(1+G'(s))\right]$. Note that if α is constant and doesn't depend on s, then $1/\beta = \left[\alpha(1+G'(s))\right]$, or
$1/\left(\alpha\beta\right) = \left[(1+G'(s))\right] \rightarrow G'(s) = \frac{1-\alpha\beta}{\alpha\beta}$. Furthermore, $\frac{1-\alpha\beta}{\alpha\beta} > \frac{1-\beta}{\beta}$ so long as $\alpha < 1$.
As such, $G'(s)$ is positive and greater in magnitude than the case where there is no hazard, suggesting that the steady state stock level is lower and exploitation is higher. The intuition, as in Polasky et al. (2011), is that the hazard term increases the discount rate.

Ambiguous Exploitation assuming Stock Effect and Endogenous, Known Hazard Model with Observability

In the case of an endogenous shift (actually, the more general case which naturally reduces to above when $\alpha'(s)$ =0), $\frac{p/\beta - \beta\alpha'(s)E\left[V(\cdot)\right] - \alpha(s)p}{\alpha(s)p} = G'(s)$, which is ambiguous compared to $\frac{1-\beta}{\beta}$. This term can be rewritten as
$\frac{1-\alpha(s)\beta}{\alpha(s)\beta} - \frac{\beta\alpha'(s)E\left[V(\cdot)\right]}{\alpha(s)p} = G'(s)$. The closer $\alpha'(s)$ is to zero, ceteris paribus, the more likely that the LHS is greater than $\frac{1-\beta}{\beta}$, and thus the exploitation incentive is greater.

Note as well that steady state value is $V(s,0) = \frac{pG(s)}{1-\beta\alpha(s)}$.

Table I: Parameter Values used in Numerical Analysis

Parameter	***Description***	***Value(s)***
δ	One-period discount factor	0.95
p	Price of resource stock	1
α	Probability of No Regime Shift (1-hazard)	(1,0.9)
	Mean of Error on State Transition	0
	Std. Dev. of Error on State Transition	0.0250
k_1	Carrying Capacity in Regime 1 [$G_1(k_1)=0$]	1
k_2	Carrying Capacity in Regime 2 [$G_2(k_2)=0$]	0.7
	Number of Grid Points in Each State Dimension	11
	Number of Grid Points in Control Dimension	201
	Number of Quadrature Points for Integration	5
	State Bounds on Physical State	(0,1)

Note that $G_1(k_2)>0$ and $G_2(k_1)<0$.

Figure 1: Optimal Value Functions for the Competing Ecosystems Model (Panel I) and Exogenous, Known Hazard Model without Observability (Panel II)

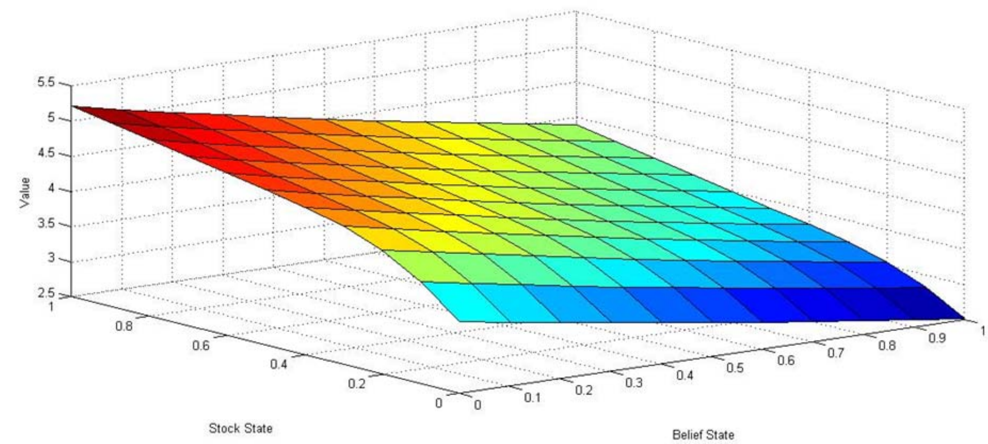

Panel I: Competing Ecosystems Model

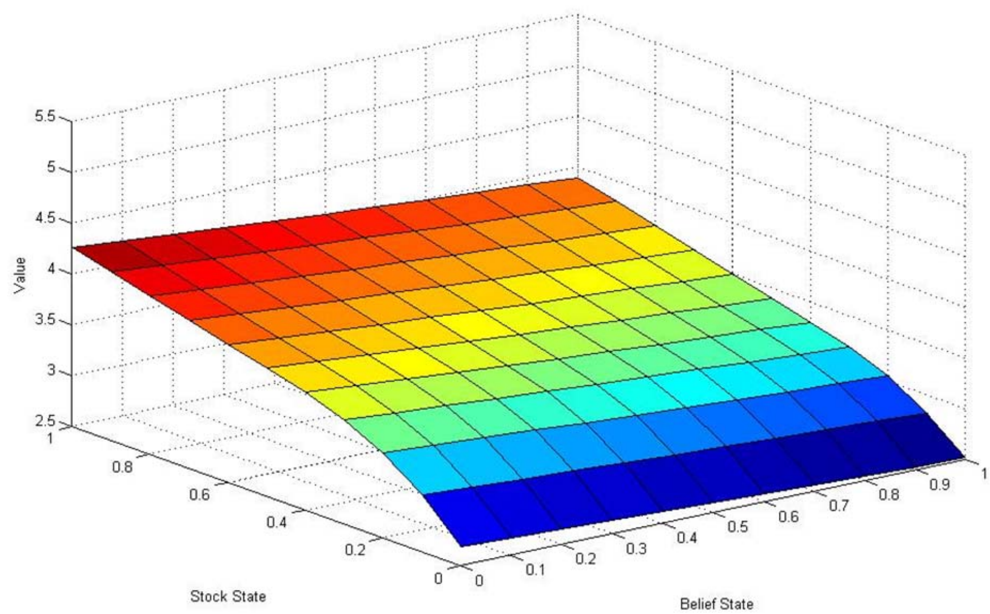

Panel II: Exogenous, Known Hazard Model without Observability

Figure 2: Simulated Optimal Paths for Competing Ecosystems Model, Error Terms on Transition Set Equal to Zero

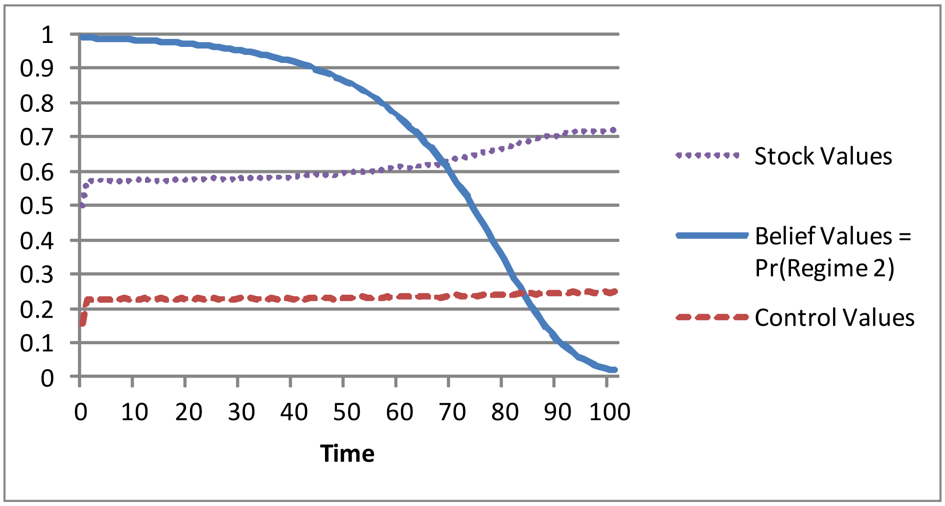

Panel I: True Growth Regime = Regime 1

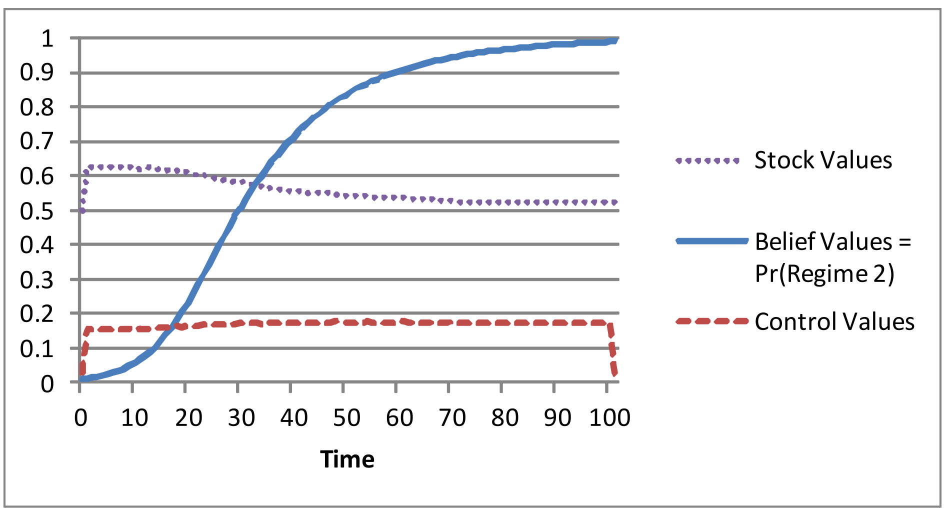

Panel II: True Growth Regime = Regime 2

Figure 3: Simulated Optimal Paths for Regime Shift Model with Known, Exogenous Hazard = 0.1, Error Terms on Transition Set Equal to Zero

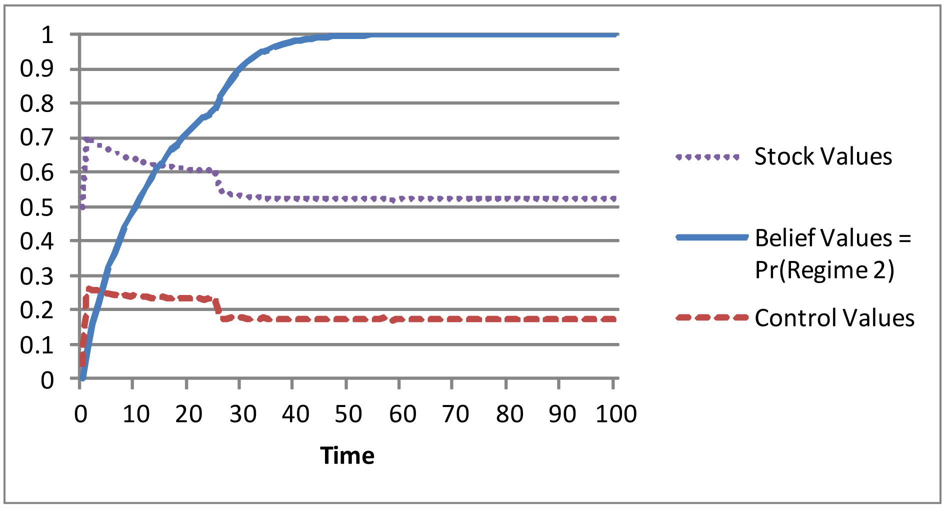

True Growth Regime = 1 for t=1..26

Figure 4: Simulated Optimal Paths for Regime Shift Model with Unknown, Exogenous Hazard = (0,0.1), Error Terms on Transition Set Equal to Zero

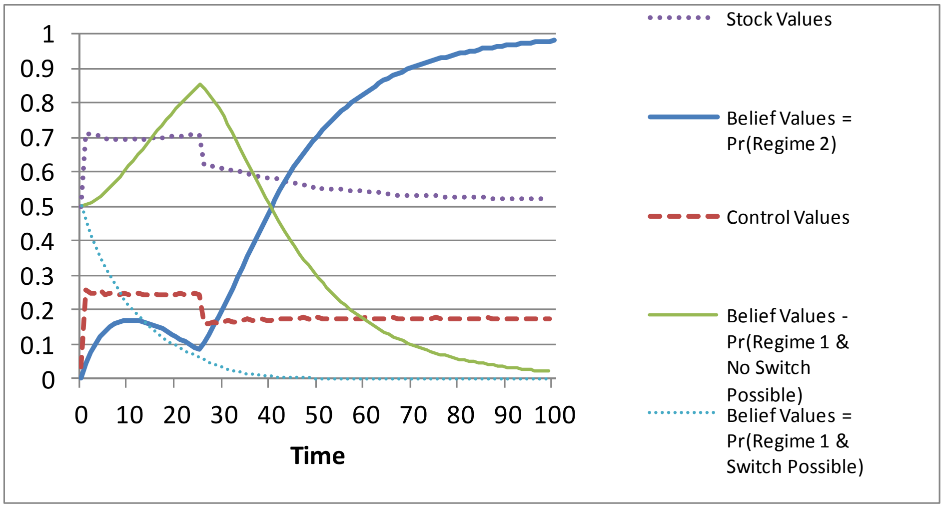

True Growth Regime = 1 for t=1..26