The Treatment of Uncertainty and Learning in the Economics of Natural Resource and Environmental Management

Jacob LaRiviere*, David Kling†, James N. Sanchirico‡, Charles Sims§, and Michael Springborn||

Introduction

Uncertainty permeates the economics of natural resource and environmental management. Well-known examples include the uncertain damages from pollution, asymmetric information regarding firms' abatement costs, random variation in resource growth (i.e., stochasticity), imperfectly observed resource stock abundance, and uncertainty regarding the functional relationship that determines how quickly pollution leaves a system (i.e., decay rates). The economic significance of uncertainty in most environmental and resource management problems raises two important questions for researchers: (1) what types of uncertainty are relevant for the problem being studied and (2) how transferable are insights from studying one type of environmental problem and one type of uncertainty to other settings

*Microsoft, Building 112/4308, 14865 NE 36th Street, Redmond, WA 98052. Tel: 425-703-3085; University of Tennessee, Knoxville, 507 Stokely Management Center, Knoxville, Tennessee 37996-0550. Tel: 865-974-3303; email: jlariv@microsoft.com.
†Oregon State University, Ballard Extension Hall 240F, 2591 SW Campus Way, Corvallis, Oregon 97331. Tel: 541-737-1418; email: david.kling@oregonstate.edu.
‡University of California, Davis, 1 Shields Ave., Davis, California 95616. Tel: 530-754-9883; email: jsanchirico@ucdavis.edu.
§University of Tennessee, Knoxville, Howard Baker Jr. Center for Public Policy, 1640 Cumberland Avem Knoxville, Tennessee 37996. Tel: 865-974-3825; email: cbsims@utk.edu.
||University of California, Davis, 1 Shields Ave, Davis, California 95616. Tel: 530-752-5244; email: mspringborn@ucdavis.edu.

This work was assisted through participation in the "Pretty Darn Good Control" Investigative Workshop at the National Institute for Mathematical and Biological Synthesis, sponsored by the National Science Foundation through NSF Award DBI-1300426, with additional support from the University of Tennessee, Knoxville. The authors acknowledge helpful discussions and input from working group co-organizers Megan Donahue, Paul Armsworth, and Carl Toews and participants Michael Bode, Carl Boettiger, Iadine Chades, and Alan Hastings, as well as comments from two anonymous reviewers and Suzanne Leonard, which strengthened this paper.

Review of Environmental Economics and Policy, volume 12, issue 1, Winter 2018, pp. 92–112
doi: 10.1093/reep/rex021
Advance Access Published on December 6, 2017
Published by Oxford University Press on behalf of the Association of Environmental and Resource Economists 2017.

(e.g., different environmental problems with similar uncertainty or similar environmental problems with different types of uncertainty)?

The current economics literature on resource and environmental management models several different classes of uncertainty, ways to reduce that uncertainty through learning, and technical solution methods to provide policy insights. This article seeks to help researchers and policymakers distinguish between these broad classes of uncertainty, learning, technical solution methods and the policy insights derived from them in this (sometimes confusing) literature. First, we introduce a simple model of optimal dynamic management to examine each type of uncertainty and learning using the model as a common framework. The framework takes the perspective of a social planner seeking to maximize the expected net present value of a generic environmental or resource stock.[1] Second, we define several forms of uncertainty in addition to stochasticity in the context of the model.[2] In doing so, we explicitly examine various ways regulators can learn about and reduce uncertainty. Third, we use the framework to evaluate four hypotheses associated with the policy implications of uncertainty and learning that have been discussed in the literature and policy circles: (1) greater levels of uncertainty result in more precautionary management (e.g., Gollier, Jullien, and Treich 2000), (2) reductions in environmental uncertainty have greater value than reductions in economic uncertainty (e.g., Sethi et al. 2005), (3) management strategies with active learning about an economic or ecological system lead to substantially higher welfare than those with passive learning (e.g., Bond and Loomis 2009), and (4) managers who try to actively learn do so faster than those who passively learn (e.g., Springborn and Sanchirico 2013). Our evaluation of these "hypotheses" is the primary contribution of this article. Most importantly, we find that while each of the hypotheses has some degree of validity, they either break down or, although true, have limited applicability across environmental or resource management contexts or different forms of uncertainty. We conclude with a discussion of promising opportunities for further research in this area.

Modeling Environmental Management Problems

Most resource or environmental management problems involve multiple types of uncertainty. To anchor our discussion and distinguish among different types of uncertainty, we first introduce a simple modeling framework of a management problem with a single decision maker whom we call the manager.[3]

The manager's goal is to maximize the expected net present value (NPV) of a stream of benefits (e.g., the discounted flow of rents from resource extraction or health improvements from pollution abatement). With this in mind, we will frame much of our discussion in the

[1]This simple model can be extended to examine common second-best management problems (e.g., profit maximization subject to minimum extinction risk or maximum allowable pollution levels). Due to space limitations we focus on models grounded in subjective expected utility theory (SEU), by far the dominant analytical framework in economics and allied quantitative disciplines (Shaw and Woodward 2008; Gilboa 2009).

[2]We follow the decomposition approach of Walters and Hilborn (1976) and Charles (1998). Fackler (2014) also notes these approaches and presents alternative ways to discuss uncertainty. For a more complete treatment of risk for policy, we refer readers to Morgan and Henrion (1990) and Pindyck (2007).

[3]Modeling a single decision maker is the traditional first step in analyzing dynamic environmental and resource economic problems (Clark 1990).

context of two classic renewable resource management problems: (1) maximizing the expected discounted flow of rents from a harvested stock of fish and (2) minimizing the expected discounted flow of damages and control costs from a stock (ambient concentration) of pollution like sulfur dioxide.

In the model, the manager decides in each period how much to manipulate a physical state variable (e.g., the stock of fish) by selecting the value of a control variable she dictates (e.g., total harvest or fishing effort in a year). For the pollution management problem, the control variable could be pollution emissions in a period, with the state variable being the stock of the pollutant.[4] The manager has a within-period reward function (e.g., industry fishing profits, monetized health benefits from pollution abatement) and future rewards are downweighted using a discount factor.

These models are made both interesting and complicated by how the state variable and control variable jointly determine benefits from management through the within-period reward function and the evolution of the state variable through the state equation. The transition of the state variable to the next period is a function of the current state and control variables. We assume that the next period's state is a function of parameters given by nature. For a fish harvest problem, these parameters could include the natural environment's carrying capacity and the intrinsic growth rate for a fish species. For a stock pollutant, these parameters would include the pollutant's decay rate. In general, the state equation is often nonlinear or features a spatial or multistate structure that adds considerable complexity to the manager's problem. For example, in the absence of fishing, a relatively large breeding stock of fish may grow more slowly than a relatively smaller breeding stock due to food abundance.

Types of Uncertainty

Next, in the context of an optimal dynamic management framework, we define four types of uncertainty common in management problems that have received significant attention in the literature: stochasticity, parametric uncertainty, model uncertainty, and state uncertainty.[5]

[4]Formally, the manager's problem is

$$\max_{\{u_t\}} E\left[\sum_{t=0}^{\infty} \delta^t \pi(x_t, u_t, \epsilon_t^{\pi}; \Theta_t^{\pi})\right] \tag{1}$$

subject to

$$x_{t+1} = f\left(x_t, u_t, \epsilon_t^{f}; \Theta_t^{f}\right),$$

where t is the time period, x_t is the physical state variable, u_t is the control variable, and $\boldsymbol{\epsilon}_t = (\epsilon_t^{\pi}, \epsilon_t^{f})$ is a vector of stochastic terms that operate on the objective function (e.g., stochastic output prices) or the state equation (e.g., stochastic shocks to the resource stock). The variable $\boldsymbol{\Theta}_t = (\Theta_t^{\pi}, \Theta_t^{f})$ is a set of economic and biophysical parameters, $\pi(\,)$ is the within-period reward benefit, and δ is the discount factor. When the manager's problem involves minimization, it may be converted to a mathematically equivalent maximization problem after making the sign of $\pi(\,)$ negative.

[5]Indeed, all four types of uncertainty may be simultaneously present in a single management problem.

Stochasticity

Natural resource and environmental management often deals with dynamic processes that are influenced by stochasticity, which is often defined as intrinsic variability, shocks, risk, or "noise."[6] In our model, stochasticity is represented by a random vector, and both within-period outcomes (e.g., economic rewards change with stochastic output or input prices) and state dynamics (e.g., excess wind leading to a lot of pollution dissipation) can be stochastic. Stochastic terms are typically drawn from a known distribution (Reed 1979).[7] The parameters dictating the stochastic component are generally fixed over time, but more complex time-dependent relationships are possible.[8]

Even within a short time scale, it is rarely technically feasible (much less cost effective) to *perfectly* forecast state variables because sources of natural variability (e.g., ocean upwelling anomalies, droughts) are difficult to predict. Moreover, there are a variety of sources of economic variability, such as resource commodity prices (Pindyck 1984). To capture natural and economic variability and account for the challenge of prediction, researchers often include *irreducible* stochasticity in the models (i.e., it cannot be resolved through learning prior to a decision period). When uncertainty is limited to stochasticity in a resource or environmental management model, the modeler implicitly assumes the manager knows all functions/distributions of stochastic variables and the manager plans as if these functional relationships precisely describe the dynamics of the management problem (Shaw and Woodward 2008).[9]

In stochastic models, managers may still have an incentive to obtain new information because of option value (Arrow and Fisher 1974; Dixit and Pindyck 1994). The concept of option value establishes that merely accounting for new information about the outcome of an uncertain stochastic shock tends to increase the value of delaying an irreversible action (Pindyck 2007). Irreversibility may represent harvesting below the minimum viable population of a renewable resource at which the resource goes extinct or incurring sunk costs to enact a policy that would preserve the resource (Sims et al. 2017). When management actions today limit the manager's options for action in the future, there is an opportunity cost of irreversible deviations from the status quo, through either *forgone information value* or *the forfeiture of the option* to react to future realizations of the random variable (Conrad 1980). In a sense, option value is the value of having a less restrictive control space—that is, more flexibility in management choices in the face of an uncertain future. Hence resolution of stochasticity is valuable, but we do not view it as learning.

[6]Unless stated otherwise, we generally consider "greater levels of uncertainty" to be a mean preserving spread in the stochastic process.

[7]This is specified in our model by the parameters in $\boldsymbol{\Theta}_t$.

[8]In keeping with standard practice in the literature, we restrict our discussion to stochastic models that are Markov processes, which means that conditional on the current state and the manager's chosen control, next period's state is independent of states and control choices that occurred prior to the current period. In addition, the appropriate distributional assumption for the stochastic component is often context dependent. For example, a lognormal distribution may be appropriate if stochasticity is to be modeled multiplicatively.

[9]In SEU theory, a decision maker assigns probability distributions to stochastic events even when she does not know the distribution precisely and then plans as if those distributions are correct and known (Shaw and Woodward 2008). To simplify our discussion and draw a clear distinction between stochasticity and other forms of uncertainty, we focus on stochasticity characterized by distributions that the manager either knows or assumes but does not learn about over time.

Parametric Uncertainty

Thus far we have assumed that although the reward function or state equation is stochastic, the manager knows the parameters of each with certainty, up to a multiplicative or additive stochastic component. Under parameter (or parametric) uncertainty, the manager plans without precise knowledge of one or more components of the system. For example, a fishery manager may specify regulations like catch limits for a stock without knowing the stock's precise intrinsic growth rate (Springborn and Sanchirico 2013), or an environmental manager may be uncertain about pollution transfer coefficients (Fowlie and Mueller 2013; Carson and LaRiviere 2017). Unlike stochasticity, parametric uncertainty (along with other types of uncertainty) is *reducible* through learning, at least conceptually. However, if such learning is not feasible or simply not pursued, then parameters are typically modeled as random draws. In either case, parametric uncertainty could lead to a less "predictable" system when combined with a purely stochastic state equation.

Model Uncertainty

Model uncertainty differs from parametric uncertainty in that the manager does not precisely know the true *functional forms* of one or more equations that characterize the management problem. For example, a fishery manager may be uncertain as to which of two possible state equations—one that features a population collapse threshold and one that does not—describes stock dynamics. In the case of air quality, there could be competing nonlinear models that dictate the formation and decay of secondary particulates like ozone.

State Uncertainty

State uncertainty arises when a physical state variable is observed imperfectly (e.g., when the precise value of the state variable is unknown). This typically occurs when the true state is difficult to measure (e.g., wild species populations or pollution conditions over space) and assessments are imperfect (e.g., when estimating fish stocks or observing coarse air quality data). State uncertainty is conceptually different from parametric and model uncertainty because the latter two typically involve uncertainty about a fixed object (e.g., the true parameter value) while state uncertainty centers on an unknown value that is changing over time at least in part due to management choices. Furthermore, when economists develop models with parameter and model uncertainty, any prediction errors by the manager concerning stock dynamics do not compound over time because the manager will observe the state variable perfectly at some point in each period in the economic model. With state uncertainty, prediction errors can compound over time.

Modeling Learning

In the context of resource and environmental management, although stochasticity is irreducible, parameter, model, and state uncertainty can often be reduced over time through learning. Modeling the costs and benefits of learning has a long history in the economics literature (Blackwell 1951; Marschak and Miyasawa 1968). For example, an air quality manager may learn from observing how reductions in emissions affect ambient pollution levels. A fisheries manager may learn about stock dynamics after a temporary fishery closure. It is important to identify and understand differences in models of learning because management incentives—and the technical solution method to the models themselves—depend on how learning is modeled.

The traditional approach to modeling learning is to introduce a parameterized distribution that characterizes the manager's beliefs over the true value or functional form of an uncertain model component, such as a Bayesian prior.[10] For example, beliefs might capture uncertainty over the true value of a parameter governing air pollution dissipation.[11]

Uncertainty may also be reduced by investing in *direct* learning. For example, the manager can monitor to reduce stock uncertainty. In the fishery management case, the manager could pay for stock assessments. Similarly, the manager could hire additional ecologists to refine parametric uncertainty. These investments need not be mutually exclusive: an investment in stock assessment could also reduce uncertainty about parameters in the state equation.

Characterizing How Learning Occurs

We next examine how the manager's strategy for maximizing the expected NPV of benefits is shaped by the modeler's assumption about how learning occurs. These three modeling approaches fall into two general categories, whereby there is either no learning (i.e., nonadaptive management) or learning is explicitly allowed and modeled (i.e., adaptive management). Within the adaptive management approach, there are two types of learning, which depend on how the manager's choices are modeled.

Nonadaptive Management

Under nonadaptive management (NAM), the decision maker treats beliefs as fixed both when determining optimal policy and when stepping forward in time. All uncertainty is treated as if it is irreducible, even if learning is technically feasible.

When all sources of uncertainty are treated as irreducible, any uncertain components are captured by fixed parameters that describe the irreducible uncertainty (rather than by parameters that change over time). Thus there is no need to introduce belief parameters. This means

[10] Formally, let the parameters of the distribution summarizing the state of beliefs in a period be given by $\boldsymbol{\beta}_t$, with information dynamics given by

$$\boldsymbol{\beta}_{t+1} = g(\boldsymbol{\beta}_t).$$

The function $g(\ )$ captures how the manager changes his/her beliefs about a system over time. Beliefs are usually assumed to be a Bayesian update performed using observed state variable dynamics in response to a management action (e.g., the chosen control variable). The updating process is usually modeled iteratively after a period's stochastic outcomes are observed and before the manager makes the next period's decisions.

[11] Since state variables are inherently dynamic (whereas parameters and models of the state equation are more often fixed), updating beliefs about an unknown state must account for both stochastic state dynamics *and* learning about the level of an uncertain state variable at a given point in time. Thus information dynamics are unlikely to be captured by simple Bayesian updating, which makes state uncertainty problems involving realistic models (i.e., beyond a handful of possible discrete states) computationally challenging.

that the stochastic dynamic programming problem for the nonadaptive manager is technically straightforward to solve: a NAM policy is solely a function of the current state variable. Although *future* management actions will be based on new realizations of a state variable, learning to reduce uncertainty does not occur. Rather, the regulator observes realizations of a stochastic random variable.

Adaptive Management

Under the adaptive management (AM) approach, beliefs are updated over time. This updating depends on a set of system characteristics, such as the history of management actions, observed state variables, and observed rewards. However, there are two ways to implement AM, which are based on whether learning is included in the identification of the optimal policy.

Passive adaptive management

In the simplest case, called passive adaptive management (PAM), the manager treats beliefs as fixed, which means that from the manager's perspective, the arrival of information is exogenous. Therefore, PAM ignores the fact that choices about the management action (e.g., the control variable) might lead to more or less learning. This means that conditional on a given set of information, the PAM and NAM policy solutions are identical. However, over time, the PAM manager updates his/her beliefs. Thus, in the long run, we would expect the PAM approach to outperform the NAM approach because PAM will become more predictable from the manager's perspective (i.e., updated beliefs more closely approximate the true underlying model).

Active adaptive management

In contrast to the PAM approach, an *active* adaptive management (AAM) decision maker accounts for learning when making management decisions. This means that AAM decisions will deviate from PAM/NAM decisions when the expected benefits of learning outweigh the opportunity cost of forgoing the optimal policy when current beliefs are held fixed due to the "explore versus exploit" trade-off (Nicol and Chadès 2012; Fackler 2014; Springborn 2014). Thus, in the short term, active learning under AAM can be viewed as a "costly" investment relative to PAM.

Alternatives for characterizing uncertainty and learning

Table 1 presents one way for modelers and real-world managers to consider stochasticity, uncertainty, and learning. The table identifies ten different ways in which uncertainty and learning can manifest in an environmental or resource management problem. The four different forms of uncertainty (columns) are matched with a sample of methods for dealing with uncertainty (rows). Since no learning occurs in a stochastic problem, the PAM and AAM learning approaches are not feasible for that type of uncertainty.[12]

[12]See the online supplementary materials for a discussion of how even this characterization of stochasticity, uncertainty, and learning is in some ways incomplete and an illustration of how PAM and AAM problems have fundamentally different solution methods.

Table 1 Characterization of possible uncertainties matched with potential solution methods/characterizations of learning in a resource manager's problem

Solution/learning	Uncertainty			
	Stochasticity	Parameter	State	Model
Nonadaptive	NAM-Stoch	NAM-P	NAM-S	NAM-M
Passive	n.a.	PAM-P	PAM-S	PAM-M
Active	n.a.	AAM-P	AAM-S	AAM-M

Notes: AAM = active adaptive management; n.a. = not applicable; NAM = nonadaptive management; PAM = passive adaptive management.

In the next four sections we examine whether and how the four general environmental and resource management "hypotheses" apply to the different uncertainties and characterizations of learning in table 1.

Hypothesis 1: Greater Levels of Uncertainty Result in More Precautionary Management

One management method for dealing with uncertainty is precaution. Examples of precautionary management include delaying action in anticipation of learning, waiting for resolution of stochasticity, or forgoing an action that increases the likelihood of an undesirable outcome above some well-defined level (Gollier and Treich 2003). Exercising more or less precaution implies deviating from some baseline management action specifically because of uncertainty. For the purposes of the discussion here, we define precaution as a change in management strategy that is aimed at avoiding bad outcomes, which means that our interpretation aligns most closely with a "strong" form of precaution discussed in the literature (Morris 2000; Foster et. al. 2000).[13] An example would be an air quality management plan that maximizes the NPV of benefits, subject to pollution levels not exceeding some maximum level more than 3 days per year (as opposed to an unconstrained maximization of the NPV of benefits). If more uncertainty in the pollution levels leads to more stringent regulation, we call it precaution. Thus, to evaluate Hypothesis 1, we examine how greater uncertainty affects the behavior of the manager in models of optimal environmental and resource management and then assess qualitatively whether the change in behavior is characterized by precaution. We first examine stochasticity and then evaluate uncertainty, paying attention to nuances within and between them.

Stochasticity and Precaution

Greater levels of uncertainty may—but not always—induce policies that could be interpreted as precautionary.[14] For example, if precaution is defined by lower permitted levels of

[13]Weak precaution is like a management strategy that seeks to maximize expected NPV. In that sense, learning in an AAM framework fits more naturally within a weak precaution framework because it is less constraining on the manager, whereas strong precaution limits management activities which *could* harm humans or the environment.

[14]We restrict the discussion here to traditional models with "thin" rather than "fat" tails and acknowledge that some of these results might not carry through to uncertainties that arise in some environmental problems, such as climate change (see Weitzman 2009).

pollution emissions or resource exploitation compared with a less volatile evolution of the state variable (e.g., pollution or the resource stock), then the general relationship between stochasticity and precaution is indeterminate (Pindyck 1984).[15] We highlight two ways in which such a result can occur.

Level of stochasticity

Consider harvesting a species to extinction: will greater stochasticity in annual growth rates lead to more precautionary management? Not necessarily. An example of *less* precaution resulting from *greater* stochasticity occurs in Olson and Roy (2000), whose model of renewable resource consumption features a nonconcave growth function and nonconsumptive utility from the standing stock. They show that introducing growth stochasticity with a distribution specified to make the stock more productive can lead to a harvest policy that is nearly guaranteed to cause extinction.[16]

Irreversible outcomes

When a bad outcome is irreversible, precaution takes on a more explicit temporal dimension. More specifically, quasi-option value (Arrow and Fisher 1974), which typically focuses on the flexibility lost from irreversible environmental degradation (e.g., the regulator loses previously available policy options), is used to argue for delaying activities that degrade the environment or overexploit resources and thus can motivate precaution. In contrast, option value (Dixit and Pindyck 1994) typically focuses on flexibility lost due to sunk costs. Option value is used to argue for delaying environmental degradation or overexploitation of resources when they specifically involve sunk costs (Conrad 2000). Insofar as irreversible environmental damage is an example of a sunk cost, both option value and quasi-option value can be used to argue for delayed intervention due to precaution (Pindyck 2000, 2002). Similarly, when both the timing and magnitude of resource exploitation or conservation can be selected (see the optimal stopping literature [Hanemann 1989]), there is a trade-off between precaution in terms of more immediate conservation and precaution in terms of less exploitation (Sims and Finnoff 2013).

Uncertainty and Precaution

Precautionary policies are sometimes justified on the grounds that it is prudent to delay action that carries a risk of harm until uncertainty is resolved. When uncertainty is reducible through either active or passive learning, is decision making generally characterized by precaution?

Parameter uncertainty

Under parameter uncertainty, the clear answer is no. A common feature of optimal policy in the AM literature is using policy tools to "explore" or "probe" for information about

[15]Note that the relationship can also be nonmonotonic, that is, small increases in the volatility of the evolution of the state variable can lead to more precaution but large increases can lead to less precaution (Saphores 2003).
[16]Note that here the increase in stochasticity is not mean-preserving. Extinction can be optimal in this model as well.

underlying parameters governing resource dynamics. Smith and Walters (1981) show that such probing may involve a decreasing resource harvest in a way that resembles precaution. On the other hand, probing may involve a more intensive harvest when parametric uncertainty is high due to a history of stock management that implemented low levels of exploitation. Using an AAM fishery management model, Springborn and Sanchirico (2013) find that greater uncertainty over a stock productivity parameter may lead to a greater harvest for reasons other than learning; this is because when uncertainty increases around a low point estimate of stock productivity, the decision maker may place more weight on the possibility that the true productivity is higher, which drives up the harvest.

Interaction of parameter uncertainty and stochasticity

Precautionary behavior may depend on the source of stochasticity and whether there is more than one form of stochasticity. If there are two sources of stochasticity (e.g., in resource dynamics and policy implementation), more stochasticity can lead to a faster reduction in polluting activities or resource exploitation (e.g., Carpenter, Ludwig, and Brock 1999; Polasky, de Zeeuw, and Wagener 2011). For example, in an analysis of fishery management decisions, Sethi et al. (2005) find that large increases in the volatility of either stock growth or implementation error do not by themselves qualitatively alter the harvest rule when both growth stochasticity and uncertainty about the resource stock level are low. However, when uncertainty about the resource stock level is high, the harvest policy may be less precautionary, involving a greater harvest over a range of possible stock levels.

Thresholds

One important type of parameter uncertainty occurs when unknown discrete thresholds divide relatively good and bad outcomes (e.g., a threshold level of atmospheric carbon dioxide concentrations leading to runaway climate change, in which case the threshold level is the parameter and atmospheric carbon dioxide concentrations are the state variable). Models of optimal management provide qualified support for a precautionary approach to this class of problems. For example, Lemoine and Traeger (2014) study optimal carbon taxation given the ability to learn about the location of a temperature threshold beyond which climate system dynamics become sharply less favorable. The existence of this threshold increases the optimal carbon tax. The authors find that when there is learning, the optimal tax is initially lower than when there is no learning, but that the learning policy ramps up the tax faster and eventually sets it higher than the no-learning policy as temperatures increase and approach the threshold. Thresholds can also occur directly with a control variable (e.g., how much use of sonar disturbs wildlife) and, in contrast to precautionary behavior, this involves setting the activity within a "risky" range of levels in order to learn (Groeneveld, Springborn, and Costello 2014).

Irreversible decisions

When decisions are irreversible, greater uncertainty often but not always leads to more precautionary optimal management. Here the manager's objective function can matter. Gollier, Jullien, and Treich (2000) show that when utility from consumption exhibits hyperbolic absolute risk aversion, and "prudence" (e.g., the propensity to forgo consumption to hedge against future risk) is sufficiently large, irreversibility does decrease first-period consumption as long as the rate of learning is high. However, the converse is also true: if learning is slow it might not be optimal to forgo consumption. As a result, the precise nature of the problem (in this case, the interaction of preferences and rates of learning) can be important.

Evaluating Hypothesis 1

Do greater levels of uncertainty result in more precautionary management? To summarize, stochasticity (irreducible uncertainty) has an ambiguous effect on the incentives for precautionary management (Pindyck 1984; Olson and Roy 2000; Saphores 2003). Under parameter uncertainty, optimal management may involve manipulating a state variable in order to learn more quickly (Smith and Walters 1981; Springborn and Sanchirico 2013). Such "probing" behavior may also run counter to precaution.[17] When state variable outcomes are irreversible, optimal management is more likely to be precautionary (Gollier, Jullien, and Treich 2000). In contrast, when *decisions* are irreversible, they are more likely to be delayed, which may suggest less precaution. Thus "greater levels of uncertainty" is actually a misnomer: the precise form of uncertainty (e.g., parameter, state, stochasticity) is what matters in determining whether greater uncertainty leads to more precaution.

Hypothesis 2: Reductions in Environmental Uncertainty Have Greater Value Than Reductions in Economic Uncertainty

Uncertainty can enter a management problem through at least two channels: economic parameters such as the costs or benefits of different actions/outcomes (e.g., resource extraction, pollution reduction) and economic states (e.g., downstream market structure) or environmental parameters (e.g., growth rates) and environmental states (e.g., resource stock levels). While the environmental and resource economics literature on environmental uncertainty is vast, there has been less work on reductions in economic uncertainty. This suggests that reductions in environmental uncertainty are viewed in the literature as having greater value than reductions in economic uncertainty. In this section we examine how reductions in economic versus environmental uncertainty affect both optimal management actions and benefits from optimal management. We first investigate the mechanics of reducing the stochasticity of environmental and resource management problems generally and then consider reductions in environmental and economic uncertainty more specifically.

General Mechanics of Reducing Stochasticity

Reductions in economic sources of stochasticity can take several forms. For example, futures markets hedge price uncertainty and the form of input (e.g., labor) contracts in resource industries can spread exposure to market fluctuations (Plourde and Smith 1989; McConnell and Price 2006). Crop, flood, and other types of insurance reduce exposure to risk and are

[17] In an article that does not consider theoretical results from economics, Doremus (2007) makes this point while advocating for "learning while doing" in resource management.

often subsidized by governments (Coble and Barnett 2013). Similarly, reductions in environmental stochasticity can result from management practices such as maintaining larger standing resource stocks (Melbourne and Hastings 2008).

The impact of reductions in stochasticity on expected profits depends on the curvature of the objective function and state equation and how stochasticity is modeled. To the first point, if profits are concave in the environmental variable, then reductions in uncertainty have value (i.e., they result in smaller reductions in expected profits), but the magnitude of this value depends on the curvature of the profit function. If the profit function is more curved in an economic state variable, then economic uncertainty will have a larger impact on expected profits ceteris parabis, and vice versa. The ceteris parabis assumption assumes that changes in uncertainty have no impact on the optimal level of the control variable.[18] To the second point, if the state equation is nonlinear or stochasticity enters the state equation multiplicatively, then optimal management actions depend on the volatility of the random variable (Hoel and Karp 2001). Therefore the solution to the manager's problem and the value of stochasticity reductions depend on whether stochasticity enters a linear–quadratic control problem additively.

One general finding in the literature that arises from the curvature of the profit function is that more stochasticity, regardless of whether its source is economic or environmental, *always* increases the conditional value of information. That is, the value of information is *conditional on being able to react to that information.*[19] This means that if deviations from the status quo are irreversible because of sunk costs or regime shifts, these actions have an opportunity cost that reflects (in part) the forfeited value of information. The value of reacting to new information has implications for a variety of problems characterized by economic and environmental stochasticity, including invasive species control where future species spread and invasion damages are stochastic (Sims and Finnoff 2013); control of a stock pollutant where pollution concentrations and damages are stochastic (Pindyck 2000, 2002); fisheries management with stochastic price and stock dynamics (Nøstbakken 2006); species conservation with stochastic species value and density (Sims et al. 2017); and exhaustible resource extraction with stochastic prices and reserve size (Almansour and Insley 2013).

Reductions in Environmental Uncertainty

Reductions in environmental uncertainty can occur through learning more about the true value of an environmental parameter, a stock level, or the model summarizing an environmental process. Valuing learning—through either PAM or AAM—offers a straightforward way to value reductions in environmental uncertainty.[20] Regardless of the particular mechanism for learning, reductions in uncertainty reflect better predictions of the future stock or parameters (Costello, Polasky, and Solow 2001; Kennedy and Barbier 2013). Furthermore,

[18]This will be the case for a linear–quadratic control problem if stochasticity enters the state equation additively (Newell and Pizer 2003).

[19]Note that the conditional value of information is not simply the *expected* value of information gained by avoiding an irreversibility (Conrad 1980).

[20]This type of learning is very different from the options literature, which sometimes refers to learning that involves observing a stochastic environmental event but not updating beliefs about true parameters, stocks, or models. Such models have no prior distribution to be updated.

valuing reductions in uncertainty through learning maps directly to being able to directly invest in uncertainty reduction (e.g., paying for better monitoring).

Stock uncertainty's impact on the benefits and costs of management actions (e.g., how uncertain pollution levels map to damages or how uncertain resource stocks map to benefits) is a very common form of environmental uncertainty in the literature. In an extension of Roughgarden and Smith (1996), Sethi et al. (2005) identify the optimal solution to the fishery manager's problem when facing three sources of uncertainty: environmental variability in fish growth (stochasticity), fish stock measurement error (stock uncertainty), and inaccurate implementation of harvest quotas (a different form of stock uncertainty due to uncertain escapement). The manager must choose a fishing quota at each point in time to maximize the discounted value of harvest subject to fish stock dynamics and each type of uncertainty. They find that an increase in stock uncertainty through measurement error has the largest impact on the propensity to close a fishery, profits, and extinction risk. However, increases in stochasticity have little impact on policy, expected profits, and extinction risk even when these sources of uncertainty are large. Thus, relative to stochasticity, environmental stock uncertainty can have a large impact on management actions and welfare.

It is important to distinguish between uncertainty that leads to changes in welfare and uncertainty that leads to changes in optimal management actions. Allowing prices and resource stock growth to evolve stochastically, Hanson and Ryan (1998) find that reductions in price uncertainty affect the level of welfare. However, reductions in environmental uncertainty influence both the level of welfare *and* the solution to the manager's problem. Furthermore, if price uncertainty and environmental uncertainty are correlated or the environmental uncertainty is sufficiently complex, optimal management (e.g., choice of a tax or quota to correct an externality) depends on the structure of the correlation or the nature of the complexity (Stavins 1996; Jensen and Vestergaard 2003; Kennedy and Barbier 2015).

Reductions in environmental uncertainty can be valued based on the conditional value of information. Assessing the value of learning is a natural way to put a price on the conditional value of information. Reductions in uncertainty matter for both welfare and optimal management, especially when there is a correlation between economic uncertainty or complex environmental processes.

Reductions in Economic Uncertainty

Economic uncertainty is often modeled as uncertainty in resource demand or the benefits of pollution reduction, both of which have explicit prices in a manager's objective function (e.g., price of fish, costs of health care due to pollution, adaption to climate change).[21] In the manager's profit function or net benefit function, prices often enter the objective function linearly while costs often enter either linearly or subject to a convex function. When they enter linearly, price or cost stochasticity leads to a vertical shift in the marginal benefit or cost of a management action (Weitzman 1974; Reed 1979). Thus the *level* of stochasticity tends not to matter as much as how benefits change in response to management decisions (e.g., the slope

[21] Although compliance cost uncertainty is prevalent in the economics of second-best policies under uncertainty (Weitzman 1974; Stavins 1996; Pizer 2002), we do not directly address this issue here. Goulder and Parry (2008) and Shogren and Taylor (2008) partially address this topic. We also do not address the issue of regulatory uncertainty.

of benefit or cost curves [see Weitzman 1974]). Golosov et al. (2014) show that optimal carbon taxes (the management action or "control variable") are only a function of discount rates, marginal costs of pollution, and pollution decay rates rather than stochastic output levels or prices.

Costs are often modeled as a nonlinear function of the control variable (Hanson and Ryan 1998). If costs are convex, the impacts of uncertainty are different than if the costs are linear. For example, fluctuations in available fishing technologies cause the costs of fishing to be stochastic as vessels decide how to fish (Squires and Vestergaard 2013). Thus economic uncertainty can affect welfare levels.

There is also evidence that output market fluctuations can affect optimal management. In an extension of a fisheries model in Weitzman (2002), Hannesson and Kennedy (2005) find that if profits vary with fishery stock size, taxes generally dominate quotas, but if profits are roughly constant over different stock sizes (e.g., for schooling fish), then quotas dominate when either economic uncertainty or resource growth stochasticity is sufficiently large.

Evaluating Hypothesis 2

Although in many circumstances both price and cost uncertainty appear to have second-order impacts, there are several specific situations in which they do qualitatively change the manager's problem. First-order impacts in both welfare and optimal management can arise when economic uncertainty enters nonlinearly in the objective function—either directly or indirectly. Thus, while reducing environmental uncertainty is certainly of first-order importance, there are cases in which reducing economic uncertainty is also of first-order importance.

Hypothesis 3: Management Strategies with Active Learning Lead to Substantially Higher Welfare Than Those with Passive Learning

The third hypothesis concerns the relative increases in welfare from two types of learning about uncertainty in resource and environmental management: AAM and PAM. As discussed earlier, the key distinction between these two types of learning is that under AAM, the regulator considers costly actions to enhance learning and considers how management decisions can affect beliefs about a parameter, state, or model. In contrast, under PAM, the optimal policy maximizes expected welfare under the assumption that the regulator's information regarding the system's parameters or structure cannot be improved.[22] Deviating from the PAM policy (i.e., using an AAM strategy) may be optimal if the expected gains from making more informed decisions in the future outweigh the costs of forgoing the higher immediate expected profits of the PAM action. From the economic modeler's perspective, some amount of learning and updating is expected under both PAM and AAM. However, the AAM regulator recognizes that different actions lead to different amounts of learning, while the PAM regulator does not. Intuitively, AAM offers welfare improvements relative to PAM because

[22]However, PAM is still adaptive because the regulator assimilates new information ex post.

the regulator has a larger set of options to consider. How much and in what ways AAM improves welfare relative to PAM is, however, an empirical question.

The net payoffs from taking an active versus passive approach to learning have been studied most intensely in the context of parametric uncertainty, with net improvements in expected welfare under AAM with parametric uncertainty often found to be modest. Using simulation approaches, Bond and Loomis (2009), Rout, Hauser, and Possingham (2009), Springborn and Sanchirico (2013), and Fackler (2014) all find that the expected welfare gains from AAM relative to PAM are relatively small, in the range of 0.1 percent to 3 percent of welfare. The likely reason is that in those models, AAM policies often lead not to fundamentally different information, but rather to faster acquisition, which generates a modest net payoff.[23]

It may well be that rather than expecting AAM policies to improve expected welfare relative to PAM policies, an active learning approach (i.e., AAM) should be viewed as a form of insurance. For example, although Springborn (2014) found that in the context of invasive species management, the most common effect of a shift from PAM to AAM was a small welfare *loss*, in the smaller (but nontrivial) share of cases in which initial beliefs were a poor reflection of reality, exploration under AAM uncovered the error and actually protected against large negative welfare impacts. Thus, not surprisingly, Springborn (2014) showed that the relative value of AAM can be much higher under risk aversion.

Evaluating Hypothesis 3

We find that AAM leads to expected welfare gains in almost all cases. However, the expected welfare gains are often small or modest (Springborn and Sanchirico 2013). Nevertheless, AAM seems to mitigate large losses relative to PAM and thus can be very useful when the regulator is risk averse (Springborn 2014).

Hypothesis 4: Managers Who Try to Actively Learn Do So Faster Than Those Who Passively Learn

The fourth hypothesis considers whether the rate of learning is faster in AAM than in PAM. With parametric, model, or state uncertainty, the key difference between AAM and PAM learning is that under AAM, the manager can choose to deviate from the PAM path with the intention of learning. Because the benefit of learning in both PAM and AAM is related to the background stochasticity in a system, we start there.[24]

Optimal learning is very much related to the level of stochasticity in the system (Hausser and Possingham 2008; Rout, Hauser, and Possingham 2009; Springborn and Sanchirico 2013). If there is no stochasticity in an environment or resource problem, the regulator can often learn everything there is to know about the system. Thus, if learning is important in a management problem, stochasticity must also be present.

[23] For example, Springborn and Sanchirico (2013) show that additional returns from an exploratory AAM fisheries policy are small unless departing from the PAM policy is required to facilitate learning.

[24] Karp and Zhang (2006) note that if the regulator anticipates exogenous learning, it leads to less precautionary behavior from the manager because the manager knows that future information will allow him/her to make better decisions. We do not consider exogenous learning rates here.

Adding uncertainty to a stochastic system effectively increases the "randomness" of future states of the world from a regulator's perspective (e.g., more uncertain fish stocks). In a PAM framework, by definition, there is no investment in learning. Rather, the regulator passively incorporates new information about a system to make better decisions. Conversely, in an AAM framework, the regulator acknowledges that each management action can affect the current and future periods' randomness through reductions in uncertainty. For example, a fishery manager could choose to drive a fish stock down close to zero, thereby decreasing the chance of an unexpectedly large stock.

The key insight from both the PAM and AAM frameworks is that observed responses to different management actions lead to learning. Some management actions might put the resource in a state in which learning occurs quickly (e.g., harvesting a resource intensely to measure growth rates). Those management actions have a large "signal strength" and lead to fast learning. Therefore, using AAM implies that more learning ought to occur because there is an incentive to perform management actions that have a high signal strength. However, Hausser and Possingham (2008) and Springborn and Sanchirico (2013)show that an AAM framework with parametric uncertainty can actually lead to less learning than a PAM approach if stochasticity is relatively large. The precise cause of this result is an open question in the environmental and resource literature.

Early theoretical work of Bar-Shalom (1981) offers a nice framework to evaluate what might lead to this result. More specifically, Bar-Shalom (1981) evaluates how a value function (e.g., welfare from resource management) could change when a decision maker learns more about the decision environment. The study decomposes the value of learning into three components, which is a very useful framework for evaluating this hypothesis. The three components, renamed here using terminology intuitive for adaptive resource and environmental management, are deterministic, stochastic, and experimentation. The deterministic component refers to the value when all uncertainties are ignored. The stochastic component describes all randomness, which is assumed to be irreducible. The experimentation component captures the expected future value of information. The PAM decision maker accounts for only the first two components, while the AAM decision maker considers all three components.

All else being equal, the experimentation effect will drive decisions with great signal strength under AAM. However, all else is not equal: there can be an interaction between the stochastic and experimentation components, because learning changes beliefs. For example, decreasing parametric or model uncertainty reduces the stochastic component associated with management outcomes. Similarly, the stochastic effect can drive changes in optimal management. For example, changes in the stochastic effect can enhance precautionary action when negative stochastic outcomes matter for welfare more than positive ones (e.g., random decreases in food leading to population collapse). Thus, even though the experimentation effect drives more "informative" management actions under AAM, anticipation of changes in the stochastic effect due to learning can either increase or decrease the value of "probing" management actions that have a greater signal strength. If the counteraction is strong enough, the PAM policy may actually result in greater learning.

Tol (2014) examines these competing effects as part of an analysis of the effect of learning about damages on optimal climate policy. He finds that accounting for uncertainty results in more stringent policy if the manager is risk averse and there is a greater likelihood of a negative rather than a positive surprise (e.g., more damages than expected from a given management action). The effect of irreversibility is not immediately clear since there are two key irreversibilities that have competing effects: the prospect of being locked into a particular level of climate change supports greater stringency, while being locked into energy/transportation capital supports less stringency. Without learning, the former effect dominates—irreversibility motivates more stringency. However, the effect of learning results in less stringency because there is less need to "hedge" against policy mistakes. Of the 17 studies on the impact of future learning on optimal short-run policy stringency reviewed by Tol (2014), one finds that learning enhances stringency, two find almost no effect, and the remaining fourteen find that optimal stringency decreases.

Evaluating Hypothesis 4

In general, we find that AAM does not always lead to increased learning relative to PAM (Hausser and Possingham 2008; Springborn and Sanchirico 2013). While the precise reason for this result is unclear, Bar-Shalom (1981) offers a potentially useful framework for investigation. Moreover, Tol (2014) suggests that accounting for learning rates can be important for policy.

Conclusions and Research Opportunities

This article has examined several different forms of uncertainty, stochasticity, and learning in resource and environmental management that have thus far been explored only in more technical surveys (e.g., Fackler 2014) and evaluated four "hypotheses" associated with uncertainty and learning in environmental management. Overall, we find that there are important differences across stochasticity and uncertainty with respect to optimal management. Management insights from the literature for different forms of uncertainty are somewhat transferable to different contexts, but it is important to be clear about the nature of the uncertainty investigated. Economic uncertainty can be first-order important, especially when it enters the objective function nonlinearly. Consistent with hypothesis 3, using AAM increases welfare, but inconsistent with hypothesis 4, it does not necessarily increase the rate of learning.[25] Our review of the economics literature has characterized several types of uncertainty (e.g., stochasticity, parameter, state, and model) and different ways to learn (e.g., PAM and AAM). However, both researchers and managers should be careful about applying general rules of thumb for "uncertainty" because the applicability of a particular management rule (e.g., precaution or importance of environmental uncertainty) depends on the specific types of uncertainty, stochasticity, and learning that characterize a particular problem.

[25]These results are also summarized in the online supplementary materials.

We conclude by identifying several topics that would benefit from further research:

1. *State uncertainty*: The economics literature regarding the welfare gains from learning with state uncertainty is less developed than the literature on parametric uncertainty, most likely because it is technically a harder problem to address.[26] One line of research has studied this issue in some depth by borrowing tools from artificial intelligence (AI) and computer science (Chadès et al. 2008, 2012; Regan, Chadès, and Possingham 2011). Tools used in AI such as partially observable Markov decision process (POMDP) solution methods (Fackler and Haight 2014) appear to offer a promising approach for addressing this issue technically.[27] Economists are increasingly using this tool to evaluate policy, but state uncertainty is more poorly understood than parameter uncertainty (Haight and Polasky 2010; Fackler and Haight 2014; Baggio and Fackler 2016; MacLachlan, Springborn, and Fackler 2016; Kling, Sanchirico, and Fackler 2017).
2. *Passive learning*: In the past 10 years, most research on adaptive management and optimal learning on the part of a regulator/manager has been in the context of natural resource management. While learning has already been examined in the environmental economics literature, insights could likely be gained from applying recent advances from that literature to, for example, the regulator's problem of imperfect knowledge of abatement costs in pollution control or imperfect knowledge of health, environmental costs, or adaption costs.
3. *Dynamic decentralized learning*: Environmental and natural resource management issues are often characterized by hierarchical or adjoining regulators. For example, because many natural resources are transboundary, multiple regulatory agencies must interact in order to maximize overall welfare. While interaction among agents under uncertainty is central to many high-profile resource and environmental management problems, with some notable exceptions (e.g., Hoel and Karp 2001; Dijkstra and Fredriksson 2010), there has been insufficient development of fully dynamic multiagent models. An important area of economics that would benefit from advancement on this front is environmental federalism (Baumol and Oates 1988). Dynamic multiagent analysis may help shed light on apparently unresolved tensions in this influential area of theory. For example, which regulatory arrangement is more efficient for a transboundary pollution problem with uncertainty but where learning is feasible: multiple nimble regulators highly informed about local context but burdened by costly coordination, or a slower federal regulator with coarser local knowledge and authority to act over a larger jurisdictional scale?

References

Almansour, A., and M. C. Insley. 2013. The impact of stochastic extraction cost on the value of an exhaustible resource: an application to the Alberta oil sands. Available at SSRN: https://ssrn.com/abstract=2287596 or http://dx.doi.org/10.2139/ssrn.2287596.

Arrow, K. J., and A. C. Fisher. 1974. Environmental preservation, uncertainty, and irreversibility. *Quarterly Journal of Economics* 88(2):312–19.

Baggio, M., and P. L. Fackler. 2016. Optimal management with reversible regime shifts. *Journal*

[26]State variables are inherently dynamic, and modeling belief dynamics becomes much more complex when the manager is learning about a moving rather than fixed target.

[27]See the online supplementary materials for a discussion of the AM approach that includes the use of POMDP solution methods.

of Economic Behavior and Organization 132(pt. B):124–36.

Bar-Shalom, Y. 1981. Stochastic dynamic programming: caution and probing. *IEEE Transactions on Automatic Control* 26(5):1184–95.

Baumol, W. J., and W. E. Oates. 1988. *The theory of environmental policy*. Cambridge: Cambridge University Press.

Blackwell, D. 1951. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, ed. J. Neyman, 93–102. Berkeley: University of California Press.

Bond, C. A., and J. B. Loomis. 2009. Using numerical dynamic programming to compare passive and active learning in the adaptive management of nutrients in shallow lakes. *Canadian Journal of Agricultural Economics/Revue canadienne d'agroeconomie* 57(4):555–73.

Carson, R. T., and J. LaRiviere. 2017. Structural uncertainty and pollution control: optimal stringency with unknown pollution sources. *Environmental and Resource Economics* https://doi.org/10.1007/s10640-017-0156-1.

Carpenter, S. R., D. Ludwig, and W. A. Brock. 1999. Management of eutrophication for lakes subject to potentially irreversible change. *Ecological Applications* 9(3):751–71.

Chadès, I., E. McDonald-Madden, M. A. McCarthy, B. Wintle, M. Linkie, and H. P. Possingham. 2008. When to stop managing or surveying cryptic threatened species. *Proceedings of the National Academy of Sciences of the United States of America* 105(37):13936–40.

Chadès, I., J. Carwardine, T. Martin, S. Nicol, R. Sabbadin, and O. Buffet. 2012. MOMDPs: a solution for modelling adaptive management problems. Presented at the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI-12), Toronto, Canada, July 2012.

Charles, A. T. 1998. Living with uncertainty in fisheries: analytical methods, management priorities and the Canadian groundfishery experience. *Fisheries Research* 37(1–3):37–50.

Clark, C. W. 1990. *Mathematical bioeconomics: the optimal management of renewable resources.* Hoboken, NJ: John Wiley & Sons.

Coble, K. H., and B. J. Barnett. 2013. Why do we subsidize crop insurance? *American Journal of Agricultural Economics* 95(2):498–504.

Conrad, J. M. 1980. Quasi-option value and the expected value of information. *Quarterly Journal of Economics* 94(4):813–20.

Conrad, J. M. 2000. Wilderness: options to preserve, extract, or develop. *Resource and Energy Economics* 22(3):205–19.

Costello, C., S. Polasky, and A. Solow. 2001. Renewable resource management with environmental prediction. *Canadian Journal of Economics/Revue canadienne d'économique* 34:196–211.

Dixit, A. K., and R. S. Pindyck. 1994. *Investment under uncertainty*. Princeton, NJ: Princeton University Press.

Dijkstra, B. R., and P. G. Fredriksson. 2010. Regulatory environmental federalism. *Annual Review of Resource Economics* 2(1):319–39.

Doremus, H. 2007. Precaution, science, and learning while doing in natural resource management. *Washington Law Review* 82:547–79.

Fackler, P. L. 2014. Structural and observational uncertainty in environmental and natural resource management. *International Review of Environmental and Resource Economics* 7(2):109–39.

Fackler, P. L., and R. G. Haight. 2014. Monitoring as a partially observable decision problem. *Resource and Energy Economics* 37:226–41.

Foster, K. R., P. Vecchia, H. Michael, and M. H. Repacholi. 2000. Science and the precautionary principle. *Science* 288(5468):979–81.

Fowlie, M., and N. Mueller. 2013. Market-based emissions regulation when damages vary across sources: What are the gains from differentiation? Energy Institute at Haas Working Paper 237.

Gilboa, I. 2009. *Theory of decision under uncertainty*, vol. 1. Cambridge: Cambridge University Press.

Gollier, C., B. Jullien, and N. Treich. 2000. Scientific progress and irreversibility: an economic interpretation of the 'Precautionary Principle'. *Journal of Public Economics* 75(2):229–53.

Gollier, C., and N. Treich. 2003. Decision-making under scientific uncertainty: the economics of the precautionary principle. *Journal of Risk and Uncertainty* 27(1):77–103.

Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski. 2014. Optimal taxes on fossil fuel in general equilibrium. *Econometrica* 82(1):41–88.

Goulder, L. H., and I. W. Parry. 2008. Instrument choice in environmental policy. *Review of Environmental Economics and Policy* 2(2):152–74.

Groeneveld, R. A., M. Springborn, and C. Costello. 2014. Repeated experimentation to learn about a flow-pollutant threshold. *Environmental and Resource Economics* 58(4):627–47.

Haight, R. G., and S. Polasky. 2010. Optimal control of an invasive species with imperfect information about the level of infestation. *Resource and Energy Economics* 32(4):519–33.

Hanemann, W. M. 1989. Information and the concept of option value. *Journal of Environmental Economics and Management* 16(1):23–37.

Hannesson, R., and J. Kennedy. 2005. Landing fees versus fish quotas. *Land Economics* 81(4):518–29.

Hanson, F. B., and D. Ryan. 1998. Optimal harvesting with both population and price dynamics. *Mathematical Biosciences* 148(2):129–46.

Hauser, C. E., and H. P. Possingham. 2008. Experimental or precautionary? Adaptive management over a range of time horizons. *Journal of Applied Ecology* 45(1):72–81.

Hoel, M., and L. Karp. 2001. Taxes and quotas for a stock pollutant with multiplicative uncertainty. *Journal of Public Economics* 82(1):91–114.

Jensen, F., and N. Vestergaard. 2003. Prices versus quantities in fisheries models. *Land Economics* 79(3):415–25.

Karp, L., and J. Zhang. 2006. Regulation with anticipated learning about environmental damages. *Journal of Environmental Economics and Management* 51(3):259–79.

Kennedy, C. J., and E. B. Barbier. Renewable resource management with environmental prediction: the importance of structural specification. *Canadian Journal of Economics/Revue canadienne d'économique* 46(3):1110–22.

———. 2015. Renewable resource harvesting under correlated biological and economic uncertainties: implications for optimal and second-best management. *Environmental and Resource Economics* 60(3):371–93.

Kling, D. M., J. N. Sanchirico, and P. L. Fackler. 2017. Optimal monitoring and control under state uncertainty: application to lionfish management. *Journal of Environmental Economics and Management* 84:223–45.

Lemoine, D., and C. Traeger. 2014. Watch your step: optimal policy in a tipping climate. *American Economic Journal: Economic Policy* 6(1):137–66.

MacLachlan, M., M. R. Springborn, and P. Fackler. 2016. Learning about a moving target in resource management: optimal Bayesian disease control. Working paper, University of California, Davis.

Marschak, J., and K. Miyasawa. 1968. Economic comparability of information systems. In *Economic information, decision and prediction*, vol. II, part II, chap. 29. Boston: D. Reidel.

McConnell, K. E., and M. Price. 2006. The lay system in commercial fisheries: origin and implications. *Journal of Environmental Economics and Management* 51(3):295–307.

Melbourne, B. A., and A. Hastings. 2008. Extinction risk depends strongly on factors contributing to stochasticity. *Nature* 454(7200):100–103.

Morgan, M. G., and M. Henrion. 1990. *Uncertainty: a guide to dealing with uncertainty in quantitative risk and policy analysis.* New York: Cambridge University Press.

Morris, J. 2000. Defining the precautionary principle. In *Rethinking risk and the precautionary principle*, ed. J. Morris, chap. 1. Boston: Butterworth Heinemann.

Newell, R. G., and W. A. Pizer. 2003. Regulating stock externalities under uncertainty. *Journal of Environmental Economics and Management* 45(2):416–32.

Nicol, S., and I. Chadès. 2012. Which states matter? An application of an intelligent discretization method to solve a continuous POMDP in conservation biology. *PLoS One* 7(2):e28993.

Nøstbakken, L. 2006. Regime switching in a fishery with stochastic stock and price. *Journal of Environmental Economics and Management* 51(2):231–41.

Olson, L. J., and S. Roy. 2000. Dynamic efficiency of conservation of renewable resources under uncertainty. *Journal of Economic Theory* 95(2):186–214.

Pindyck, R. S. 1984. Uncertainty in the theory of renewable resource markets. *Review of Economic Studies* 51(2):289–303.

———. 2000. Irreversibilities and the timing of environmental policy. *Resource and Energy Economics* 22(3):233–59.

———. 2002. Optimal timing problems in environmental economics. *Journal of Economic Dynamics and Control* 26(9):1677–97.

———. 2007. Uncertainty in environmental economics. *Review of Environmental Economics and Policy* 1(1):45–65.

Pizer, W. A. 2002. Combining price and quantity controls to mitigate global climate change. *Journal of Public Economics* 85(3):409–34.

Plourde, C., and J. B. Smith. 1989. Crop sharing in the fishery and industry equilibrium. *Marine Resource Economics* 6:179–93.

Polasky, S., A. de Zeeuw, and F. Wagener. 2011. Optimal management with potential regime shifts. *Journal of Environmental Economics and Management* 62(2):229–40.

Reed, W. J. 1979. Optimal escapement levels in stochastic and deterministic harvesting models. *Journal of Environmental Economics and Management* 6(4):350–63.

Regan, T. J., I. Chadès, and H. P. Possingham. 2011. Optimally managing under imperfect detection: a method for plant invasions. *Journal of Applied Ecology* 48(1):76–85.

Roughgarden, J., and F. Smith. 1996. Why fisheries collapse and what to do about it. *Proceedings of the National Academy of Sciences of the United States of America* 93(10):5078–83.

Rout, T. M., C. E. Hauser, and H. P. Possingham. 2009. Optimal adaptive management for the translocation of a threatened species. *Ecological Applications* 19(2):515–26.

Saphores, J. D. 2003. Harvesting a renewable resource under uncertainty. *Journal of Economic Dynamics and Control* 28(3):509–29.

Sethi, G., C. Costello, A. Fisher, M. Hanemann, and L. Karp. 2005. Fishery management under multiple uncertainty. *Journal of Environmental Economics and Management* 50(2):300–318.

Shaw, W. D., and R. T. Woodward. 2008. Why environmental and resource economists should care about non-expected utility models. *Resource and Energy Economics* 30(1):66–89.

Shogren, J. F., and L. O. Taylor. 2008. On behavioral-environmental economics. *Review of Environmental Economics and Policy* 2(1):26–44.

Sims, C., and D. Finnoff. 2013. When is a "wait and see" approach to invasive species justified? *Resource and Energy Economics* 35(3):235–55.

Sims, C., D. Finnoff, A. Hastings, and J. Hochard. 2017. Listing and delisting thresholds under the Endangered Species Act. *American Journal of Agricultural Economics* 99(3):549–70.

Smith, A. D. M., and C. J. Walters. 1981. Adaptive management of stock-recruitment systems. *Canadian Journal of Fisheries and Aquatic Sciences* 38(6):690–703.

Springborn, M. R. 2014. Risk aversion and adaptive management: insights from a multi-armed bandit model of invasive species risk. *Journal of Environmental Economics and Management* 68(2):226–42.

Springborn, M., and J. N. Sanchirico. 2013. A density projection approach for non-trivial information dynamics: adaptive management of stochastic natural resources. *Journal of Environmental Economics and Management* 66(3):609–24.

Squires, D., and N. Vestergaard. 2013. Technical change and the commons. *Review of Economics and Statistics* 95(1):1769–87.

Stavins, R. N. 1996. Correlated uncertainty and policy instrument choice. *Journal of Environmental Economics and Management* 30(2):218–32.

Tol, R. S. 2014. Correction and update: the economic effects of climate change. *Journal of Economic Perspectives* 28(2):221–25.

Walters, C. J., and R. Hilborn. 1976. Adaptive control of fishing systems. *Journal of the Fisheries Board of Canada* 33(1):145–59.

Weitzman, M. L. 1974. Prices vs. quantities. *Review of Economic Studies* 41(4):477–91.

———. 2002. Landing fees vs harvest quotas with uncertain fish stocks. *Journal of Environmental Economics and Management* 43(2):325–38.

———. 2009. On modeling and interpreting the economics of catastrophic climate change. *Review of Economics and Statistics* 91(1):1–19.