A density projection approach for non-trivial information dynamics: Adaptive management of stochastic natural resources

Michael Springborn [a,*], James N. Sanchirico [b]

[a] Department of Environmental Science and Policy, University of California, Davis, United States
[b] Department of Environmental Science and Policy, University of California, Davis and University Fellow, Resources for the Future, Washington DC, United States

ARTICLE INFO

ABSTRACT

We demonstrate a density projection approximation method for solving resource management problems with imperfect state information. The method expands the set of partially-observed Markov decision process (POMDP) problems that can be solved with standard dynamic programming tools by addressing dimensionality problems in the decision maker's belief state. Density projection is suitable for uncertainty over both physical states (e.g. resource stock) and process structure (e.g. biophysical parameters). We apply the method to an adaptive management problem under structural uncertainty in which a fishery manager's harvest policy affects both the stock of fish and the belief state about the process governing reproduction. We solve for the optimal endogenous learning policy—the active adaptive management approach—and compare it to passive learning and non-learning strategies. We demonstrate how learning improves efficiency but typically follows a period of costly short-run investment.

1. Introduction

The traditional method of attempting to resolve structural uncertainty on ecosystem dynamics is to invest in natural science research (Moore and McCarthy, 2010). Because management decisions are made on a repeated basis, however, there is also the possibility of using these actions to learn about the ecosystem, i.e. to practice adaptive management. In one of the foundational works on adaptive management, biologist Carl Walters argued that the primary means of reducing uncertainty in environmental models is "through experience with management itself rather than through basic research or the development of (theory)" (Walters, 1986).

Optimal investment in learning depends on the expected value of information generated (Hartmann et al., 2007; Springborn et al., 2010). All else equal, the faster a resource manager resolves structural uncertainty, the sooner the manager is able to reduce management error resulting from the difference between the true structure of a system and the manager's beliefs about that structure. However, accelerating the process of learning, relative to the optimal non-learning policy, comes with an opportunity cost, in the form of short run losses due to learning-driven deviations. An optimal endogenous

* Corresponding author. Fax. +1 530 752 3350.
  E-mail addresses: mspringborn@ucdavis.edu (M. Springborn), jsanchirico@ucdavis.edu (J.N. Sanchirico).

learning approach—sometimes called active adaptive management (AAM)—involves balancing these tradeoffs by pursuing informative actions, but only when the investment is justified by the value of reducing uncertainty.

We employ an adaptive control framework to model a fishery manager choosing the optimal harvest in each period subject to both irreducible uncertainty (stochasticity) and reducible (structural) uncertainty in the population dynamics as a partially observed Markov decision process (POMDP) (see also Zhou et al., 2010; McDonald-Madden et al., 2010, and Bond and Loomis, 2009).[1] Learning to reduce structural uncertainty is accomplished by observing the population, which a manager can influence by altering harvest policy.[2] A POMDP is a sequential decision problem in which at least one underlying state is not known with certainty. Using the concept of a *belief state*, i.e. encoding the decision-maker's probabilistic beliefs about the uncertain state, the POMDP can be converted to a continuous-state Markov decision process (MDP), or *belief* MDP (Bertsekas, 1995).

Our paper makes three important contributions to the growing literature on adaptive control in natural resource economics. First, we use a hierarchical uncertainty model to flexibly represent uncertainty. Second, we relax constraints imposed in the existing AAM literature and much of the economics literature that limit the action space or belief space to a few discrete alternatives. Finally, we employ for the first time in the resource economics and adaptive management literature (to our knowledge) a *density projection* approximation method that enables us to solve an otherwise intractable high-dimensional control problem. The approach is general: the density projection method results in concise analytical conditions for determining information dynamics which hold across a broad set of applications (exponential family distributions).

Hierarchical modeling—also called multilevel or state-space modeling—is a framework in which complex processes, such as uncertain fish population dynamics, are broken down into a set of conditional sub-models that are coherently tied together using simple probability structures (Wikle, 2003).[3] We use a hierarchical model of uncertainty to allow for a more realistic and flexible representation of uncertainty in ecosystem dynamics and therefore the beliefs of the managers. Practical usage of hierarchical models in dynamic resource management models, however, has been limited by state space dimensionality concerns.

To address the challenge of the growing dimensionality of the belief state, researchers often discretize and constrain the belief state space. For example, extending a shallow lake pollution model to incorporate learning, Bond and Loomis (2009) limit uncertainty over a critical nutrient threshold to two possible levels. Another recent example is Johnson (2011), who describes adaptive management of waterfowl harvest by the U.S. Fish and Wildlife Service, who aims to discern between four competing models of population dynamics.

Rather than constrain the belief state space that could limit the ability to map model outputs to management actions, we use a density projection method to solve for the optimal AAM policy in the presence of a hierarchical model of uncertainty. A limited set of hierarchical models are, in Bayesian terminology, conjugate, meaning that an explicit closed form solution exists for the belief state dynamics. However, hierarchical models are typically non-conjugate, as is the case here. We gain tractability for our non-conjugate hierarchical model by approximating the posterior distribution of the manager's beliefs using a function that minimizes the Kullback–Leibler (KL) divergence between true updated beliefs and an approximation (Zhou et al., 2010). The approximation of the belief state through this density projection[4] method converts the *belief* MDP to a *projected-belief* MDP with lower dimensionality. Relatively simple analytical expressions can be derived to identify the parameters of the projected or approximate beliefs for a wide range of applications. A further advantage of this conversion is the ability to solve for the AAM solution using standard dynamic programming methods, such as value or policy iteration (Zhou et al., 2010).

Finally, a number of studies consider applications similar to ours (though not with a hierarchical uncertainty model) where a manager can learn about a system by observing population dynamics (e.g., Hauser and Possingham (2008), Rout et al. (2009) and McDonald-Madden et al. (2010)). However, in these studies, the researchers constrain the set of management actions to a set of 2–3 discrete choices for tractability (e.g., high and low harvest levels). While the limited set of control actions can be informative for how learning might change management responses and it reduces the dimensionality of the problem, we allow for a continuous choice control variable that is more consistent with how fish populations (and many other natural resources) are managed. A notable exception to the use of limited action and/or belief spaces is Wieland (2000) who allows for continuous beliefs and actions in a generic, linear regression-style learning model (Kalman filter). However, Wieland assumes a convenient, conjugate structure using normal distributions for tractability, in contrast to the framework here which does not require selecting from a limited set of conjugate models.

---

[1] These two types of uncertainty are sometimes referred to as aleatory uncertainty, which refers to randomness in samples, and epistemic uncertainty stemming from basic lack of knowledge about particular processes (Paté-Cornell and Elisabeth, 1996).

[2] Learning by observing population dynamics without the requirement of perturbing the system is similar to the information dynamics in Rout et al. (2009). An alternative approach often presented in the engineering literature on adaptive dual control (see Filatov and Unbehauen (2000) for an overview) and in the foundational text of Walters (1986) is to model information as emerging from "exciting" or perturbing the system from the status quo (e.g., perturbing the system by changing the stock of fish). Our core methodological contributions are adaptable to either setting, though specific numerical results will obviously differ depending on the nature of the learning process.

[3] A common hierarchical model used in econometrics is the random effects model. More than simply a technical approach, Royle, Dorazio (2008) argue that hierarchical modeling is a conceptual framework for doing science that is flexible, fosters the fundamental activity of model construction and elucidates the nature of inference. Several recent reviews remark on the increasing popularity of the framework (Gelfand, 2012; Schaub and Kéry, 2012). Halstead et al. (2012) state that "we have likely only seen the tip of the iceberg with regard to the utility of these models in ecology" (p. 134).

[4] This general approach is known by other terms, e.g. *assumed density filtering* (Maybeck, 1982) and *belief compression* (e.g. Roy et al., 2005).