

University of Nebraska - Lincoln

DigitalCommons@University of Nebraska - Lincoln

USGS Staff -- Published Research

US Geological Survey

2011

Adaptive resource management and the value of information

Byron K. Williams

U.S. Geological Survey, Cooperative Research Units, Byronkenwilliams@usgs.gov

Mitchell J. Eaton

U.S. Geological Survey, North Carolina State University, meaton@usgs.gov

David R. Breininger

Innovative Health Applications

Follow this and additional works at: <https://digitalcommons.unl.edu/usgsstaffpub>




Part of the [Geology Commons](#), [Oceanography and Atmospheric Sciences and Meteorology Commons](#), [Other Earth Sciences Commons](#), and the [Other Environmental Sciences Commons](#)

Williams, Byron K.; Eaton, Mitchell J.; and Breininger, David R., "Adaptive resource management and the value of information" (2011). *USGS Staff -- Published Research*. 1155.

<https://digitalcommons.unl.edu/usgsstaffpub/1155>

This Article is brought to you for free and open access by the US Geological Survey at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in USGS Staff -- Published Research by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.



Adaptive resource management and the value of information

Byron K. Williams^{a,*}, Mitchell J. Eaton^b, David R. Breininger^c

^a U.S. Geological Survey, Cooperative Research Units, Reston, VA 20192, United States

^b U.S. Geological Survey, Patuxent Wildlife Research Center, Laurel, MD 20708, United States

^c Innovative Health Applications, IHA-300, Kennedy Space Center, FL 32899, United States

ARTICLE INFO

Article history:

Received 24 April 2011

Received in revised form 1 July 2011

Accepted 5 July 2011

Available online 30 July 2011

Keywords:

Value of information
Adaptive management
Structural uncertainty
Decision making

ABSTRACT

The value of information is a general and broadly applicable concept that has been used for several decades to aid in making decisions in the face of uncertainty. Yet there are relatively few examples of its use in ecology and natural resources management, and almost none that are framed in terms of the future impacts of management decisions. In this paper we discuss the value of information in a context of adaptive management, in which actions are taken sequentially over a timeframe and both future resource conditions and residual uncertainties about resource responses are taken into account. Our objective is to derive the value of reducing or eliminating uncertainty in adaptive decision making. We describe several measures of the value of information, with each based on management objectives that are appropriate for adaptive management. We highlight some mathematical properties of these measures, discuss their geometries, and illustrate them with an example in natural resources management. Accounting for the value of information can help to inform decisions about whether and how much to monitor resource conditions through time.

Published by Elsevier B.V.

1. Introduction

In recent years there has been a growing concern about uncertainty in natural resources management, and a recognition of the impact uncertainty can have on effective resource management. Vulnerability analysis, risk assessment, and adaptive management are three of many approaches that explicitly account for uncertainty and attempt to factor it into resource assessments and management policies. A particularly important source of uncertainty in natural resources concerns a lack of understanding about the factors and processes influencing resource dynamics (Williams, 2010b). This “structural” or “epistemic” uncertainty is often expressed through models that incorporate different hypotheses about how the resource system works. Here incomplete knowledge of natural systems is distinguished from other sources of irreducible uncertainty such as demographic and environmental variation.

Structural uncertainty can limit management effectiveness because in its presence decision making relies on a less than complete understanding of system responses. Conversely, the reduction or elimination of that uncertainty can result in improved management that is better tailored to actual resource dynamics. It often is useful to measure the loss associated with structural uncertainty,

or equivalently, to determine the potential gain in resource value that is possible with its reduction or elimination.

The “value of information” is a generic term for potential management value that is foregone under uncertainty. Among other things, the value of information can inform an assessment of the potential effectiveness of monitoring and analysis to reduce that uncertainty. For example, the value gained by reducing uncertainty can be compared against opportunity and other costs that are associated with collecting and analyzing information, to determine whether monitoring should be undertaken.

The concept of a value of information has been around for several decades and is now well developed. Raiffa and Schlaifer (1961) provided one of the first seminal treatments, coining the name and developing many of its key expressions. Since then it has been applied in economics, finance, medicine, engineering and many other fields (e.g., Frauendorfer, 1992; Bontems and Thomas, 2000; Karnon, 2002; Koerkamp et al., 2006; Eidsvik et al., 2008). However, there have been relatively few applications in ecology and natural resources management, and almost all of them involve an assessment of the utility of information with non-iterative decision making.

Here we describe and discuss the value of information in a context of adaptive management with iterative decision making. A useful framework for adaptive management recognizes management objectives, a range of potential actions, models that forecast resource changes, and measures of confidence in those models (Williams et al., 2007). In adaptive management, actions are taken sequentially over a timeframe, with the choice of an action at any

* Corresponding author. Tel.: +1 703 648 4260; fax: +1 703 648 4269.
E-mail address: Byron.ken.williams@usgs.gov (B.K. Williams).

point guided by its potential consequences for immediate returns and future value. There is an explicit recognition of uncertainty about the biological processes influencing resources dynamics, and the uncertainty is tracked through time along with the status of the managed resource. The generic idea is to reduce uncertainty by monitoring and assessing resource responses to management actions, so that management can be improved over time based on the accumulation of learning (Walters, 1986; Williams et al., 2007).

The objectives of this paper are to describe the value of reducing or eliminating uncertainty through adaptive decision making. We describe several measures of the value of information in terms of management objectives that are appropriate for adaptive decision making. The focus initially is on the expected value of perfect information, and later we include the expected value of partial and sample information. We highlight the mathematical properties of these values, develop the geometry of the expected value of perfect information, and provide a natural resource example to illustrate the value of information.

2. Resource management under uncertainty

A general statement of the natural resource problem under consideration here involves a dynamic resource system that changes over a discrete timeframe $\{0, 1, \dots, T\}$ in response to changing environmental conditions and time-specific management actions. A policy A_t expresses state-specific actions to be taken at each point in time, starting at time t and extending over the remainder of the timeframe. For notational convenience we assume in what follows that there are only finitely many states, so that the system state can be represented with counting indices. We begin with a description of optimal decision making on assumption that resource dynamics are fully known and understood. Then we allow for uncertainty about the system processes and/or parameters in the decision framework. In what follows we use the following notation:

- i and j denote process state.
- a denotes an action.
- q denotes a model state consisting of a probability or confidence measure $q(k)$ for each model in a set of potential process models.
- $P_k(j|i, a)$ is the probability of transition for model k from state i to state j at the next time if action a is taken.
- $R(a|i)$ is the return for action a if the process is in state i .
- $V_k(A_t|i)$ is the process value for model k corresponding to policy A_t , given that the process is in state i at time t .
- $V_t[i, q]$ is the optimal process value at time t over all policies given the process and model states are i and q respectively.
- $Q_t(a|i, q)$ is the process value at time t for a strategy consisting of action a at time t , followed by actions that are optimal over the remainder of the time frame.

Structural certainty. The control of a fully known resource system has been comprehensively treated in the optimal control literature (Puterman, 1994; Bertsekas, 1995). The data consist of realizations $\{i_0, i_1, \dots, i_T\}$ of states, with transition structure

$$j = F(i, a, z), \tag{1}$$

for single-step transitions. Control a at time t is an element of a policy A_t of state- and time-specific actions, and $\{z_t\}$ is a white noise process representing environmental variation. Demographic stochasticities as well as randomness from z induce Markovian transition probabilities $P(j|i, a)$ over the timeframe. First-order Markovian dynamics is a key feature underlying the following development. The problem of optimal decision making is consider-

ably more complex with non-Markovian processes (e.g., Williams, 2007).

The reward structure for this problem is based on returns $R(a|i)$ that depend on system state i and action a . A process value function

$$V(A_t|i) = E \left[\sum_{\tau=t}^T R(a_\tau|i_\tau)|i \right], \tag{2}$$

accumulates returns over time for a particular policy A_t , starting in state i at time t . The value function can be expressed recursively in terms of current and future expected returns (see Williams, 2009):

$$V(A_t|i) = R(a|i) + \sum_j P(j|i, a)V(A_{t+1}|j). \tag{3}$$

Eq. (3) can be used to determine a policy-specific process value for any state at any time, with an optimal policy A_t^* and value function $V_t[i] = V(A_t^*|i)$ given by (Puterman, 1994; Bertsekas, 1995)

$$V_t[i] = \max_a \left\{ R(a|i) + \sum_j P(j|i, a)V_{t+1}[j] \right\}. \tag{4}$$

Structural uncertainty. A key variation on the above framework allows for uncertainty about system functions and their influencing parameters. Here we characterize uncertainty with models incorporating different functional forms (or a discrete set of parameter values), along with a measure of confidence in the ability of each model to describe resource dynamics. Thus, system transitions are described by

$$j = F_k(i, a, z), \tag{5}$$

with the transition from i to j depending on the particular model k . Randomness from z induces a Markovian probability structure $P_k(j|i, a)$ that also is model-specific.

The uncertainty about which model is the most appropriate is represented by a time-varying distribution of confidence measures for the models, denoted here by q and referred to as a *model state*. The distribution specifies a probability mass $q(k)$ for each model k , and averaging the transition probabilities over the model state produces

$$\bar{P}(j|i, a, q) = \sum_k q(k)P_k(j|i, a). \tag{6}$$

The model state evolves through time, with the single-step transition from q to q' expressed via Bayes' theorem (Lee, 1989) as

$$q'(k) = \frac{q(k)P_k(j|i, a)}{\bar{P}(j|i, a, q)}, \tag{7}$$

where the denominator $\bar{P}(j|i, a, q)$ is the average probability of transition to state j . From Eq. (7) the transition from q to q' depends on the state i , the action a taken, and the resulting state j .

The reward structure for this problem also can be model-specific, with state- and action-specific returns $R_k(a|i)$. Accumulating these returns over time produces

$$V_k(A_t|i) = E \left[\sum_{\tau=t}^T R_k(a_\tau|i_\tau)|i \right], \tag{8}$$

and averaging the accumulated returns over the model state produces a process value function

$$V(A_t|i, q) = \sum_k q(k)V_k(A_t|i). \tag{9}$$

Denoting the average return across index values by

$$\bar{R}(a|i, q) = \sum_k q(k)R_k(a|i), \tag{10}$$

the value function can be expressed recursively in terms of current and future expected returns:

$$V(A_t|i, q) = \bar{R}(a|i, q) + \sum_j \bar{P}(j|i, a, q) \sum_k q'(k)V_k(A_{t+1}|j), \tag{11}$$

or

$$V(A_t|i, q) = \bar{R}(a|i, q) + \sum_j \bar{P}(j|i, a, q)V(A_{t+1}|j, q') \tag{12}$$

(Williams, 2009). Eq. (12) can be used to determine an optimal policy A_t^* and value function $V_t[i, q]$:

$$V_t[i, q] = \max_a \left\{ \bar{R}(a|i, q) + \sum_j \bar{P}(j|i, a, q)V_{t+1}[j, q'] \right\}. \tag{13}$$

Including model state in the argument of the value function extends the optimization described by Eq. (4) to cover both the discrete space of system states and the continuous space of model states.

Recognition of uncertainty and the opportunity to reduce it is the basis of an adaptive approach to resource management, which emphasizes the influence of actions not only on resource status but also on learning through time (Walters, 1986). In this particular case, learning is represented by changes in the model state that occur in response to management actions, as in Eq. (7). At each time the state of the system is identified through monitoring, and the model state is updated. The model state reveals the residual uncertainty about resource structure, and systematic change in that uncertainty through time represents incremental learning based on the accumulation of monitoring data. Learning through management, with the use of what is learned to guide future management actions, is definitive of adaptive management (Williams et al., 2007).

3. Derivation of the expected value of perfect information

Q-functions. In developing the expected value of perfect information (EVPI) we use the concept of a “Q-function” representing the expected reward for taking action a at time t in state i and then acting optimally over the remainder of the timeframe (Kaelbling et al., 1998). Under structural certainty, the Q-function is the bracketed term in Eq. (4),

$$Q_t(a|i) = R(a|i) + \sum_j P(j|i, a)V_{t+1}[j], \tag{14}$$

and its optimization produces the optimal value function:

$$V_t[i] = \max_a Q_t(a|i). \tag{15}$$

A process Q-function under structural uncertainty utilizes model-specific Q-functions. Let

$$Q_t^k(a|i) = R_k(a|i) + \sum_j P_k(j|i, a)V_{t+1}^k[j], \tag{16}$$

denote the Q-function for model k , representing the expected reward under model k for taking action a at time t in state i , and then acting optimally over the remainder of the timeframe. A pro-

cess Q-function can be obtained by averaging the model-specific Q-functions:

$$\begin{aligned} Q_t(a|i, q) &= \sum_k q(k)Q_t^k(a|i) \\ &= \bar{R}(a|i, q) + \sum_k q(k) \left\{ \sum_j P_k(j|i, a)V_{t+1}^k[j] \right\}. \tag{17} \\ &= \bar{R}(a|i, q) + \sum_j \bar{P}(j|i, a, q)V_{t+1}[j, q'] \end{aligned}$$

Thus, the process Q-function is simply the bracketed term in Eq. (13), and its optimization produces the process optimal value function:

$$V_t[i, q] = \max_a Q_t(a|i, q). \tag{18}$$

Expected value of perfect information. The Q-function provides an intuitive way to derive EVPI. For a particular model k , consider the loss in value that comes from the use of action a

$$\begin{aligned} l_t^k(a|i) &= Q_t^k(a^*|i) - Q_t^k(a|i) \\ &= V_t^k[i] - Q_t^k(a|i) \end{aligned} \tag{19}$$

with a^* the optimal action for model k when the process is in state i at time t . Losses are non-negative for different actions, and the loss vanishes when $a = a^*$. The average loss over the model state is

$$\begin{aligned} \bar{l}_t(a|i, q) &= \sum_k q(k)l_t^k(a|i) \\ &= \sum_k q(k)[V_t^k[i] - Q_t^k(a|i)] \\ &= \sum_k q(k)V_t^k[i] - \sum_k q(k)Q_t^k(a|i) \end{aligned} \tag{20}$$

and the smallest average loss over the actions is

$$\begin{aligned} \min_a \bar{l}_t(a|i, q) &= \min_a \left\{ \sum_k q(k)V_t^k[i] - \sum_k q(k)Q_t^k(a|i) \right\} \\ &= \sum_k q(k)V_t^k[i] - \max_a \sum_k q(k)Q_t^k(a|i) \\ &= \sum_k q(k)V_t^k[i] - V_t[i, q] \end{aligned} \tag{21}$$

The latter form defines the expected value of perfect information:

$$EVPI(i, q) = \sum_k q(k)V_t^k[i] - V_t[i, q]. \tag{22}$$

From Eq. (21), $EVPI(i, q)$ can be seen as expressing the smallest average loss in value associated with model state q when the system is in system state i . In essence, the presence of uncertainty (as expressed by the model state q) leads to an expected loss in value that is at least as much as indicated by the EVPI metric. Alternatively, EVPI can be interpreted as a measure of the “importance” of eliminating that uncertainty.

4. Properties of EVPI

EVPI is essentially a comparison of 2 terms that utilize expectation and optimization of the model value functions, but in reverse order. Thus, the summation term $\sum_k q(k)V_t^k[i]$ in Eq. (22) consists of (i) optimization over the available actions to produce model-specific value functions, followed by (ii) an expectation that averages these functions over the model state. On the other hand,




Fig. 1. Value functions for a structurally uncertain process involving 2 models with model state $[q(1), q(2)] = [\tilde{q}, 1 - \tilde{q}]$, where $\tilde{q} \in [0, 1]$. A value function $V_t(a|i, \tilde{q}) = \tilde{q}V_t^1(a|i) + (1 - \tilde{q})V_t^2(a|i)$ is defined for each of 4 actions. An optimal partition of belief space is defined by the intersection points \tilde{q}^* and \tilde{q}^{**} , with the optimal value function shown in bold. The average of the 2 model-specific optimal value functions for $\tilde{q} \in [0, 1]$ is shown with a dashed line. The expected value of perfect information is the distance between the bold and dashed lines.

the term $V_t[i, q]$ is the result of (1) an expectation that averages the model-specific value functions over the model state, followed by (2) optimization over the available actions, as in Eq. (13). EVPI is simply the difference between these 2 treatments of the model value functions. That it is non-negative follows from the fact that the average of functional maxima is necessarily greater than the maximum of an average of functional values (Williams et al., 2002).

The geometry of EVPI is inherited from the terms in Eq. (22). Following the same arguments that demonstrate the optimal value function of a partially observable Markov decision process is piecewise linear convex in belief state (Sondik, 1971; Smallwood and Sondik, 1973), the optimal value function $V_t[i, q]$ under structural uncertainty also can be seen to be piecewise linear convex in model state (Williams, 2010a). Since the summation $\sum_k q(k)V_t^k[i]$ is linear in q , EVPI, which is simply the difference between the two value expressions, is piecewise linear concave in model state (concavity follows from negation of the convex value function in EVPI).

Piecewise linearity simplifies the interpretation of the value of information. For example, Fig. 1 displays the optimal value function for a problem involving 2 models and 4 actions. Because there are 2 models, the model state can be expressed by a single confidence value \tilde{q} with $q(1) = \tilde{q}$ and $q(2) = 1 - \tilde{q}$. From the figure optimization leads to a policy with actions that are specific to the model state, with a partition of the model space into regions over which the same action is optimal. Fig. 1 also exhibits the average of the




Fig. 2. Expected value of perfect information as a function of model state \tilde{q} . EVPI is piecewise linear concave in the model state $\tilde{q} \in [0, 1]$. Maximum EVPI is in the interior of the $[0, 1]$ interval, with convergence to 0 at the endpoints.

model-specific optimal values as a function of the model state. The vertical distance between the optimal value function and the average of model-specific optimal values represents the expected value of perfect information, as shown in Fig. 2.

Some patterns in EVPI are evident from Figs. 1 and 2. One is that EVPI vanishes as the confidence in any model approaches unity (e.g., as \tilde{q} converges to 0 or 1 in Fig. 2). Another is that EVPI is maximal for a model state somewhere in the interior of the model space (e.g., for \tilde{q} distant from 0 and 1). Yet another is that EVPI changes linearly within the interior of a partition region, but not across regions (EVPI is linear on either side of an intersection point, but not across that point). In essence, EVPI exhibits a constant gradient for all model states within a partition segment. Finally, EVPI is specific to system state, in that the value of information recorded in EVPI varies depending on the state at which it is evaluated (Fig. 3). The dependence on system state is clear from the notation in Eq. (22), where state i is shown as an argument in $EVPI(i, q)$.

These same patterns are evident for an optimal value function involving 3 models. As shown in Fig. 4, the model space is a two-dimensional simplex within which model states can vary. From the figure, EVPI is seen to be piecewise linear in the two-dimensional model state. It goes to zero at the 3 corners of the state space where confidence in one of the models is unity, and it is maximal somewhere in the interior of the state space. Finally, the optimal value function partitions state space into regions within which the function changes linearly. In fact, an examination of the terms in Eq. (22) makes clear that these same properties hold irrespective of the number of models under consideration. Because $V_t[i, q]$ is piecewise linear and convex for any time t in the timeframe (Williams, 2010a), the properties also hold irrespective of time t .

5. Example

We consider the value of information for an adaptive management project to promote the recovery of Florida scrub-jay (*Aphelocoma coerulescens*) populations in Brevard County, Florida, USA. The scrub-jay is Florida’s only endemic bird species, and it




Fig. 3. Optimal value function $V_t(a|i, \tilde{q}) = \tilde{q}V_t^1(a|i) + (1 - \tilde{q})V_t^2(a|i)$ for different system states. (a) Value function for state $i=1$. (b) Value function for state $i=2$.




Fig. 4. Value functions for a structurally uncertain process with model state $[q(1),q(2),q(3)]$. Model space consists of the 2-dimensional simplex defined by $q(1)+q(2)+q(3)=1$. The optimal value function is piecewise linear convex, and the average of the 3 model-specific optimal value functions is linear. The expected value of perfect information is the distance between the linear and piecewise linear forms. An optimal partition of belief space is defined by the intersection points and lines.

is an indicator for the rare and endangered Florida scrub ecosystem (Noss et al., 1997). Scrub habitat (and by extension scrub-jays, which are restricted to this ecosystem) is threatened by fragmentation and the disruption of natural succession patterns. Scrub habitat composition and density were maintained historically by fire, but fragmentation and active fire suppression have resulted in unchecked scrub succession to the point where scrub-jays are unable to sustain viable populations (Breininger and Carter, 2003). Land managers therefore rely on prescribed fires and mechanical thinning of overgrown vegetation to produce conditions believed to be optimal for maintaining stable scrub-jay populations (Johnson et al., in press).

Management actions are performed at the scale of individual fire management units (FMU) within a management reserve. Reserve managers have expressed a desire to base annual decisions on a combination of habitat and occupancy states of scrub-jay territories in an FMU. Florida scrub-jays are highly territorial, with a family (generally a breeding pair and one or more juvenile ‘helpers’) defending a territory of approximately 10-ha when the landscape is at carrying capacity (Carter et al., 2006; Breininger et al., 1995). To incorporate the scales of both management and the biological system, we define habitat and occupancy states on the basis of a regular grid of 10-ha cells within an FMU.

If an FMU contains N 10-ha scrub-jay territory patches, its state is described at the beginning of each management cycle by the number of territories in each of 4 possible succession stages – short (n_1), optimal height with open sandy patches (n_2), optimal height with no sandy patches (n_3), and tall (n_4) – and the number of territories occupied by scrub-jays (c), which can take any value from zero to the total number of territories. At any given time step, each territory of the FMU can be in only one of the 4 habitat types and is either occupied or not. The number of possible states an FMU can take is the product of the combination of 4 habitat types distributed across N territories and the number of territories that are occupied. Let $\underline{N} = [n_1, n_2, n_3, n_4]$ represent habitat state, consist-

ing of the number of territories of each type. The total count of all occupied territories is denoted by $c \in [0, 1, \dots, N]$.

Building directly on the work of Breininger and colleagues (unpublished data) we developed an integrated habitat-occupancy model in which habitat transitions are predicted as a function of management action, and occupancy dynamics depend on FMU habitat state. Annual changes in habitat state are modeled via a Markovian process, with probabilities $P(j|i, a)$ that a territory in state i' at a particular time t is in state j' at time $t+1$, given that action a is taken between t and $t+1$. Reserve managers hope to influence habitat transitions to favor scrub-jays by selecting among four possible actions each year: controlled burning, light mechanical cutting followed by a controlled burn, heavy mechanical cutting followed by a controlled burn, or a “do nothing” action. As in nearly all natural resource management problems, significant uncertainty exists in the habitat response expected after implementing any action. Uncertainty regarding the rate of scrub regeneration and the effects on future burning success of reducing highly combustible palmetto following restoration actions diminishes confidence in selecting the most appropriate management policies. We capture this uncertainty by considering multiple transition models that reflect differing beliefs regarding system dynamics. This model or ‘structural’ uncertainty is represented by competing versions of the state transition matrix model, in which the matrix elements $P_k(j'|i', a)$ describing the transition of a territory from state i' to state j' can vary among models. A Markov model for habitat state transitions is given by

$$E(\underline{N}_{t+1}) = \underline{N}_t \times \underline{P}_k(a), \tag{23}$$

where \underline{N} is a row vector of territory counts for each of the four habitat types and $\underline{P}_k(a) = [P_k(j'|i', a)]$ is a model-specific 4×4 matrix of model-specific transition probabilities for action a .

Annual probabilities of extinction and colonization for each territory are modeled as linear-logistic functions of the composition of anticipated habitat types (\underline{N}_t) in the FMU following implementation of a specified management action. Parameter coefficient values for these functions were derived from a 14-year data set containing observations on habitat state transitions and scrub-jay occupancy dynamics across 40 Brevard County scrub reserves (D. Breininger, unpub. data).

The annual state transitions for the integrated model are the combined transitions for the habitat and occupancy states. Letting i and j represent particular combinations of \underline{N} and c at times t and $t+1$, the probability of transition from i to j is expressed as

$$\begin{aligned} \tilde{P}_k(a) &= [\tilde{P}_k(j|i, a)] \\ &= [P_k(\underline{N}_{t+1}|\underline{N}_t, a)p(c_{t+1}|c_t, \underline{N}_{t+1}, a)] \end{aligned} \tag{24}$$

where c_{t+1} is predicted as a function of action a , model k and states \underline{N}_{t+1} and c_t .

The reward structure for this problem is model-specific and based on expected scrub-jay occupancy:

$$R_k(a_t|i) = E^k(c_{t+1}|a_t, i). \tag{25}$$

For the purpose of this example, we represent a portion of the uncertainty in system dynamics by considering 2 models that differed in their predictions of the effects on habitat transitions under the 4 available management alternatives. Model 1 hypothesized slower scrub growth of the optimal states in the absence of management, relative to Model 2. Model 1 also anticipated higher retention of palmetto scrub under burning, resulting in increase flammability and higher probability of remaining in a short state. However, Model 1 posited that light mechanical cutting resulted in a greater reduction of palmetto relative to Model 2, leading to

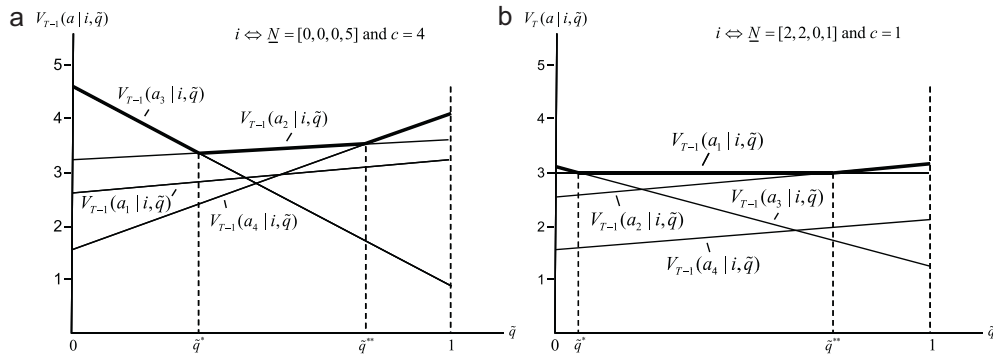


Fig. 5. Value functions for structural uncertainty represented by two models that predict the consequences of management actions on scrub habitat succession dynamics and Florida scrub-jay occupancy. System state (i) is the number of potential bird territories in each of 4 possible habitat types and the number of those territories that are occupied. (a) System state dominated by tall scrub, $N=[0,0,0,5]$, and 4 of 5 territories occupied. (b) Heterogeneous system with a mix of habitat succession states, $N=[2,2,0,1]$ and only 1 of 5 territories occupied by jays. The state-dependent differences in EVPI is evident from the geometry of the optimal piecewise linear value functions, with a greater gain expected by reducing uncertainty when the system is dominated by tall shrub.

reduced flammability and higher probabilities of transitioning to taller scrub classes. The value function at time $T - 1$ is

$$V(a_{T-1}|i) = \sum_k q(k)R_k(a_{T-1}|i) = \tilde{q}E^1(c_T|a_{T-1}, i) + (1 - \tilde{q})E^2(c_T|a_{T-1}, i) \tag{26}$$

with $\tilde{q} \in [0, 1]$. Fig. 5 displays value functions at $T - 1$ for two initial system states in an FMU of size $N=5$, for the 4 management actions. The system state with $\underline{N} = [0, 0, 0, 5]$ and $c = 4$ represents a predominance of tall vegetation in the FMU. For this state the partitioning of belief space is defined by the intersection at $\tilde{q}^* = 0.33$, where the optimal action switches from light mechanical cutting (action 3) to burning (action 2). This result is explained by the belief under Model 1 that reduced palmetto and increased growth rate moves post-treatment habitat more quickly into optimal conditions, whereas cutting under Model 2 greatly reduces growth relative to burning. The intersection at $\tilde{q}^{**} = 0.82$ reverts the optimal decision to heavy mechanical cutting (action 4) due to the hypothesis that heavy cutting is required to stimulate growth under Model 2. The ‘do nothing’ action (action 1) is dominated by the other actions across model space, and thus is not included in the optimal strategy. In comparison, the system state with $\underline{N} = [2, 2, 0, 1]$ and $c = 1$ represents earlier successional stages in 4 of the 5 territories in the FMU. Here, light mechanical cutting (action 3) is optimal only with strong belief in model 1 ($\tilde{q} = 0.0-0.06$) because of increased growth of optimal scrub classes. The “do nothing” option is optimal between $\tilde{q}^* = 0.06$ and $\tilde{q}^{**} = 0.76$, and the burn option (action 2) is optimal when there is confidence in model 2 ($\tilde{q} = 0.76-1.0$) which predicts suppression of palmetto by fire and thus higher rate of transition into optimal vegetation height. The geometries in Fig. 5 demonstrate that the potential average loss in value due to model uncertainty is dependent on system state, with a greater loss expected for the system state corresponding to $\underline{N} = [0, 0, 0, 5]$ and $c = 4$. This is seen by the greater vertical distance between the piecewise optimal value function and the average model-specific optimal value for $\underline{N} = [0, 0, 0, 5]$ and $c = 4$.

6. Extensions

Expected value of partial information. From Eqs. (19)–(22), EVPI is described in terms of a loss of value in the model-specific Q-function $Q_t^k(a|i)$, with each loss function based on a Q-function that is specified unconditionally. An alternative measure relaxes this requirement to obtain the expected value of partial information, sometimes written as EVPXI (Yokota and Thompson, 2004).

The underlying idea with EVPXI is that if there are multiple sources of uncertainty, one can describe a loss in value for some

uncertainty factors conditional on others. To illustrate, assume there are 2 sources of uncertainty, expressed by indices k and k' . For example, the index k might represent uncertainty about survival rates for a biological population, with k' representing uncertainty about recruitment. In this situation the model state is effectively a joint distribution $q(k, k')$ of the 2 indices, and it can be written in terms of conditional distributions as

$$q(k, k') = q(k|k')q(k') = q(k'|k)q(k) \tag{27}$$

Let q_k and $q_{k'}$ represent marginal distributions over the indices k and k' , with $q_{k|k'}$ and $q_{k'|k}$ the associated conditional distributions. By conditioning on k , one can compute an average of Q-functions over k' ,

$$Q_t^{jk}(a|i) = \sum_{k'} q(k'|k)Q_t^{k,k'}(a|i) \tag{28}$$

thereby incorporating uncertainty about recruitment into the Q-function for a given survival model k . Letting a^* represent the action that maximizes $Q_t^{jk}(a|i)$, the loss corresponding to a less than optimal action is

$$l_t^{jk}(a|i) = Q_t^{jk}(a^*|i) - Q_t^{jk}(a|i) = V_t[i, q_{k'|k}] - \sum_{k'} q(k'|k)Q_t^{k,k'}(a|i) \tag{29}$$

Then the average loss over the marginal model state q_k is

$$\begin{aligned} \bar{l}_t(a|i, q_k) &= \sum_k q(k)l_t^{jk}(a|i) \\ &= \sum_k q(k) \left\{ V_t[i, q_{k'|k}] - \sum_{k'} q(k'|k)Q_t^{k,k'}(a|i) \right\}, \tag{30} \\ &= \sum_k q(k)V_t[i, q_{k'|k}] - \sum_k \sum_{k'} q(k, k')Q_t^{k,k'}(a|i) \end{aligned}$$

and the smallest average loss over the actions is

$$\begin{aligned} \min_a \bar{l}_t(a|i, q) &= \min_a \left\{ \sum_k q(k)V_t[i, q_{k'|k}] - \sum_k \sum_{k'} q(k, k')Q_t^{k,k'}(a|i) \right\} \\ &= \sum_k q(k)V_t[i, q_{k'|k}] - \max_a \sum_k \sum_{k'} q(k, k')Q_t^{k,k'}(a|i) \tag{31} \\ &= \sum_k q(k)V_t[i, q_{k'|k}] - V_t[i, q] \end{aligned}$$

The latter form defines the expected value of partial information:

$$EVPXI(i, q_k) = \sum_k q(k) V_t[i, q_{k|k}] - V_t[i, q]. \tag{32}$$

Thus, EVPXI expresses the minimum loss in value associated with q_k , except in this instance the loss function for the value of information incorporates residual uncertainty about k' . In the example, EVPXI averages over the distribution of survival rates, using a loss function that includes uncertainty about recruitment. $EVPXI(i, q_k)$ has the same general computing form as $EVPI(i, q)$ except it uses a different loss function, one that includes the conditional distribution $q_{k'|k}$. Thus, $EVPXI(i, q_k)$ is “conditional,” in that the value of information it expresses will vary depending on the form of this distribution. Because the terms in Eq. (32) are piecewise linear, the expected value of partial information is as well.

The expected value of partial information offers a way to address system complexity that might otherwise be problematic for adaptive optimization. By investigating the components of uncertainty with EVPXI, one can compare the relative importance of different uncertainty sources and use the comparison to decide which uncertainty factor is most important to focus on with adaptive management.

Expected value of sample information. It often is useful to consider the gain in value that could result from a sample of observations that produce less than perfect information. This “expected value of sample information” is sometimes written as EVSI (Yokota and Thompson, 2004)

To simplify notation, consider a situation at time t in which the system state is i and model state is q . Optimization as in Eq. (13) produces $V_t[i, q]$, along with a stochastic transition to a new system state j and an updated model state q' at time $t+1$. Now imagine that somehow one could know the new system state j and updated model state q' at time t . With this new information, optimization could then produce $V_t[j, q']$. The difference between $V_t[i, q]$ and $V_t[j, q']$ represents the information value for the data that produced the updated system state j and derived model state q' . Of course, this difference is conditional on the state j that is produced. By averaging over the all system states using the average transition probabilities $\bar{P}(j|i, a^*, q)$, one produces the *expected* value of sample information:

$$\begin{aligned} EVSI(i, q) &= \sum_j \bar{P}(j|i, a^*, q) \{V_t[j, q'] - V_t[i, q]\} \\ &= \sum_j \bar{P}(j|i, a^*, q) V_t[j, q'] - V_t[i, q] \end{aligned} \tag{33}$$

Like EVPI, the expected value of sample information is a function of the system state i and model state q .

In the context of adaptive management, the key difference between EVSI and EVPI is seen by comparing Eqs. (22) and (33). Thus, EVPI in Eq. (22) uses an average of model-specific optimal values across model probabilities, and EVSI in Eq. (33) uses an average of weighted combinations of these values

$$\begin{aligned} \sum_j \bar{P}(j|i, a^*, q) V_t[j, q'] &= \sum_j \bar{P}(j|i, a^*, q) \sum_k q'(k) V_t^k[j] \\ &= \sum_j \sum_k P_k(j|i, a^*, q) q(k) V_t^k[j] \\ &= \sum_k q(k) \left\{ \sum_j P_k(j|i, a^*, q) V_t^k[j] \right\} \end{aligned} \tag{34}$$

The notation makes clear that EVSI and EVPI both are functions of model state q , and both are conditional on system state i . Like EVPI, the expected value of sample information is piecewise linear

in q , a result of the fact that the value terms in Eq. (33) are piecewise linear. However, EVSI produces different and somewhat more complicated patterns in the value of information.

7. Discussion

The value of information is a potentially useful but infrequently utilized construct in natural resources management. The few examples documented in the literature typically are limited to decision making at a single point in time, and do not account formally for the future impacts of decisions (see Runge et al., 2011 for a recent example). Yet a great many important problems in natural resources management are fundamentally sequential in nature, with the opportunity to improve resource understanding as decisions are made through time.

The value of information as discussed here links naturally to economics through the framing of VOI in terms of potential losses, including opportunity costs, and potential benefits to be accrued through the reduction of uncertainty. It is straightforward to frame the decision making problem as aggressively seeking to eliminate uncertainty when the benefit for doing so exceeds the cost, but adopting less informative strategy when the cost in information is high relative to benefits.

We have discussed 3 different measures for the value of information, namely the expected value of perfect, partial and sample information. In each case the notation makes clear that the measure is a function of the model state, and is parameterized by the system state. The value of information will vary with the amount of uncertainty (i.e., the model state), and decline to zero as the most appropriate model is recognized. It also will vary with the state of the resource – for a given model state the value of information may be high in one resource state and low in another (Fig. 5 and the above example).

The value of information will exhibit general trends through time. Because adaptive management tends to identify the most appropriate model as decision making progresses, the model state can be expected to converge to a unit vector representing that model. In turn, the value of information is expected to decline as uncertainty is reduced. One implication is that the value of eliminating any residual uncertainty is expected to decrease over time.

A few points about differences among the measures of information value are noteworthy. One is that the expected values of partial information do not sum to the expected value of perfect information. For example, if uncertainty is represented by the indices k and k' , then the measures $EVPXI(i, q_k)$ and $EVPXI(i, q_{k'})$, representing the expected values of partial information for each uncertainty factor, need not sum to the value of perfect information:

$$EVPXI(i, q_k) + EVPXI(i, q_{k'}) \neq EVPI(i, q). \tag{35}$$

On reflection this is expected; the 2 expected values of partial information are constructed with the conditional distributions $q_{k|k'}$ and $q_{k'|k}$, whereas the expected value of perfect information is constructed with a joint distribution over both uncertainty factors. Unless k and k' are independent, the partial distributions do not relate in a straightforward way to the joint distribution:

$$q(k|k')q(k'|k) \neq q(k, k'). \tag{36}$$

For that reason alone it seems reasonable that the 2 EVPXI values would not necessarily relate to EVPI in a straightforward way.

A second point of comparison concerns the value of sample information. From Eqs. (22) and (34) it is clear that the only substantive difference between EVPI and EVSI are the multiplicands for $q(k)$ in the 2 expressions. Thus, EVPI uses model-specific value functions, whereas EVSI uses a linear combination of these functions. The net effect of this difference is that EVPI is effectively an upper

bound on EVSI for any particular system and model state (Yokota and Thompson, 2004).

A possible application of the value of information for adaptive management involves a decision about whether to conduct monitoring or not (McDonald-Madden et al., 2010). The amount of monitoring (including the possibility of foregoing monitoring altogether) can be incorporated in adaptive management as a decision to be made sequentially. For example, a decision about whether to monitor might depend on a combination of resource conditions and the amount of uncertainty, with monitoring deferred if the resource is abundant and well understood (Hauser et al., 2006). One of the metrics for the value of information could be used in such an assessment to decide if resource understanding is adequate to defer monitoring, through a comparison of the value gained by additional monitoring against the cost of acquiring the information. In the context of adaptive decision making, a decision to forego monitoring at a given time step could be handled naturally, simply by using the information about system state from the previous time for current decision making. In consequence, the model state from the previous time would remain unchanged, i.e. the same model state would be used in successive time periods.

Finally, a natural extension of the above treatment of VOI would focus on systems that include both partial observability and structural uncertainty (Williams, 2009). Assuming partial observability but complete certainty as to model form, a VOI function that is analogous to Eq. (22) would be a function of belief state only:

$$EVPI(b) = \sum_i b(i)V_t[i] - V_t[b], \quad (37)$$

where the “belief state” b represents a distribution of likelihoods over the possible system states, $V_t[i]$ is the optimal value function corresponding to state i , and $V_t[b]$ is based on the average of state-specific value functions (Williams, 2009). A measure of the true value of information for processes that include both forms of uncertainty should account for both uncertainty components. Useful investigations might include comparative assessments of a “partial value of information” from including one source of uncertainty but not the other, and a comparison of these partial values against the value of information resulting from the inclusion of both.

References

- Bertsekas, D.P., 1995. *Dynamic Programming and Optimal Control*, vol. 1. Athena Scientific, Belmont, MA.
- Bontems, P., Thomas, A., 2000. Information value and risk premium in agricultural production: the case of split nitrogen application for corn. *American Journal of Agricultural Economics* 82, 59–70.

- Breining, D.R., Carter, G.M., 2003. Territory quality transitions and source-sink dynamics in a Florida Scrub-Jay population. *Ecological Applications* 13, 516–529.
- Breining, D.R., Larson, V.L., Duncan, B.A., Smith, R.B., Oddy, D.M., Goodchild, M., 1995. Landscape patterns in Florida Scrub Jay habitat preference and demography. *Conservation Biology* 9, 1442–1453.
- Carter, G.M., Stolen, E.D., Breining, D.R., 2006. A rapid approach to modeling species-habitat relationships. *Biological Conservation* 127, 237–244.
- Eidsvik, J., Bhattacharjya, D., Mukerji, T., 2008. Value of information of seismic amplitude and CSEM resistivity. *Geophysics* 73, R59–R69.
- Frauentorfer, K., 1992. On the value of perfect information and approximate solutions in convex stochastic two-stage optimization. *System Modelling and Optimization* 180, 564–573.
- Hauser, C.E., Pople, A.R., Possingham, H.P., 2006. Should managed populations be monitored every year? *Ecological Applications* 16, 807–819.
- Johnson, F.A., Breining, D.R., Duncan, B.W., Nichols, J.D., Runge, M.C., Williams, B.K. A Markov decision process for managing habitat for Florida Scrub-Jays. *Journal of Fish & Wildlife Management*, in press.
- Kaelbling, L.P., Littman, M.L., Cassandra, A.R., 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101, 99–134.
- Karnon, J., 2002. Planning the efficient allocation of research funds: an adapted application of a non-parametric Bayesian value of information analysis. *Health Policy* 61, 329–347.
- Koerkamp, B.G., Myriam Hunink, M.G., Stijnen, T., Weinstein, M.C., 2006. Identifying key parameters in cost-effectiveness analysis using value of information: a comparison of methods. *Health Economics* 15, 383–392.
- Lee, P.M., 1989. *Bayesian Statistics: An Introduction*. Edward Arnold, London, UK.
- McDonald-Madden, E., Baxter, P.W.J., Fuller, R.A., Martin, T.G., Game, E.T., Montambault, J., Possingham, H.P., 2010. Monitoring does not always count. *Trends in Ecology and Evolution* 25, 547–550.
- Noss, R.F., O’Connell, M.A., Murphy, D.D., 1997. *The Science of Conservation Planning: Habitat Conservation under the Endangered Species Act*. Island Press, Washington, DC, USA.
- Puterman, M.L., 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley and Sons, New York, NY.
- Raiffa, H., Schlaifer, R.O., 1961. *Applied Statistical Decision Theory*. Graduate School of Business Administration, Harvard University, Cambridge, MA, USA.
- Runge, M.C., Converse, S.J., Lyons, J.E., 2011. Which uncertainty? Using elicitation and expected value of information to design an adaptive program. *Biological Conservation* 144, 1214–1223.
- Smallwood, R.D., Sondik, E.J., 1973. The optimal control of partially observable Markov processes over a finite horizon. *Operations Research* 21, 1071–1088.
- Sondik, E.J., 1971. *The optimal control of Markov decision processes*. Ph.D. Thesis. Stanford University, Palo Alto, CA.
- Walters, C.J., 1986. *Adaptive Management of Renewable Resources*. Blackburn Press, Caldwell, NJ.
- Williams, B.K., 2007. Optimal management of non-Markovian biological populations. *Ecological Modelling* 200, 234–242.
- Williams, B.K., 2009. Markov decision processes in natural resources management: observability and uncertainty. *Ecological Modeling* 220, 830–840.
- Williams, B.K., 2010a. Resolving structural uncertainty in natural resources management using POMDP approaches. *Ecological Modelling* 222, 1092–1102.
- Williams, B.K., 2010b. Adaptive management of natural resources: framework and issues. *Journal of Environmental Management*, 1–8.
- Williams, B.K., Nichols, J.D., Conroy, M.J., 2002. *Analysis and Management of Animal Populations*. Academic Press, San Diego, CA.
- Williams, B.K., Szaro, R.C., Shapiro, C.D., 2007. *Adaptive Management: The U.S. Department of the Interior Technical Guide*. Adaptive Management Working Group, U.S. Department of the Interior, Washington, DC.
- Yokota, F., Thompson, K.M., 2004. Value of information literature analysis: a review of applications in health risk management. *Medical Decision Making* 24, 287–298.