A Cost-Optimization Framework for Planning Applied Environmental Science

MATTHEW R. FALCY

Natural resource managers need information about the status and structure of complex environmental systems to meet society's dual demand for natural resource use and conservation. Applied scientists supply this information through monitoring and research efforts that can be expensive. The relationship between increasing investment in information and the attendant decline in probability of making bad decisions can be optimized to ensure that applied environmental research is cost effective. In the present article, I use straightforward analytical and numerical methods to explore conditions in which the value of information is less than the cost of obtaining it. As pressures on Earth's natural resources grow, managers will be called on increasingly to do more with less, which can be achieved with formal optimization. Implementing such optimizations will help ensure that formal decision-making tools are used to translate science into decisions, which will often require collaboration between biological scientists and social scientists.

Keywords: value of information, decision-making, environmental monitoring, optimization

The importance of understanding the connections between natural biophysical systems and civil governance (Ehrlich et al. 2012) grows with human population size and per capita consumption. Municipal, state, and federal government officials must make difficult decisions about the use, allocation, and conservation of natural resources. These decisions are guided in part by knowledge of the status and structure of natural systems, as is revealed by applied scientific monitoring and research. It would be counterproductive, even ironic, if applied science initiatives intended to support environmental decision-makers were inefficient or wasteful of financial resources. However, when researchers, managers, and bureaucrats do not communicate well, data may accumulate without ever affecting decisions (Field et al. 2007, Lindenmayer and Likens 2010).

The gap between publications and public actions is well known to many conservation biologists (Anonymous 2007, Arlettaz et al. 2010). Although improving communication is a necessary first step, simply getting people together does not ensure efficient evidence-based decision-making. To avoid improvident spending of public money, a formal optimization framework is needed. A formal optimization framework will not only help ensure that the right amount of effort is spent on the acquisition of new scientific information, but it will also entreat practitioners to implement decision-theoretic approaches (Westgate et al. 2013).

Statistical power analysis is a form of optimization. Power analyses are used by researchers to determine the sample size that is needed to detect a treatment effect with a specified degree of confidence. This is optimization in the sense that there is a quantitative method to determine when too much or too little is spent on data collection. However, power analysis may have little utility in applied environmental sciences because they are based on a priori specification of the magnitude and variance of a treatment effect, but estimating treatment effects can be a primary goal of applied science. Furthermore, statistical hypothesis tests that power analyses support are often inadequate in applied settings because they do not consider the costs of alternative actions (Mapstone 1995, Johnson 1999). A more effective way to plan applied environmental science initiatives is to know what the decision options are and how they depend on the status (i.e., population size, contaminant load) and structure (i.e., population dynamics, biogeochemical cycles) of the natural system. Just as Platt (1964) emphasized the importance of focusing biological research on discrimination among competing a priori hypotheses to improve the rate of learning, Nichols and Williams (2006) note that formally integrating decision-making with biological monitoring improves management efficiency.

Most decision-making methods require an explicit statement about ultimate values and objectives (Keeney 1992).

BioScience 68: 912–922.
doi:10.1093/biosci/biy109
Advance Access publication 10 October 2018

Because values and objectives are often intangible, it can be difficult to derive explicit decision criteria. However, failure to explicitly define criteria does not make decision points go away, just as the absence of an explicit model does not imply robustness to the unstated assumptions. Instead, the unstated becomes a hidden psychophysical reality of the decision-maker that manifests itself as a gut feeling. Making decisions with uncertainty using unaided intuition is problematic because several decades of psychological research have identified numerous cognitive biases that can negatively affect decisions (Tversky and Kahneman 1974, Denes-Raj and Epstein 1994, Hilbert 2012, Iftekhar and Pannell 2015). Uncertainty about the status of nature and the effects of management never go away, but uncertainty and perceived risks can be incorporated directly into decision-making models (Regan et al. 2005, Addison et al. 2013, Falcy 2016), so that decisions are more coherent, transparent, replicable, and free from bias.

In the present article, I describe a cost-optimization framework for planning future environmental science and emphasize conditions in which it is optimal to spend nothing on new data. A complete theoretical framework for applied environmental science would extend beyond this, including many other decision-making methods that can be applied to existing data (e.g., risk analysis, Burgman 2005, Hope 2006) and methods to cope with decisions with no data or severe uncertainty (e.g., info-gap decision theory, Ben-Haim 2006, but see also Hayes et al. 2013). An introduction to the elements of decision-making that constitute a broader theoretical framework guiding applied environmental science are given by Gregory and colleagues (2012) and Conroy and Peterson (2013).

I focus on evaluating the cost effectiveness of acquiring new information for two reasons. First, failure to integrate costs and benefits can lead to wasted efforts and missed opportunities. Second, and perhaps more important, being explicit about supposed costs and benefits during the research or monitoring design phase will beg questions that lead into the broader set of decision-analysis tools intended to translate science into decisions. Therefore, implementing the cost optimizations described in the present article will help ensure that the results of research and monitoring have unambiguous applications.

Evaluating the cost effectiveness of acquiring scientific information is an emerging framework for prioritizing and evaluating the usefulness of applied environmental science (Naidoo et al. 2006, Bode et al. 2008, Chadès et al. 2008, McDonald-Madden et al. 2010, Fackler and Haight 2014, Canessa et al. 2015, Maxwell et al. 2015, Williams and Johnson 2015). The leading-edge research on cost optimization of information is often mathematically sophisticated, which can prevent appreciation by practitioners that do not regularly work with math. I illustrate the cost optimization of scientific information using generic examples to highlight core concepts without being encumbered by the nuances of the real world and the attendant mathematical complexity. Although intentionally abstract, the examples will cover fundamental concepts that that have broad applications in the real world. Box 1 provides three real-world problems that leverage the concepts described below.

Example 1: A formula for optimal expenditure on information

Just as increasing sample size decreases the risk of overlooking treatment effects that truly exist, let us assume that the probability of making a bad decision declines exponentially at a rate r, with increasing investment in scientific information (e.g., data collection), x. An exponential decline in the probability of making a management mistake with increasing expenditure on information will not be strictly accurate if sampling startup costs are high (e.g., fixed costs of equipment), and the precise value of r will depend on several features of the system (e.g., complexity and adequacy of approximating models). However, the exponential-decline function e^{-rx} is relatively simple, will apply to many situations, and will be replaced with a more complex function in the next example.

We can multiply the exponential-decline function into the probability of making a management mistake in the absence of data, p. If no money is spent on data collection ($x = 0$), then $e^{-rx} = 1$ for any value of r, and so $pe^{-rx} = p$. This essentially recovers the definition of p. However, when some amount of money is spent on information ($x > 0$), then the function pe^{-rx} returns a number less than p and declines exponentially as x increases. We can multiply the output of this function by the cost of the mistaken decision, m, to get the expected cost of an error. The word *expected* contains the notion that there is a fixed cost of a mistaken decision, m, that will occur with some probability, pe^{-rx}. The total cost is the cost of data collection plus the expected cost of a mistaken decision (Falcy et al. 2016):

$$\textit{total cost} = x + mpe^{-rx} \tag{1}$$

The *total cost* will often approximate a U shape over different levels of expenditure on information, x. The exponential decline in the probability of making a bad decision as sampling intensity increases is responsible for the descending left side of the U. However, as more and more money is spent on information, the reward of minimizing errors begins to reach an asymptote, but sampling still presents its own cost (first term on RHS of equation 1). This causes the ascending right side of the U shape of the *total cost* function. Because the cost of a bad decision is already built in, a fiscally responsible decision-maker should spend just enough on information so that the *total cost* function is minimized.

The minimum of the *total cost* function can be found with an analytical optimization. This is done by taking the first derivative of the *total cost* function with respect to x,

$$\frac{\partial}{\partial x}(x + mpe^{-rx}) = 1 - e^{-rx}mpr,$$

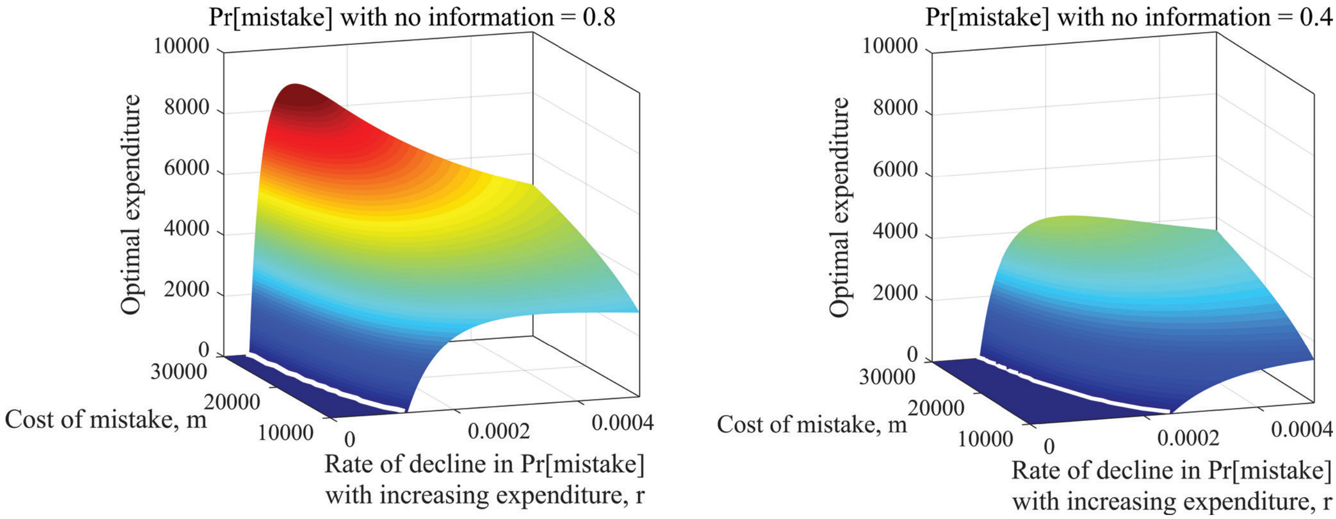

Figure 1. The optimal expenditure on information depends on the cost of mistaken decision, the rate of decline in the probability of making a mistaken decision with increasing expenditure on information, and the inherent probability of making a mistake in the absence of scientific information. The hotter colors represent higher optimal expenditure. The area to the left of the white curved line is a region of parameter space where the optimal expenditure on scientific information is zero because costs are not outweighed by benefit.

setting the right-hand side equal to zero, and then solving for x:

$$x^{opt} = \frac{\log(mpr)}{r}.$$

This analytical procedure finds the value of x, where a straight line making a tangent with the U shape has a slope of zero. This occurs at the trough of the U and therefore gives the minimal (optimal) expenditure on information given values for m, p, and r. The combinations of parameter values where the optimal expenditure on sampling is zero (figure 1) is noteworthy because any expenditure on information is wasteful. In the region of the parameter space where the optimal expenditure is zero, the gain never exceeds the costs. Unsurprisingly, not collecting any data is optimal when the effect of collecting more data results in a small decline in the probability of making a management mistake (small values for r). The parameter space where no money should be spent on data expands when the probability of making a management error with no information, p, decreases. The optimal expenditure can be relatively high when the probability of making a mistake with no information is also high but decreases dramatically when the probability of making a mistake with no information goes down. See Wintle and colleagues (2010) for an application of a closed-form analytical assessment of budget allocation to omnibus surveillance monitoring versus hypothesis-driven monitoring.

Example 2: Expected value of information

In many cases a decision-maker can make a management mistake in many different ways. For example, decisions often involve choosing among a suite of options that have different costs and benefits. Furthermore, the costs and benefits themselves can change depending on the status of an environmental system. In such a scenario, the decision-maker can use scientific information to infer the status of the system and then act accordingly.

Suppose that the state of nature is represented with a variable that can be either good or bad. An environmental decision-maker must make recurrent decisions about whether to take a conservative action. The decision could be to cease harvest or disallow the transport of a toxic substance. Let the conservative action be appropriate when the state of nature is bad (poor environmental conditions) but inappropriate if nature is good (robust environmental conditions). The ability to anticipate the state of nature would help the decision-maker. How much should she spend to acquire this knowledge? The answer depends, in part, on the payout associated with all four combinations of the two states of nature and the decision to act conservatively or not (table 1). If the decision-maker knew that the state of nature was good then she would not take a conservative action and receive 4 units of value. In contrast, if she knew that the state of nature was bad then she would act conservatively and receive 5 units of value. Let us suppose that our hypothetical decision-maker could determine from historical records that there is a .3 and .7 probability that nature could be in good and bad states, respectively. The decision-maker could then compute the expected payout of not acting conservatively ($.3 \times 4 + .7 \times -5 = -2.3$) and the expected payout of acting conservatively ($.3 \times -2 + .7 \times 5 = 2.9$). As a rational decision-maker, she would see that $2.9 > -2.3$ and would choose to always act conservatively and expect a long-term gain of 2.9 for each iteration of the decision. But if she could perfectly

Table 1. Decision-making payout matrix.

Human action	Nature: Good	Nature: Bad
Don't act	4	–5
Act	–2	5

Note: An environmental decision-maker must decide whether to implement a conservation action (rows) while not knowing the true state of nature (columns). The decision-maker wants to implement conservation actions only when the state of nature is bad. Positive and negative numerical entries in this payout matrix are the perceived values of making correct and incorrect decisions, respectively. The goal is to make the decision that maximizes the payout.

predict the state of nature, then she would never make an incorrect decision. With perfect information about the state of nature, she would expect to gain $.3 \times 4 + .7 \times 5 = 4.7$ units of value for each iteration of the decision. The difference between the payout with perfect ability to predict nature and the payout associated with rational evaluation of the past, $4.7 - 2.9 = 1.8$, is the maximum amount she should be willing to spend to acquire perfect information. This is known as expected value of perfect information (EVPI, Raiffa and Schlaifer 1961). Spending more than the EVPI is not rational because once the expense of information is subtracted from the payout, then the amount left over is less than the payout associated with the rational evaluation of the past, which is available at no cost.

Information is rarely perfect. Instead, knowledge gradually improves with more information. Let us now suppose that knowledge improves only slightly with minimal expenditure, then improves rapidly, and then finally reaches an asymptote such that further expenditure yields minimal improvements to knowledge. This can be modeled with a sigmoidal function, such as the Gompertz function,

$$k(x) = ae^{-be^{-cx}}, \qquad (2)$$

where k is the degree of perfection of knowledge, x is the expenditure on knowledge, a is an upper asymptote, b sets the initial lag in improved knowledge, and c is the rate of learning (figure 2a). Setting $a = 1$ yields a function for knowledge bounded between zero and one. When the knowledge function returns zero, the decision-maker has no ability to predict nature but still makes a rational decision on the basis of the frequencies of the states of nature found in an historical record. When the knowledge function returns one, the decision-maker has perfect ability to predict the state of nature and act accordingly. Now we must decide how much money to spend on knowledge. Given the payout described in table 1 and a knowledge function, $k(x)$, what is the optimal x? When is no expenditure ($x = 0$) optimal?

The payout from table 1 can be computed when the decision-maker chooses the correct action for the state of nature with a probability of k, and the rational choice given the historical record with probability $1 - k$ (see the supplemental material online). Subtracting the expense of knowledge, x, from the payout yields a function in which the optimal expenditure on knowledge, x^{opt}, occurs at the peak (figure 2b). For a .5 probability that nature is good and a knowledge function defined by $a = 1$, $b = 5$, $c = 2$, there is a slight increase in optimal expenditure from $x^{opt} = 1.6$ to $x^{opt} = 1.9$ when the values in the payout matrix in table 1 changes from $\begin{bmatrix} 4 & -5 \\ -2 & 5 \end{bmatrix}$ to $\begin{bmatrix} 6 & -7 \\ -4 & 7 \end{bmatrix}$ (figure 2b). However, if the knowledge function is defined by $a = 1$, $b = 6$, $c = 1$, then the optimal expenditure is $x^{opt} = 0$ when the payout is $\begin{bmatrix} 4 & -5 \\ -2 & 5 \end{bmatrix}$. The optimal expenditure then jumps up to $x^{opt} = 3.14$ when the payout changes to $\begin{bmatrix} 6 & -7 \\ -4 & 7 \end{bmatrix}$. The optimal expenditure on monitoring is nothing ($x^{opt} = 0$) when the EVPI is relatively low and the knowledge function is more shallow (b large and c small), meaning that information is relatively expensive. With a .5 probability that nature is good, the EVPI of the payout matrices $\begin{bmatrix} 4 & -5 \\ -2 & 5 \end{bmatrix}$ and $\begin{bmatrix} 6 & -7 \\ -4 & 7 \end{bmatrix}$ is 3 and 5, respectively. Therefore, increasing the EVPI can make the optimal expenditure on information become relatively large when the knowledge function is shallow. This may be somewhat counterintuitive. If the EVPI is ignored, a person may conclude that the optimal expenditure on knowledge should go down as it becomes more expensive to acquire knowledge. However, the EVPI is important and if it is sufficiently large and knowledge is expensive to acquire, then large expenditures on monitoring and research will be optimal.

The optimal expenditure on monitoring given a knowledge function and a payout matrix can be found with numerical optimization methods (see the supplemental material online) and then plotted as a function of the probability that nature is in a good state (figure 2c). The optimal expenditure on knowledge is zero when the probability that nature is good is either very high or very low. Under such conditions, there is little utility associated with being able to predict future states of nature, because deviation from the state most commonly observed in the historical record occurs rarely. It is not until the state of nature is likely to switch states that expenditure on information is warranted. Payout matrices with lower EVPI are associated with a narrower range of probabilities that nature is good where expense on knowledge is warranted. As was discussed under figure 2b, very high expenses on information are associated with relatively high EVPI and shallow knowledge functions.

Runge and colleagues (2011) and Williams and Johnson (2011) apply EVPI to specific conservation problems and discuss broader implications for linking applied science to decisions. See Maxwell and colleagues (2015) and Canessa and colleagues (2015) for more in-depth introductions to EVPI.

Example 3: Recurring decisions about dynamic systems

Many management scenarios involve a sequence of decisions through time. Recurrent decision-making in dynamic

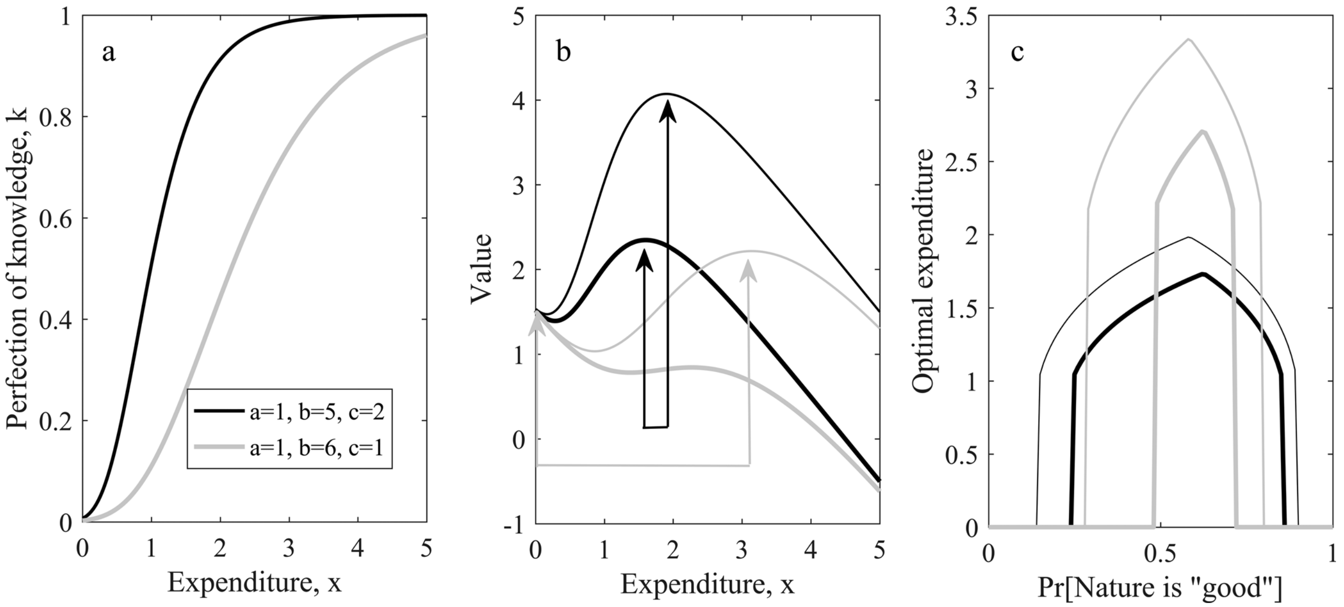

Figure 2. A Gompertz function can be used to model the relationships between expenditure on scientific information and the resulting degree of perfection in knowledge (a). The optimal expenditure on scientific information occurs when the value of the decision is maximized (b). The colors in (b) correspond to the knowledge functions in (a). The thick lines in (b) use the payout matrix given in table 1 $\begin{bmatrix} 4 & -5 \\ -2 & 5 \end{bmatrix}$ whereas thin lines use $\begin{bmatrix} 6 & -7 \\ -4 & 7 \end{bmatrix}$ in the payout matrix. Note how grey goes from spend nothing to spend a lot, whereas black only changes a little. The results in (b) assume that nature has a .5 probability of being in a good state. In (c), the optimal expenditure is plotted for different probabilities that nature is in a good state. The line colors and thicknesses in (c) are identical to those in (b).

systems is complicated by the fact that a decision at a given point in time will affect the future state of the system, which, in turn, will affect future decisions. Identifying an optimal sequence of decisions through time must consider the immediate payout at the current time as well as future consequences. Imagine running a marathon and making pace decisions every kilometer. One strategy would be to run each kilometer as fast as possible. However, sprinting the first kilometer will severely constrain the possible paces over subsequent segments of the race. We know that sprinting a marathon is a bad strategy because we have intimate knowledge of the limitations of our own bodies. However, when it comes to sequential decision making about other dynamic systems, human intuition can lead to decision policies that are analogous to sprinting a marathon. An analytical method for solving this kind of problem, known formally as a Markov decision process (MDP), is stochastic dynamic programming (SDP, Clark and Mangel 2000, Marescot et al. 2013). I will elaborate on how the solutions to such problems can yield information about cost-effective environmental management.

Consider a system in which each year there is some probability that an invasive species will arrive (γ_1), or habitat degrades (γ_2), or both (γ_1 γ_2). Once these problems emerge they do not go away unless a manager treats them. However, treatments are imperfect, with probabilities θ_1 and θ_2 of fixing the invasive and habitat problems, respectively. Furthermore, treatments have a cost to implement ($\$_1$ and $\$_2$ for invasive and habitat, respectively), and allowing the system to exist with problems also entails costs because the manager's duty is to maintain a healthy system. Let the annual costs of allowing the system to manifest invasion and habitat problems be $\$_3$ and $\$_4$ respectively, with the possibility that allowing both problems to exist simultaneously entails a cost that is greater than the sum of the parts $\$_5 = 1.25(\$_3 + \$_4)$ because the problems may interact synergistically. Assume the manager knows the state of the system perfectly (more on this later). The manager must decide among four actions: do nothing, treat for the invasive species, treat for degraded habitat, or treat for both the invasive and habitat. What is the optimal, state-dependent decision to take so that costs are minimized over 20 years? Is it possible that, even with perfect knowledge about the presence of a problem, doing nothing is part of an optimal policy?

This problem can be expressed compactly using two kinds of matrices: a transition matrix that gives probabilities of the system moving from one state to another under all possible actions and a reward matrix that gives the payout associated with all possible combinations of actions and states (box 2). These matrices can be analyzed with canned MDP solvers that relieve users from the need for custom-coding SDP solutions (Fackler 2011, Chadès et al. 2014). Therefore, the greatest impediment to solving an MDP today lies in

Box 1. Real environmental problems that leveraged concepts described in each hypothetical example.

The gypsy moth (*Lymantria dispar*; top; correspondence with example 1) is an exotic invasive forest pest in North America. Epanchin-Niell and colleagues (2012) used numerical simulations to optimize the tradeoff between surveillance effort and management costs. They demonstrated how the minimum total cost varies with biological parameters (e.g., establishment rate) and management parameters (e.g., sample costs). Importantly, Epanchin-Niell and colleagues (2012) found substantial long-term financial savings by accounting for spatial heterogeneity of invasion risk and surveillance cost. The whooping crane (*Grus americana*; middle; correspondence with example 2) is listed as endangered by the federal governments of Canada and the United States. Runge and colleagues (2011) worked with stakeholders to identify biological hypotheses and key elements of recovery decisions. They then used several formulations of expected value of perfect information (EVPI) to determine which hypotheses and actions are most fruitful to pursue further. Resolving the hypothesis that nest abandonment is caused by parasitic black flies was shown to provide more information than resolving all seven alternative hypotheses combined. Runge and colleagues (2011) also found that, if measurement error is relatively low, then killing black flies would generate the most actionable information for resolving future decisions, otherwise either the discontinuation of egg salvage or egg swapping was recommended. Mesquite (*Prosopis* spp.; bottom; correspondence with example 3) is an exotic invasive weed in western Australia. Pichancourt and colleagues (2012) developed a life cycle model with dispersal and add control actions targeting specific life stages. They allowed costs of control actions to vary with effort, and then used stochastic dynamic programming (SDP) to identify cost-optimal management strategies. Pichancourt and colleagues (2012) found simple and robust management rules, or heuristics, that depended on the underlying biological parameters (e.g., dispersal, relative abundance of juveniles and adults). They also identified the kinds of new knowledge that would be unimportant.

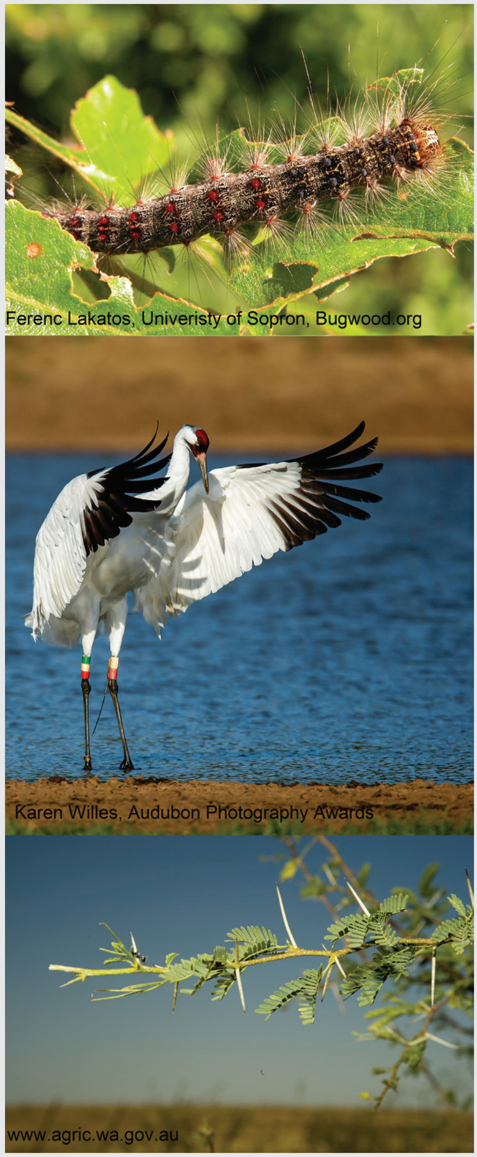

specifying the management problem as matrices (box 2). Specifying the management problem as matrices is itself a useful exercise because it will likely impose clarity on otherwise seemingly intractable problems.

Applying MDPToolbox (Chadès et al. 2014) to the invasive species and habitat management problem (see the supplemental material online) reveals that the optimal, state-dependent sequence of decisions can be quite sensitive to parameter settings (figure 3). When the probability of a species invasion and habitat degradation problem are relatively low (figure 3a), then the optimal action generally matches common sense: Only act when a problem is known to exist, and match the action to the known problem. However, toward the end of the decision process the action *do nothing* becomes part of an optimal solution even when problems are known to exist. System states beyond step 20 are not considered, so the immediate reward of cutting costs is not counterbalanced with a penalty for putting the system into a problematic state. Of course, a responsible manager should not manage for optimality only over his tenure but, rather, over much longer time horizons. A method known as *policy iteration* is available in MDP solvers and will overcome the problem of finite time horizons. In the finite time scenario of figure 3a, the specific time at which inaction becomes prudent even when a problem is known to exist (steps 19, 18,

Box 2. Markov decision problems.

Markov decision problems can be stated compactly in two kinds of matrices. The first kind of matrix is a transition matrix. The example below comes from example 3 of the main text. In this case, the transition matrix is three dimensional. The first dimension (rows) gives the state of the system at time t. The second dimension (columns) gives the state of the system at time t + 1. The third dimension repeats the first two dimensions for all possible actions (displayed vertically). Entries give the probabilities of transitioning from one state to another. Note that each row must sum to 1. Consider transitioning from *no problem* to having a habitat problem when the second action (treat for invasive) is undertaken (bold). This transition can happen in two ways: Either a habitat problem can arise (γ_2), and an invasive does not arise ($1 - \gamma_1$), or a habitat problem (γ_2) and an invasive problem (γ_1) both arise, and the invasive treatment is effective (θ_1). (See tables 2–5.)

Table 2. State transition matrix for action 1 (do nothing) at time t.

	Time t + 1			
	No problem	**Invasive (I)**	**Habitat (H)**	**I and H**
No problem	$(1 - \gamma_1)(1 - \gamma_2)$	$\gamma_1(1 - \gamma_2)$	$(1 - \gamma_1)\gamma_2$	$\gamma_1\gamma_2$
Invasive (I)	0	$(1 - \gamma_2)$	0	γ_2
Habitat (H)	0	0	$(1 - \gamma_1)$	γ_1
I and H	0	0	0	1

Table 3. State transition matrix for action 2 (treat for invasive) at time t.

	Time t + 1			
	No problem	**Invasive**	**Habitat**	**I and H**
No problem	$\gamma_1\theta_1(1 - \gamma_2)$	$\gamma_1(1 - \theta_1)(1 - \gamma_2)$	$\boldsymbol{\gamma_2(1 - \gamma_1) + \gamma_2\gamma_1\theta_1}$	$\gamma_1(1 - \theta_1)\gamma_2$
Invasive (I)	$\theta_1(1 - \gamma_2)$	$(1 - \theta_1)(1 - \gamma_2)$	$\theta_1\gamma_2$	$(1 - \theta_1)\ \gamma_2$
Habitat (H)	0	0	$(1 - \gamma_1) + \gamma_1\theta$	$\gamma_1(1 - \theta_1)$
I and H	0	0	θ_1	$(1 - \theta_1)$

Table 4. State transition matrix for action 3 (treat for habitat).

	Time t + 1			
Time t	**No problem**	**Invasive (I)**	**Habitat (H)**	**I and H**
No problem	$(1 - \gamma_1)\ \gamma_2\theta_2$	$\gamma_1\gamma_2\theta_2 + \gamma_1(1 - \gamma_2)$	$\gamma_2(1 - \theta_2)$	$\gamma_1\gamma_2(1 - \theta_2)$
Invasive (I)	0	$\gamma_2\theta_2 + (1 - \gamma_2)$	0	$\gamma_2(1 - \theta_2)$
Habitat (H)	$(1 - \gamma_1)\theta_2$	$\gamma_1\theta_2$	$(1 - \gamma_1)(1 - \theta_2)$	$\gamma_1(1 - \theta_2)$
I and H	0	θ_2	0	$(1 - \theta_2)$

Table 5. State transition matrix for action 4 (treat for habitat and invasive).

	Time t + 1			
Time t	**No problem**	**Invasive (I)**	**Habitat (H)**	**I and H**
No problem	$\gamma_1\theta_1\gamma_2\theta_2 + (1 - \gamma_1)(1 - \gamma_2) + \gamma_1\theta_1(1 - \gamma_2) + \gamma_2\theta_2(1 - \gamma_1)$	$\gamma_1(1 - \theta_1)(1 - \gamma_2) + \gamma_1(1 - \theta_1)\gamma_2\theta_2$	$\gamma_2(1 - \theta_2)(1 - \gamma_1) + \gamma_2(1 - \theta_2)\gamma_1\theta_1$	$\gamma_1(1 - \theta_1)\gamma_2(1 - \theta_2)$
Invasive (I)	$\gamma_2\theta_2\theta_1 + \theta_1(1 - \gamma_2)$	$(1 - \gamma_1)(1 - \theta_1) + \gamma_2\theta_2(1 - \theta_1)$	$\theta_1\gamma_2(1 - \theta_2)$	$(1 - \theta_1)\gamma_2(1 - \theta_2)$
Habitat (H)	$(1 - \gamma_1)\theta_2 + \gamma_1\theta_1\theta_2$	$\theta_2\gamma_1(1 - \theta_1)$	$\gamma_1\theta_1(1 - \theta_2) + (1 - \gamma_1)(1 - \theta_2)$	$\gamma_1(1 - \theta_1)\gamma_1(1 - \theta_1)$

Box 2. Continued.

The second kind of matrix is a reward matrix. The reward matrix specifies the payout associated with all possible state-action combinations. The reward matrix for example 3 is given below. $\$_1$ and $\$_2$ are the costs of treating invasive and habitat degradation, respectively. $\$_3$ and $\$_4$ are the costs of having invasive and habitat degradation problems, respectively. $\$_5$ is the cost of having both problems at once. All costs are negative numbers. Stochastic dynamic programming will maximize the reward (minimize costs). (See table 6.)

Table 6. Reward matrix.

	Action			
State	Do nothing	Treat invasive (I)	Treat habitat (H)	Treat I and H
No problem	0	$\$_1$	$\$_2$	$\$_1 + \$_2$
Invasive (I)	$\$_3$	$\$_1 + \$_3$	$\$_2 + \$_3$	$\$_1 + \$_2 + \$_3$
Habitat (H)	$\$_4$	$\$_1 + \$_4$	$\$_2 + \$_4$	$\$_1 + \$_2 + \$_4$
I and H	$\$_5$	$\$_1 + \$_5$	$\$_2 + \$_5$	$\$_1 + \$_2 + \$_5$

and 20 for states with both problems, habitat, and invasive, respectively) is related to the costs of treating the problems. When both problems exist, treating only the invasion becomes part of an optimal strategy because treating for a habitat problem is more expensive than treating an invasive species.

Seemingly slight alterations in parameter values of an MDP problem can lead to radically different decisions. Increasing the probability that an invasive species arrives (γ_1) or habitat degrades (γ_2) from $\gamma_1 = \gamma_2 = .2$ to $\gamma_1 = \gamma_2 = .4$ and decreasing the cost of having a degraded habitat from $\$_3 = -12$ to $\$_3 = -10$ causes the differences in optimal decisions visible across figure 3a and 3b. With increased probability of developing invasive species and habitat degradation problems, and slight decrease in the cost of having an invasive species problem, it becomes imprudent to simultaneously treat for both problems even when both problems are known to exist. Treating both problems simultaneously is not part of an optimal allocation of money because the problem will just reemerge and treatment efficacy is imperfect (θ). Removing the synergist effect of concurrently having problems with invasive species and habitat degradation ($\$_5$) makes treatment for habitat degradation become imprudent (compare figure 3b and 3c).

If seemingly slight alterations in parameter values of an MDP problem can lead to radically different decisions, then we might be willing to pay for monitoring or research activities that estimate the parameters. One method of evaluating financial investment in monitoring or research entails the application of EVPI principles described in example 2. Because MDP solvers find sequences of state-dependent decisions that minimize costs, the solvers can return the expected cost of managing a system while making optimal decisions. For example, if the system given in figure 3b starts in the state *no problems* and is managed with sequentially optimal decisions over 20 years, then the total expected management cost is –174.47 (note that the negative sign indicates a cost payed out, and is consistent with values $\$_1$, $\$_2$, … $\$_5$). However, if the system does not have synergistic effects of contemporaneous invasion and habitat degradation (figure 3c), then the expected cost of managing the system beginning with the state *no problems* will be –143.81. The difference 174.47 – 143.81 = 30.66 is the maximum amount of that should be spent on monitoring or research into synergism, assuming that synergism cannot exceed 125% of the sum of the parts.

Recent development in the analysis of partially observable Markov decision processes (POMDP, Chadès et al. 2017) provide another method of evaluating financial investment in monitoring or research. Unlike the MDP analyzed in the present article, a POMDP does not assume perfect knowledge of the system. Instead, a POMDP could consider the state of the system to be observed imperfectly, with probability distributions used to relate all possible pairs of states and actions to observations. The analysis of POMDP systems is more technical and recent applications have focused on adaptive management (Chadès et al. 2017). Explicit consideration of the financial costs of management in POMDP models (Fackler and Haight 2014) constitutes a leading edge in the assessment of applied environmental research and monitoring.

A role for social scientists

Most environmental decisions involve tradeoffs between things that are not obviously commensurable. How much extinction risk of an animal population are we willing to trade for harvest opportunity? This is a complex problem. Natural science can inform the relationship between extinction risk and harvest, but natural science cannot address how much risk is acceptable. Rather, it is a social question. Therefore, social scientists can be helpful partners in

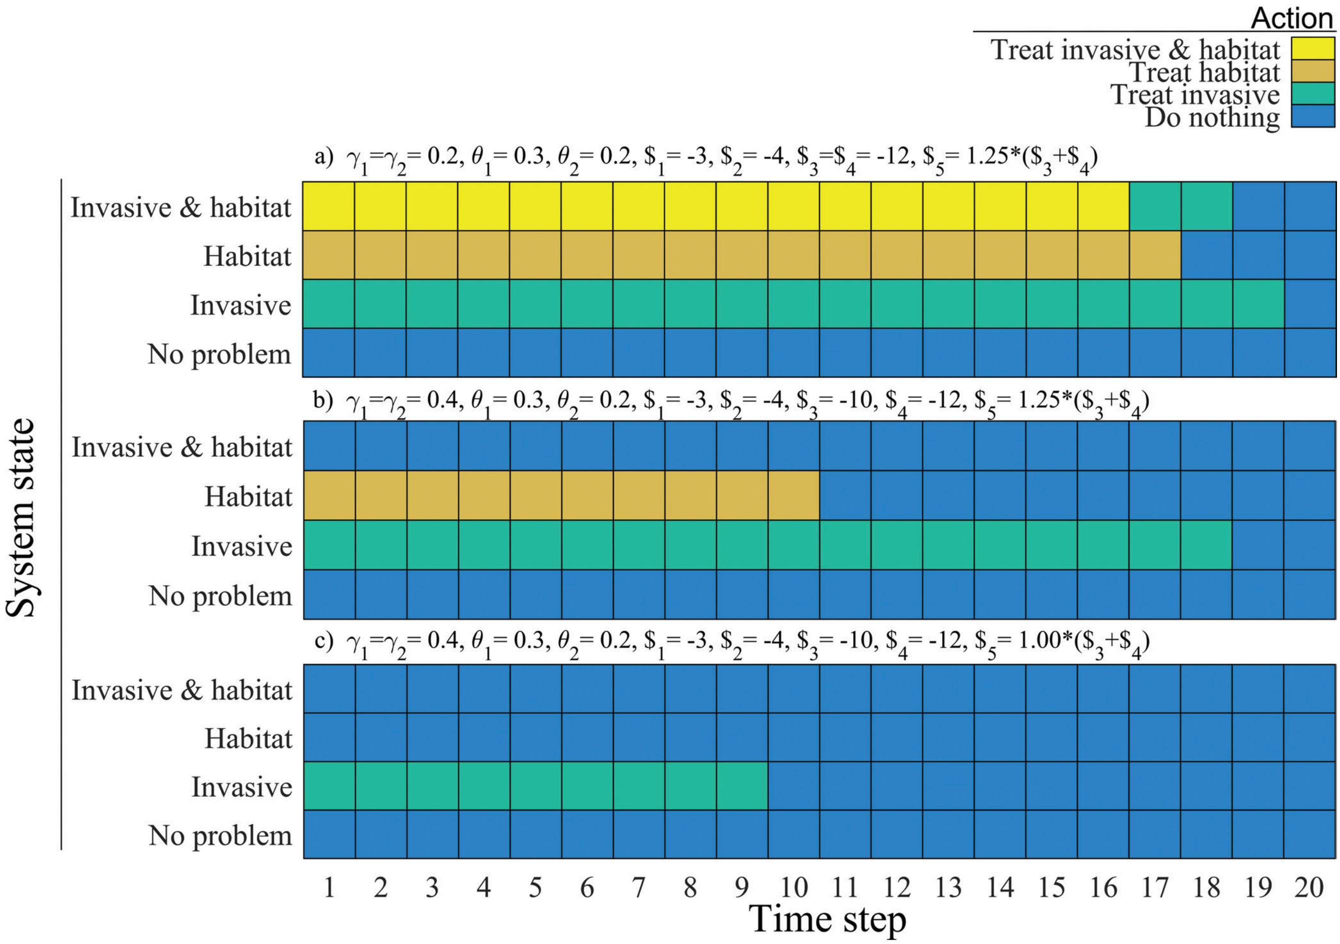

Figure 3. The results of an application of stochastic dynamic programming to a Markov decision problem. The rows within each panel give the system state, the columns give the time step, and the colored cells indicate the action that should be taken to minimize costs over the entire 20-step period. Panels (a), (b), and (c) present different scenarios defined by the given parameter values. The probabilities of developing an invasive species or habitat degradation during one step are γ_1 and γ_2, respectively. The probabilities of remedying invasive and habitat problems if they are treated are θ_1 and θ_2, respectively. The cost of treating invasive and habitat problems are \$$_1$ and \$$_2$, respectively. The cost of having invasive and habitat problems are \$$_3$ and \$$_4$, respectively, and the cost of having both problems at once is \$$_5$.

environmental decision-making processes by measuring social values, individual perceptions, or computing a common currency for otherwise nonfungible items (Gregory et al. 2012). Tools for doing this include choice experiments (Hanley, Wright and Adamowicz 1998) and value elicitation (Keeney 1992, Wright and Goodwin 1999). Martin and colleagues (2012) and Hanea and colleagues (2017) describe expert elicitation processes that can be used to parameterize decision-making models when data do not yet exist.

The costs used in the foregoing examples can be difficult to measure. The cost of a management mistake, *m*, used in example 1 will include not just actual costs or missed-opportunity costs, but also intangible costs such as public dissatisfaction. Similarly, the numerical entries in the payout matrices used in example 2 could involve intangible perceptions of the decision-maker. In example 3, the costs of allowing the system to have an invasive species or habitat degradation problem need to be specified. Social scientists capable of capturing these values could play an important role in defining future environmental science through the cost-optimization framework discussed in the present article.

Conclusions

The cost of performing applied environmental research can exceed the benefit to decision makers. The three examples used in the present article highlight multiple ways this might happen. The optimal expenditure on information, including the point at which no expenditure on information is warranted, is a function of the probability of making a management mistake without further information, the extent to which increasing investment in information translates into improved ability to anticipate the state of nature, and the likelihood that undesirable disturbances recur in the future. Viewing applied environmental problems through the lens of cost optimization entreats practitioners to focus on how information and decisions are linked. Despite broad recognition for a need to translate environmental science into practice there is not broad recognition of the existing decision-making tools for doing it (Enquist et al. 2017).

In the environmental sciences, it may be helpful to augment the commonly understood continuum between basic and applied research. Applied environmental research that unveils facts about nature without an explicit connection to decision-makers might be better described as *potentially applied research*. *Applied research* could be reserved for cases in which research activities were identified in conjunction with decision-makers and in which models are in place that map different data sets into different decisions. Although results from potentially applied research may eventually influence someone's thinking, the same can also be said of basic science.

A core tenet of sustainability is to meet the needs of the present without compromising the ability of future generations to meet their own needs. However, humanity may not be on a sustainable trajectory (Schramski et al. 2015) and profligate spending entails missed opportunity costs. Therefore, spending public money on applied science that is never applied is singularly ironic. There should be no question about the importance of training students and making scientific discoveries, but more can be done to ensure that all the goals of applied research are achieved. Cost optimization provides a framework for planning future applied environmental science that is efficient and effective. Although efforts to guarantee that financial resources are spent optimally may involve their own costs, environmental scientists and managers should nonetheless be prepared to justify expense allocation using optimality concepts. It is my hope that we continue developing a cost-optimization framework to plan, evaluate, and leverage applied environmental science so that future generations can be more sustainable.

Acknowledgments
Michael Runge provided uniquely constructive commentary that helped me refine and focus this article. Bob Steidl, Tom Stahl, and two anonymous reviewers also contributed helpful feedback.

Supplemental material
Supplemental data are available at *BIOSCI* online.

References cited
Addison PFE, Rumpff L, Bau SS, Carey JM, Chee YE, Jarrad FC, McBride MF, Burgman MA. 2013. Practical solutions for making models indispensable in conservation decision-making. Diversity and Distributions 19: 490–502.

Anonymous. 2007. The great divide. Nature 450: 135–136.

Arlettaz RA, Schaub M, Fournier J, Reichlin TS, Sierro A, Watson JEM, Braunisch V. 2010. From publications to public actions: When conservation biologists bridge the gap between research and implementation. BioScience 60: 835–842.

Ben-Haim Y. 2006. Info-Gap Decision Theory: Decision Under Sever Uncertainty. Elsevier.

Bode M, Wilson KA, Brooks TM, Turner WR, Mittermeier RA, McBride MF, Underwood EC, Possingham HP. 2008. Cost-effective global conservation spending is robust to taxonomic group. Proceedings of the National Academy of Sciences 105: 6498–6501.

Burgman M. 2005. Risks and Decisions for Conservation and Environmental Management. Cambridge University Press.

Canessa S, Guillera-Arroita G, Lajoz-Monfort JJ, Southwell DM, Armstrong DP, Chadès I, Lacey RC, Converse SJ. 2015. When do we need more data? A primer on calculating the value of information for applied ecologists. Methods in Ecology and Evolution 6: 1219–1228.

Chadès I, Chapron G, Cros M-J, Farcia F, Sabbadin R. 2014. MDPtoolbox: A multi-platform toolbox to solve stochastic dynamic programming problems. Ecography 37:916–920.

Chadès I, McDonald-Madden E, McCarthy MA, Wintle B, Linkie M, Possingham HP. 2008. When to stop managing or surveying cryptic threatened species. Proceedings of the National Academy of Sciences 105: 13936–13940.

Chadès I, Nicol S, Rout TM, Péron M, Dujardin Y, Pichancourt J-B, Hastings A, Hauser CE. 2017. Optimization methods to solve adaptive management problems. Theoretical Ecology 10: 1–20.

Clark CC, Mangel M. 2000. Dynamic State Variable Models in Ecology: Methods and Applications. Oxford University Press.

Conroy MJ, Peterson JT. 2013. Decision Making in Natural Resource Management: A Structured, Adaptive Approach. Wiley-Blackwell.

Denes-Raj M, Epstein S. 1994. Conflict between intuitive and rational processing: When people behave against their better judgment. Journal of Personality and Social Psychology 66: 819–829.

Ehrlich PR, Kareiva PM, Daily GC. 2012. Securing natural capital and expanding equity to rescale civilization. Nature 486: 68–73.

Enquist CAF et al. 2017. Foundations of translational ecology. Frontiers in Ecology and the Environment 15: 541–550.

Epanchin-Niell, RS, Haight RG, Berec L, Kean JM, Liebhold AM. 2012. Optimal surveillance and eradication of invasive species in heterogeneous landscapes. Ecology Letters 15: 803–812.

Fackler PL. 2011. MDPSolve: A MATLAB toolbox for solving Markov decision problems with dynamic programming. https://sites.google.com/site/mdpsolve/user-s-guide

Fackler PL, Haight RG. 2014. Monitoring as a partially observable decision problem. Resource and Energy Economics 37: 226–241.

Falcy MR. 2016. Conservation decision making: Integrating the precautionary principle with uncertainty. Frontiers in Ecology and the Environment 14: 499–504.

Falcy MR, McCormick JL, Miller SA. 2016. Proxies in practice: Calibration and validation of multiple indices of animal abundance. Journal of Fish and Wildlife Management 7: 117–128.

Field SA, O'Connor PJ, Tyre AJ, Possingham HP. 2007. Making monitoring meaningful. Austral Ecology 32: 485–491.

Gregory R, Failing L, Harstone M, Long G, McDaniels T, Ohlson D. 2012. Structured Decision Making: A Practical Guide to Environmental Management Choices. Wiley-Blackwell.

Hanea AM, McBride MF, Burgman MA, Wintle BC, Fidler F, Flander L, Twardy CR, Manning B, Mascaro S. 2017. Investigate Discuss Estimate Aggregate for structured expert judgement. International Journal of Forecasting 33: 267–279.

Hanley N, Wright RE, Adamowicz V. 1998. Using choice experiments to value the environment. Environmental and Resource Economics 11: 413–428.

Hayes KR, Barry SC, Hosack GR, Peters GW. 2013. Sever uncertainty and info-gap decision theory. Methods in Ecology and Evolution 4: 601–611.

Hilbert M. 2012. Toward a synthesis of cognitive biases: How noisy information processing can bias human decision making. Psychological Bulletin 138: 211–237.

Hope B. 2006. An examination of ecological risk assessment and management practices. Environment International 32: 983–995.

Iftekhar MS, Pannell DJ. 2015. "Biases" in adaptive natural resource management. Conservation Letters 8: 388–396.

Johnson DH. 1999. The insignificance of statistical significance testing. The Journal of Wildlife Management 63: 763–772.

Keeney RL. 1996. Value Focused Thinking: A Path to Creative Decisionmaking. Harvard University Press.

Lindenmayer DB, Likens GE. 2010. The science and application of ecological monitoring. Biological Conservation 143: 1317–1328.

Mapstone B. 1995. Scalable decision rules for environmental impact studies: Effect size, type 1 and type 2 errors. Ecological Applications 5: 401–410.

Marescot L, Capron G, Chadès I, Fackler PL, Duchamp C, Marboutin E, Gimenez O. 2013. Complex decisions made simple: A primer on stochastic dynamic programming. Methods in Ecology and Evolution 4: 872–884.

Martin TG, Burgman MA, Fidler F, Kuhnert PM, Low-Choy S, McBride M, Mengersen K. 2012. Eliciting expert knowledge in conservation science. Conservation Biology 26: 29–38.

Maxwell SL, Rhodes JR, Runge MC, Possingham HP, Ng CF, McDonald-Madden E. 2015. How much is new information worth? Evaluating the financial benefits of resolving management uncertainty. Journal of Applied Ecology. 52: 12–20.

McDonald-Madden E, Baxter PWJ, Fuller RA, Martin TG, Game ET, Montambault J, Possingham HP. 2010. Monitoring does not always count. Trends in Ecology and Evolution 25: 547–550.

Naidoo R, Balmford A, Ferraro PJ, Polasky S, Ricketts TH, Rouget M. 2006. Integrating economic costs into conservation planning. Trends in Ecology and Evolution 21: 681–687.

Nichols JD, Williams BK. 2006. Monitoring for conservation. Trends in Ecology and Evolution 21: 688–673.

Pichancourt J-B, Chadès I, Firn J, van Klinken RD, Martin TG. 2012. Simple rules to contain an invasive species with a complex life cycle and high dispersal capacity. Journal of Applied Ecology 49: 52–62.

Platt JR. 1964. Strong inference. Science 146: 347–353.

Raiffa H, Schlaifer R. 1961. Applied statistical decision theory. Graduate School of Business Administration. Harvard University.

Regan HM, Ben-Haim Y, Langford B, Wilson WG, Lundberg P, Andelman SJ, Burgman MA. 2005. Robust decision-making under severe uncertainty for conservation management. Ecological Applications 15: 1471–1477.

Runge MC, Converse SJ, Lyons JE. 2011. Which uncertainty? Using expert elicitation and expected value of information to design an adaptive program. Biological Conservation 144: 1214–1223.

Schramski JA, Gattie DK, Brown JH. 2015. Human domination of the biosphere: Rapid discharge of the earth-space battery foretells the future of humankind. Proceedings of the National Academy of Sciences 112: 9511–9517.

Tversky A, Kahneman D. 1974. Judgment under uncertainty: Heuristics and biases. Science 185:1124–1131.

Westgate MJ, Likens GE, Lindenmayer DB. 2013. Adaptive management of biological systems: A review. Biological Conservation 158: 128–139.

Williams BK, Johnson FA. 2015. Value of information in natural resource management: Technical developments and application to ping-footed geese. Ecology and Evolution 5: 466–474.

Wintle BA, Runge MC, Bekessy SA. 2010. Allocating monitoring effort the face of unknown unknowns. Ecology Letters 13: 1325–1337.

Wright G, Goodwin P. 1999. Rethinking value elicitation for personal consequential decisions. Journal of Multi-Criteria Decision Analysis 8: 3–10.

Matt Falcy (matt.falcy@oregonstate.edu) is the fish conservation biologist for the Oregon Department of Fish and Wildlife.