Panaceas and diversification of environmental policy

William A. Brock* and Stephen R. Carpenter†‡

*Department of Economics, and †Center for Limnology, University of Wisconsin, Madison, WI 53706

Edited by Elinor Ostrom, Indiana University, Bloomington, IN, and approved June 29, 2007 (received for review March 7, 2007)

We consider panacea formation in the framework of adaptive learning and decision for social–ecological systems (SESs). Institutions for managing such systems must address multiple timescales of ecological change, as well as features of the social community in which the ecosystem policy problem is embedded. Response of the SES to each candidate institution must be modeled and treated as a stochastic process with unknown parameters to be estimated. A fundamental challenge is to design institutions that are not vulnerable to capture by subsets of the community that self-organize to direct the institution against the overall social interest. In a world of episodic structural change, such as SESs, adaptive learning can lock in to a single institution, model, or parameter estimate. Policy diversification, leading to escape from panacea traps, can come from monitoring indicators of episodic change on slow timescales, minimax regret decision making, active experimentation to accelerate model identification, mechanisms for broadening the set of models or institutions under consideration, and processes for discovery of new institutions and technologies for ecosystem management. It is difficult to take all of these factors into account, but the discipline that comes with the attempt to model the coupled social–ecological dynamics forces policy makers to confront all conceivable responses. This process helps induce the modesty needed to avoid panacea traps while supporting systematic effort to improve resource management in the public interest.

adaptive learning | governance | institutional design | minimax regret | model uncertainty

Ostrom (1) describes a tendency for analysts to prescribe singular solutions for environmental problems. Examples include centralized government regulation, grassroots governance, civic environmentalism, privatization, and so forth. There is a tendency in the literature to take very strong positions concerning specific policy instruments, which are advocated as a panacea while alternatives are derided as worse than useless. Yet panaceas create significant problems (1). Despite the tendency toward panaceas, many different policy instruments have been used to address society's environmental problems (2). Often it is not clear which policy is optimal, and in many cases a mix of policies may perform better than a single policy (2).

The articles in this special feature consider diverse explanations for the origins, consequences, and solutions of panaceas. We discuss the hypothesis that panaceas can derive from a failure to properly address model uncertainty in sensible pragmatic practice. Although model uncertainty is one of many potential causes of panaceas, the management of social–ecological systems (SESs) must depend to some extent on models, whether formal or informal, for the processes to be managed. Embedded in each model are institutions, the set of nonphysical constraints on economic behavior (3) that are rules of the game in a society or, more formally, are the humanly devised constraints that shape human interaction (4). We represent the institutional design problem as a choice among alternative institutions in an uncertain world. A set of models, each representing institutions as well as other relevant aspects of the SES, is compared with respect to estimated welfare outcomes, weighted according to consistency with data. Panaceas arise when people think they have arrived at the truth regarding how the world would work under different institutional designs and the design they propose is optimal. We discuss how such panaceas arise and how they can be avoided.

We motivate the theoretical discussion with a description of a specific SES, the Northern Highland Lake District (NHLD) of Wisconsin in the United States; this SES exemplifies the multi-scaled, multivariate challenges of ecosystem management (5, 6). In this setting, panaceas can form in several ways: adherence of individuals or managers to a data-inconsistent model; Bayesian lock-in through the weight of history; failure to address gradual change in slow variables; lack of tools to anticipate or detect regime shifts; and failure to build resilience to cope with unforeseeable regime shifts. This example suggests a theoretical framework for adaptive management, which we use to demonstrate ways in which apparently rational policy processes can lead to panaceas. This mechanism of panacea formation derives from dynamic statistical properties of adaptive processes. We close with a discussion of ways to avoid, or escape from, some types of panaceas.

Northern Highland Lake District

The NHLD is a rural region of rolling moraines, second-growth forest, extensive wetlands, and numerous lakes. It encompasses 5,300 km^2, >7,600 lakes, and 65,000 permanent residents (in 2000) (5, 7, 8). Outdoor recreation and forest products are mainstays of the economy. The NHLD is famous for sport fishing, which is an important component of the outdoor recreation industry.

Scales of some major social and ecological processes of the NHLD range from a few days to ≅12,000 years (time since the last ice age) and from a single lake to the entire Great Lakes region of North America (Fig. 1). A full analysis of the SES would address multiple, interacting ecosystem services on a heterogeneous landscape [see supporting information (SI) Table 1]. To simplify the analysis, we narrowed our focus to the interaction between shoreland vegetation, fishing, and fish dynamics (6). Lakeshore residents have tended to remove living and fallen trees, which provide crucial habitat for fishes (9, 10). Once shoreline forests or fallen trees in lakes are removed, regrowth or replacement takes decades or even centuries. Habitat loss causes fishes to grow more slowly, while fish stocks become more vulnerable to collapse from overfishing, invasion of exotic species, or other causes (6, 7, 9, 10, 11). Spatially, lakes are linked by the mobility of anglers; if fishing is poor on a given lake, anglers can move (6). Spatial patterns of lakeshore residence and shoreland management can also change, but these changes occur on timescales of years, whereas angler movements can occur in a matter of hours.

We focus on three levels of decision making by resource users (6): selecting a region in which to spend time (NHLD vs. other regions), selecting a base lake on which to buy or rent housing (note that buy vs. rent implies longer vs. shorter residence time on site), and

Author contributions: W.A.B. and S.R.C. performed research and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Abbreviations: NHLD, Northern Highland Lake District; SES, social–ecological system; IID, independent and identically distributed.

‡To whom correspondence should be addressed. E-mail: srcarpen@wisc.edu.

This article contains supporting information online at www.pnas.org/cgi/content/full/0702096104/DC1.

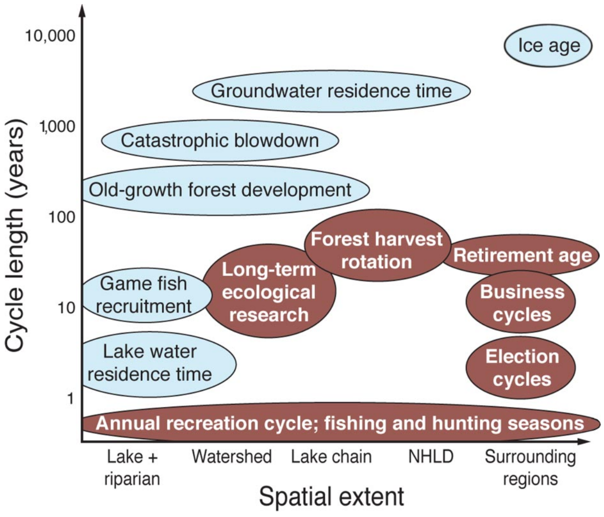

Fig. 1. Temporal scale (measured as return time) vs. spatial scale (measured as extent of spatial patterns) for some key ecological (blue ellipses) and social phenomena (brown ellipses) in the NHLD.

selecting a lake on which to fish given the choice of base lake (Fig. 2). Fishery dynamics on each lake depend on lakeshore management practices of the residents and fishing practices of the residents and visitors (Fig. 2). In this system, large rapid shifts in fish stocks of individual lakes can result from decadal changes in habitat interacting with annual changes in the housing market, daily decisions of anglers, and exogenous shocks such as storms (6). Under some circumstances, angler movement can trigger waves of fish stock reduction across the landscape. Management must therefore account for slow changes in habitat, somewhat faster changes in lakeshore housing and vegetation, and even faster movements of anglers among lakes. At any given time, optimal policies are heterogeneous among lakes on the landscape (6). Moreover, policies must continually adapt to the changing mosaic of users, habitat, and fish stocks cocreated by tens of thousands of people interacting with resources of thousands of lakes (6). Policies that are constant in space or static in time either underexploit resources or

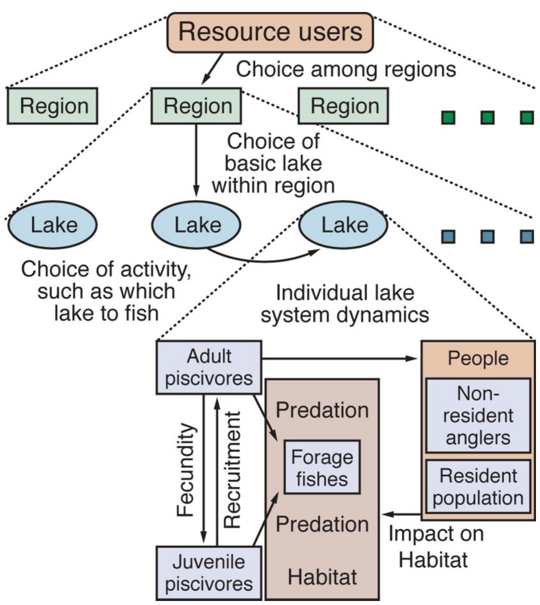

Fig. 2. Scales of decision making by recreational users, and dynamics of people and fish stocks in lakes.

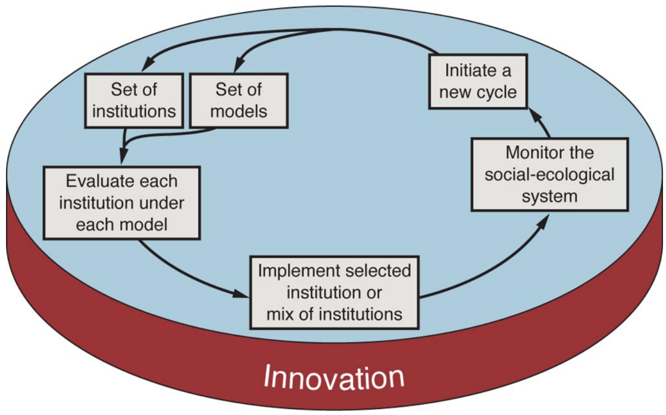

Fig. 3. Cycle of adaptive learning and decision. The set of institutions and set of models for the SES change from cycle to cycle. These dynamics depend, in part, on introduction of innovations by people or by emergence of new ecological or social phenomena.

trigger cascading waves of resource breakdown. To succeed, policy must evolve in time and diversify over space.

The NHLD illustrates three key features of change in SESs that must be addressed by any framework that seeks to foster an adaptive approach: (*i*) Change is ongoing at multiple scales, from the slow weathering of rock or percolation of groundwater to life cycles of microorganisms measured in minutes. This change has both repeated and novel components as captured by the title *Time's Arrow, Time's Cycle* (12). The future will bring patterns not seen in the past. Thus adaptive-learning processes must continually invent and evaluate new institutions and models. (*ii*) Social–ecological change, although usually gradual and predictable, is sometimes rapid, extensive, and unpredictable (13–16). These extensive changes, or regime shifts, may be difficult or impossible to reverse. Some regime shifts have significant consequences for ecosystem services, livelihoods, and human well-being (2). (*iii*) Any institution that is implemented evokes a reaction from people and ecosystems, which then cocreate a system that may become different from the one that the institution was intended to manage. The reaction may take days or decades, and the resulting cocreated system may be substantially new. For any proposed institution, an analysis of the reaction of the SES must be considered, and this dynamic must be estimated and predicted, before the ultimate outcome of the proposed institution can be evaluated. This analysis is the heart of adaptive management, and the source of potential panacea traps.

Adaptive Learning and Decision

We employ a general framework of statistical adaptive learning and decision for governing natural resources of a complex region such as the NHLD. We do not imply that such a framework is the only way for society to deal with complex resource management issues. However, the framework is a useful reference point for several reasons. The analysis is repeatable, in the sense that two decision makers working from the same information would reach the same conclusion. Many real-world resource management agencies claim to be managing adaptively. A significant body of rigorous research exists for statistical adaptive learning (17–20). Finally, if any management approach were panacea-proof, one would think that a data-driven adaptive-learning approach would be it. Yet, we will show how easily panaceas can emerge from adaptive learning if parameters are changing on a slow timescale relative to the timescale of the learning process.

The fundamental unit of adaptive learning and decision is a cycle (Fig. 3). In each cycle, the analyst considers a set of institutions for managing the region and a set of models for dynamics of the SES, together with available relevant data. The models address ecological and social dynamics, including social responses to the institutions. Forecasts of social welfare are developed for each candidate institution and model, and they are weighted according to their data-based credibility. This comparison of institutions leads to selection of an institution, or mix of institutions, to be implemented. Following implementation, the SES is monitored. Note that decisions regarding whether and what to monitor, and responses to changes revealed by monitoring, are embedded in the set of institutions to be considered. Deliberate experimentation to determine the effects of particular policies on the SES is also among the institutions that could be considered (19).

The decision to initiate a new cycle is a critical step; failure to take this step can create a panacea. The new cycle could be prompted by monitoring data (e.g., information that suggests a regime shift is occurring). It could be triggered by discovery of new institutions, technologies, or ecosystem phenomena. Note that policies to invest in exploration leading to discoveries are part of the set of institutions under consideration. Thus, the outcome of adaptive learning and decision rests on a foundation of innovation (Fig. 3); change depends on assimilation of discoveries into the set of institutions and models under consideration. When a new cycle is initiated, the sets of institutions and models under consideration include new inventions, and the available data include new information.

Evaluation of Alternative Policies

This section takes a small step toward formalizing the discussion to show how panaceas can arise from failure to address model uncertainty in adaptive-learning processes. First we analyze a simple optimization problem where different models are represented by different values of a parameter and the decision maker has beliefs regarding the parameter. Second we show how the decision maker can learn the true value of the parameter by using data, but the decision maker can persist in the wrong belief for a long time if the true value of the parameter changes. We discuss Bayesian learning on a fast timescale, and we assume that estimation converges to a reasonable estimate of social–ecological dynamics. We then introduce a second, slower timescale on which the SES shifts to a new regime. On this slower timescale, Bayesian learning on the fast timescale can easily stabilize on the wrong model, a form of panacea.

Assume that ecosystem i produces service potential $x_{t+1,i}$, which is a scalar that follows (on the fast timescale) the one-dimensional dynamic system

$$x_{t+1,i} = f(x_{t,i}, r_{t+1,i}, u_{t+1,i}),\ i = 1, 2, \ldots N. \tag{1a}$$

Here, $x_{t,i}$ denotes the service potential or state (e.g., capacity to produce ecosystem services, or biomass, of a renewable resource) of ecosystem i at time t; $f(x, r)$ is the law of motion, where r denotes a scalar stochastic shock process; $u(.)$ is a control function to be specified below that denotes some impact of humans (e.g., harvest of a renewable resource, damage to habitat, pollution, or management intervention); and N is the number of ecosystems on the landscape.

We will specialize below to case **1b**,

$$x_{t+1,i} = f(x_{t,i}, r_{t+1,i}) - h_{t+1,i},\ i = 1, 2, \ldots N, \tag{1b}$$

where the intertemporal control choice, $\{u\} = \{h\}$, is some impact, such as harvesting. We assume the stochastic process $\{r_{t+1,i}\}$ is independently and identically distributed (IID) with finite mean and finite variance across all points in time and across all ecosystems $i = 1, 2, \ldots N$. Assume the managers of ecosystem i wish to maximize the mathematical expectation given their beliefs of long-run, steady-state utility

$$E[U(h)] \tag{2}$$

produced by the ecosystem by choosing $h_{t+1,i} = u[f(x_{t,i}, r_{t+1,i})]$ where $u(.)$ is a function with sufficient regularity that the result (**1b**) is a first-order nonlinear autoregression that possesses a unique invariant measure and converges exponentially fast to this unique invariant measure (ref. 21, chap. 3.6). Here, "long run" is a period long enough on the fast timescale so that **1b** has converged close enough to the unique invariant measure, or steady-state distribution, to use it as an approximation. E denotes the expectation taken with respect to the assumed unique stationary distribution of the state, x, determined by control function $u(.)$.

The fast-scale learning problem is that the policy makers do not know the parameters of the dynamics and must learn their true values by experience. In addition, the planner faces the problem that the true values of the parameters change on the slow timescale, which poses a challenge for any scheme for learning unknown parameters, including Bayesian learning.

Focus now on a single ecosystem. Let the set U of control functions $u(.)$ be linear-in-output rules of the form $h_{t+1} = u[f(x_t, r_{t+1})] = (1 - s) f(x_t, r_{t+1})$, where s is in the closed interval [0, 1]. Because from **2** $x_{t+1} + h_{t+1} = f(x_t, r_{t+1})$ we have

$$x_{t+1} = sf(x_t, r_{t+1}) \equiv sy_{t+1},\ h_{t+1} = (1 - s)y_{t+1}. \tag{3}$$

Think of s as the fraction of ecosystem output y_{t+1} that is saved while $1 - s$ is the fraction harvested. Let the objective be to choose s in [0, 1] to maximize

$$E\{\ln[(1 - s)y_{t+1}]\}, \tag{4}$$

where the mathematical expectation is computed over the long-run steady-state distribution determined by fixed s, and $\ln(x)$ denotes the natural logarithm of x.

Here we present an example that can be solved in closed form (ref. 21, chaps. 3 and 6, and refs. 22 and 23). Although this example uses linear savings rules as in **3**, it is not necessary to restrict the optimization to linear savings rules. On the fast timescale, the law of motion of the ecosystem is

$$f(x_t, r_{t+1}) = r_{t+1}x_t^{a_{t+1}}, \tag{5}$$

where $\{\ln[r_t]\}$ is IID with mean 0 and finite variance and $\{a_t\}$ is IID binary with $0 < a_1 < 1/2 < a_2 < 1$ the two possible values, with probabilities p_1 and p_2, respectively. We also assume that $\{\ln[r_t]\}$ and $\{a_t\}$ are independent of each other. In the context of this model, panacea advocates believe $\{a_t\}$ is a deterministic process with $a_t = a_1$ or a_2 when in truth it is a stochastic process.

In **5**, a measures the elasticity of response of ecosystem service potential tomorrow to the service potential of the ecosystem today. It is defined by

$$d\ \ln(f)/d\ \ln(x) = a. \tag{6}$$

If we take natural logs of **3** and calculate the steady-state limit distribution, it is easy to show that the optimal rule s^* that maximizes **4** is

$$s^* = E(a) = p_1a_1 + p_2a_2. \tag{7}$$

We define "panacea belief of type $j = 1, 2$" to be a belief that $a = a_j$ with probability 1. A panacea j believer will optimize **3** and choose $s^*_j = a_j$, but the true dynamics (if no harvest) is given by

$$x_{t+1} = r_{t+1}x_t^{a_{t+1}}, \tag{8}$$

where $\{r\}$ and $\{a\}$ are IID stochastic processes independent of each other, as defined above. Thus, panacea believers are similar to fundamentalist ideologues, people who believe they know the truth with an unwarranted degree of certainty.

If a_1 is close to 0 and a_2 is close to 1, both panacea believers do poorly in the long run compared with one who knows the true process [under the assumption that $E(a) = 1/2$]. Notice that the panacea 1 believer exploits the ecosystem too heavily, whereas the panacea 2 believer does just the opposite. Note also that each panacea believer should, after several periods of observing actual output, learn that a is actually random and, by using sensible statistical procedures, eventually learn the true process $\{a_t\}$ as well as learn the process $\{r_t\}$. Notice that it is difficult to learn that $\{a_t\}$ is random, because $\{r_t\}$ is also random.

Panaceas 1 and 2 are analogs of certain strategies in classical decision theory. Panacea 1 is closely related to a maximax strategy, and panacea 2 is closely related to a maximin strategy (ref. 24, p. 282). Of course, we are not saying that maximax and maximin are always panaceas; indeed, in some settings these strategies are optimal (24). To compute these strategies, prepare a table of payoffs to each possible act for each possible state of the world (ref. 24, pp. 277–280). Each row of this table gives the list of payoffs across states of the world for a particular action. A maximax (maximin) strategy is an action that maximizes the best-case state-of-the-world payoff (worst-case payoff) over the set of possible actions. Find a maximax (maximin) strategy by finding the value of the largest (smallest) payoff across each action's row and selecting the action that gives the largest value.

How might a planner who is not a panacea believer attempt to learn the truth in the setting above? There are two main sets of approaches taken in modern decision theory. The first approach is Bayesian. In a Bayesian approach, the decision maker assigns prior probabilities to the different states of the world and chooses the action that maximizes the expected payoff.

The second set of approaches is non-Bayesian and includes maximax, maximin, and minimax regret. Roughly speaking, minimax regret minimizes the maximum loss when one is ignorant of the true situation (ref. 24, p. 280).

Beyond the scope of formal decision theory, panacea thinking might be avoided by expanding the set of institutions and models under consideration, or initiating new cycles of adaptive learning (Fig. 3). Here we focus on the outcome of formal decision methods in the context of the simple panacea model above. Specifically, when the true value of parameter a changes infrequently (but changes are unknown to the decision maker), Bayesian approaches are susceptible to panaceas 1 and 2 defined above. We argue that panaceas of this type can potentially be mitigated by using non-Bayesian approaches.

Bayesian Analysis on the Fast Timescale. We assume that process $\{a_t\}$ is deterministic and either a_1 or a_2 is the true state of $\{a_t\}$ but the decision maker does not know which it is. Let $p_{1,t}$ be the belief probability of the decision maker that the true state is a_1 at date t. To ease the calculations below, we assume that $\{\ln(r_t)\}$ is IID normal with mean zero and finite, known variance v. Bayes's updating formula gives us

$$p_{1,t+1} = p_{1,t} f_1(x_{t+1}|x_t, u_t)/D_{t+1} \qquad \textbf{[9a]}$$

$$p_{2,t+1} = p_{2,t} f_2(x_{t+1}|x_t, u_t)/D_{t+1} \qquad \textbf{[9b]}$$

$$D_{t+1} = p_{1,t} f_1(x_{t+1}|x_t, u_t) + p_{2,t} f_2(x_{t+1}|x_t, u_t), \qquad \textbf{[10]}$$

where u_t is the control rule applied at date t, which, in our case, is parameterized by s. Here the f_j values denote the likelihoods of the data under belief system j.

We may write, after shortening the notation in an obvious manner,

$$\ln\left(\frac{p_{1,t+1}}{p_{2,t+1}}\right) = \ln\left(\frac{p_{1,t}}{p_{2,t}}\right) + \ln\left(\frac{f_{1,t+1}}{f_{2,t+1}}\right) \qquad \textbf{[11]}$$

and observe that if we sum **11** from $t = 1, 2, \ldots T$ observations and divide by T, and take the probability limit, if it exists, we can simply evaluate $E\{\ln(f_1/f_2)\}$ to tell whether the ratio of p_1 to p_2 goes to 0 or to infinity. E denotes the mathematical expectation under the true density function. It is unknown, but it can be estimated from observable data by taking the sample average. Although an intermediate value is possible, the generic case is 0 or 1 for the limiting ratio of p_1/p_2 as we shall sketch below.

Proposition: *If the probability limit of*

$$\frac{1}{T}\sum_{t=1}^{T} \ln\left(\frac{f_{1,t}}{f_{2,t}}\right) \Rightarrow E\left\{\ln\left(\frac{f_{1,t}}{f_{2,t}}\right)\right\} \equiv D_{1,2}, \quad T \Rightarrow \infty \qquad \textbf{[12a]}$$

exists and is positive (negative) then $p_{1,t}/p_{2,t}$ converges to infinity (0). Furthermore, this proposition generalizes to settings where there are a finite number of possible values of a, i.e., $a_1, a_2, \ldots a_n$, and we can get an expression for the speed of convergence from the limit in **12**.

For our example with $\{r_t\}$ IID lognormal, if we put $r_{t+1} \equiv \exp(v^{0.5} e_{t+1})$ and take natural logs of both sides of the equation $x_{t+1} = s\, r_{t+1} x^a$, we obtain $y_{t+1} = \ln(s) + v^{0.5} e_{t+1} + a\, y_t$, where $\{e_t\}$ is IID normal with mean zero and variance v. Now suppose $a = a_1$ is true. Then

$$\ln\left(\frac{f_{1,t}}{f_{2,t}}\right) = \frac{1}{2v}\{[y_{t+1} - \ln(s) - a_2 y_t]^2 - [y_{t+1} - \ln(s) - a_1 y_t]^2\}. \qquad \textbf{[12b]}$$

By the assumption that $a = a_1$ is the true value of a, we have

$$y_{t+1} = \ln(s) + a_1 y_t + v^{0.5} e_{t+1}. \qquad \textbf{[12c]}$$

Sufficient conditions for a law of large numbers (LLN) to hold for the left side of **12a** are modest. Under the assumption that an LLN holds, compute the right side of **12a** to obtain

$$D_{1,2} = E\{[(a_1 - a_2)y_t)]^2\}, \qquad \textbf{[12d]}$$

where the expectation is computed over the long-run steady-state distribution determined by **12c**. $D_{12} > 0$ if a_1 is not equal to a_2. Notice that the control s drops out in this derivation. A similar calculation may be done for all harvesting functions of output that give $x_{t+1} = g(r_{t+1} x^a)$, where $g(.)$ is monotone increasing and has an inverse. Calculate by putting $y = \ln(x)$ to get $y_{t+1} = \ln\{g[\exp(a\, y_t + v^{0.5} e_{t+1})]\}$, then form the right side of **12a** and use the inverse of g to find an expression analogous to **12d**.

Sufficient conditions for **12** to converge can be found in Vuong (25) and Sin and White (26), who study the problem of selecting between two or more models. They show, for example, that when the choice is between two models, the one closest to the true data-generating process is chosen in the sense that it attracts all of the limiting probability in **12**. Although they do not study what happens when control functions are inserted into the dynamics, the worked example above suggests the conjecture that sufficient conditions may be found on the set of controls U such that their results can be applied. This analysis is beyond the scope of this short article.

A general optimization and learning framework can be applied to problems of managing ecosystems where there are unknown parameters to be estimated as well as possible alternative stable states for some values of those parameters (27). Discounting has an important impact on learning in this setting (19, 27) but is also beyond the scope of this article.

The Slow Timescale. The preceding section shows how a Bayesian learning mechanism on a fast timescale might be too sluggish in adapting to parameter changes on a slow timescale, if the time between changes on the slow timescale is long enough that convergence of Bayes's learning is almost complete on the fast timescale. It is beyond the scope of this short article to formalize a serious model of what happens on the slow timescale and how the panacea can be mitigated. Other literature points the way. Expression **12a** is closely related to recursive algorithms that can make the learning more responsive to structural changes (28). Beck and Wieland (20) study incentives to experiment to learn slowly changing parameters. Colacito *et al.* (17) study how long it takes for an optimal Bayesian learning mechanism to correct a specification error in the case of two models that are almost equally well fitted to economic data. They also compare passive and active learning in a dynamic setting analogous to passive and active adaptive management of ecosystems (19).

Summary of the Example. Our example is a kind of panacea. It implies that $p_{1,T}/p_{2,T}$ acts like $\exp(T\,D_{1,2})$ for large T, so if $D_{1,2}$ is positive (negative) then $p_{1,T}$ is converging exponentially fast to 1 (0). Thus if a_1 is true for the time interval $[1, N]$ where N is large, the decision maker will come to believe that the parameter a is at a_1 with high probability. If the value of a then switches to a_2 on $[N + 1, 2N]$, it will take some time for the updating mechanism to leave a small neighborhood of placing high weight on a_1. This inertia of Bayesian updating mechanisms can lead to a false sense of certainty that a is at a_1 when a is really at a_2.

This panacea resembles the extreme simplification described for failing corporations (29) or government policies (30) and the rigidity trap described for ecosystem management (31). Peterson *et al.* (32) present a simulation model of lake management that illustrates this phenomenon.

Introducing New Options

Adaptive learning easily converges to beliefs that support panaceas. How can panacea formation be countered? In the NHLD, spatially heterogeneous fishery regulations and zoning are emerging to counter tendencies toward panaceas (7). In the context of adaptive learning, several kinds of processes can diversify policy. Some are Bayesian. For example, the planner can monitor regime-shift indicators and use these data to determine when a new cycle of adaptive learning is needed. Elsewhere, we discuss construction of regime detectors in certain cases in which the regime shift is caused by a bifurcation in the underlying dynamics (33, 34). De Lima (35) provides a general test for regime shift. Other approaches are non-Bayesian. This section briefly describes two non-Bayesian approaches, minimax regret decision making and regulatory tiering, and points to a general need for research on innovation and diversification of environmental policy.

Minimax Regret. Some approaches for decision making under uncertainty, e.g., maximin and minimax regret, have vastly different implications for the speed of learning an unknown structure (36). In the context of a bank making loans to lendees with unknown payoff probabilities, a maximin lender will assume the worst-case scenario and will refuse to lend to a large group of people no matter how much information there is on their payoff probabilities, because the belief probability on the possibility that a member of the group will default is always positive (36). Hence, the worst-case scenario is default and the maximin lender will never know the true payoff probability of loans to members of this particular group. However, the mixing induced by minimax regret ensures that, typically, the lender will lend to a positive fraction of this group. Hence, in just one period, the true probability will be revealed by the LLN, and the lender will have learned the true probability, whereas the maximin lender will never learn the true probability.

This parable suggests that if we have N ecosystem dynamics being managed and they are all alike, we would get a similar result when we compare maximin learning vs. minimax regret learning. Even on the fast timescale, the maximin learner will always assume the worst-case scenario and hence always plan to guard against the worst case, because even under Bayesian updating, the date t posterior probability of the worst-case parameter value is still positive (even though it becomes small exponentially fast if the worst-case parameter value is not the true parameter value). Minimax regret mixes options in many situations. This diversification allows one to use a panel of case studies to learn more effectively how to improve performance in any one of the particular situations. Active adaptive management, which involves deliberate experiments to accelerate model identification (19), goes even further than minimax regret in exploring alternative institutions or policies. Deliberate whole-ecosystem experiments have helped frame policy options for lake management in the NHLD (7, 8, 10).

Regulatory Tiering. Regulatory tiering (37) is the practice of making the burden of regulations lighter on smaller entities. A kind of panacea in regulation of negative externalities (such as pollution) is the argument that all businesses must be treated equally. But often equal treatment is inefficient. Brock and Evans (37) provide an example of regulating polluters. The size distribution of business firms is approximately lognormal. Suppose businesses in a particular industry (e.g., farms in agriculture) pollute the environment (e.g., by eutrophication of surface waters). If the classical remedy of taxation of negative externalities is implemented, there are administrative costs at both the firm level and the regulatory authority level. There are also uncertainties on how bureaucrats will behave. These costs act as fixed costs plus variable costs at the firm level. Variable costs appear because, for example, auditing is more intensive as tax rates are increased because avoidance attempts tend to increase with higher taxes. Small firms cannot average down these fixed costs. It is more efficient to introduce regulatory tiering in which small firms might be entirely exempted, middle-sized firms be taxed more lightly than larger firms, and the largest firms be taxed at full social cost per unit of pollutant emission (37).

Responses of the Community to Management Institutions. Certain issues must be addressed whenever an institution is proposed to solve a resource management problem. If one proposes, for example, a government authority, one must take into account the response of the community. Although we are dealing here with resource management, we can learn much from management of international trade (38). In this case, the very existence of a trade management authority causes lobbies to form for and against trade restrictions. Organization and maintenance of an effective lobby is analytically similar to organization and maintenance of community-based management of a natural resource. Therefore, the same structural conditions that facilitate or hamper community-based management (39) also facilitate or hamper formation of effective lobbying efforts. In the case of community-based management, it is in society's interest that effective social action self-organize. In the case of international trade (assuming no negative externalities), it is just the opposite. Hence, in some cases (e.g., when there are no negative externalities) it is in the social interest not to interfere with trade and, hence, to shut down government intervention. The unpredictability of subgroup responses to the advent of any institution is another good reason not to be dogmatic concerning what will work well and what will not (35). Indeed, the surprising results in game theory, where cooperation emerges in settings where one would expect noncooperation, underscores the difficulty of predicting collective responses (40). Colander (41) addresses rent-seeking activities of subgroups in general economic settings and what can be done to tame this damaging activity.

For any proposed institution, an analysis of the reaction of the community must be conducted, and this dynamic must be estimated and predicted before the ultimate outcome of the proposed institution can be evaluated. For example, the policy maker might want to implement taxation of negative externalities. At the same time, it might be helpful to have a bit of government regulation in setting standards, to catalyze people to work out a community-based management scheme, and to work to align property rights with the public interest. Such a combined approach, rather than a panacea approach of focusing strictly on taxation of negative externalities, could deliver better results in some settings, even in situations where most experts would agree that taxation of negative externalities is called for.

Discussion

We studied a very simple, stylized problem under an assumption that one could separate the hierarchy of temporal speeds into two timescales. In regions like the NHLD, learning and adaptation in pursuit of goals takes place in a hierarchy of timescales and spatial extents (Figs. 1 and 2 and SI Table 1). Furthermore, the systems being managed are coupled with feedbacks causing coevolution across space and time scales, and across ecological and social domains. Thus, the actual problem is significantly more complex than we have studied here. Nevertheless, our simplified analysis reveals important points.

Even relatively sophisticated adaptive-learning approaches fall easily into panaceas in the face of potential regime shifts, especially during periods when the system appears stationary in a statistical sense. This mechanism provides a rational explanation for panaceas (32). We suspect that a wide range of other mechanisms, such as social or political processes that create stability despite changing environments, are also important for formation of panaceas in practice (14, 42).

Although we have emphasized how fast and slow timescales can lead to panaceas, failure to address spatial heterogeneity may also lead to panaceas. In the NHLD, uniform regional fishery regulations increase risk of fishery collapse across the landscape (6). This panacea can arise if average fishery parameters across the diverse landscape are used to set policies for all of the lakes. Fishing regulations are sometimes uniform in space, and it is plausible that the uniformity is explained by inappropriate spatial averaging. However, it is equally plausible that the panacea derives from political pressure from anglers to simplify and standardize regulations, from resort owners who believe they will lose business if fishery regulations are tightened on "their" lakes, or a host of other social and political phenomena. Such explanations imply social or political processes that drive policy choice away from the social optimum. We think that our model is only one explanation for panaceas and that more work is needed to develop a comprehensive theory.

A new cycle of adaptive learning is an opportunity for consideration of new institutions, new models for dynamics of the SES, and new data relevant to performance of the system (Fig. 3). Innovation depends in part on past decisions on monitoring the SES and investment in discovery of new options for institutions or models of social–ecological dynamics. The brevity of this article does not allow for adequate discussion of innovation and its role in overcoming the tendency to form panaceas, which is a crucial topic that needs greater attention.

Gunderson and Holling (31) describe an adaptive cycle that oscillates between phases of discovery and innovation, then phases of consolidation and efficiency. Adaptive learning, in contrast, focuses on a systematic cycle of adaptive learning based on an underlying matrix that supports innovation (Fig. 3). Thus, we do not imply any particular temporal relationship between phases of innovation, analysis, and decision, but we recognize that choice of institutions affects future discoveries and innovations, and the possibilities for considering them in future decisions. Linked adaptive cycles across a range of timescales (panarchies) play the same role for Gunderson and Holling (31).

Avoidance of panaceas seems to be characteristic of resilient SESs and a key to maintaining ecosystem services in a shrinking world (43). Therefore, more work is needed on practices for avoiding panaceas. All available approaches rest on the active discovery and evaluation of new options for governing SESs. Institutions that evoke such discoveries may be the ones best able to avoid panaceas.

We thank participants in the Beyond Panaceas Project and three referees for helpful comments, Joanna Broderick for editing, Bill Feeny for graphics, and the National Science Foundation for support.

1. Ostrom E (2007) *Proc Natl Acad Sci USA* 104:15181–15187.
2. Millennium Ecosystem Assessment (2005) *Ecosystems and Human Well-Being* (Island, Washington, DC).
3. Greif A (1994) *J Polit Econ* 102:912–950.
4. North D (1991) *J Econ Perspect* 5:97–112.
5. Peterson GD, Beard TD, Jr, Beisner BE, Bennett EM, Carpenter SR, Cumming GS, Dent CL, Havlicek TD (2003) *Ecol Soc* 7(3):1. Available at www.ecologyandsociety.org/vol7/iss3/art1/print.pdf.
6. Carpenter SR, Brock WA (2004) *Ecol Soc* 9(1):8. Available at www.ecologyandsociety.org/vol9/iss1/art8/print.pdf.
7. Carpenter SR, Benson BJ, Biggs R, Chipman JW, Foley JA, Golding SA, Hammer RB, Hanson PC, Johnson PTJ, Kamarainen AM, *et al.* (2007) *BioScience* 57:323–335.
8. Magnuson JJ, Kratz TK, Benson BJ, eds (2006) *Long-Term Dynamics of Lakes on the Landscape* (Oxford Univ Press, Oxford).
9. Christensen DL, Herwig BR, Schindler DE, Carpenter SR (1996) *Ecol Appl* 6:1143–1149.
10. Sass GD, Kitchell JF, Carpenter SR, Hrabik TR, Marburg AE, Turner MG (2006) *Fisheries* 31:321–330.
11. Schindler DE, Geib SI, Williams MR (2000) *Ecosystems* 3:229–237.
12. Gould SJ (1987) *Time's Arrow, Time's Cycle* (Harvard Univ Press, Cambridge, MA).
13. Holling CS (1973) *Annu Rev Ecol Syst* 4:1–23.
14. Scheffer M, Carpenter SR, Foley JA, Folke C, Walker B (2001) *Nature* 413:591–596.
15. Repetto R, ed (2006) *Punctuated Equilibrium and the Dynamics of US Environmental Policy* (Yale Univ Press, New Haven, CT).
16. Brock W, Durlauf S, West K (2007) *J Econometrics* 127:629–664.
17. Colacito R, Cogley T, Sargent T (2007) *J Money Credit Bank* 39:67–99.
18. Easley D, Kiefer N (1988) *Econometrica* 56:1045–1064.
19. Walters CJ (1986) *Adaptive Management of Renewable Resources* (MacMillan, New York).
20. Beck G, Wieland V (2002) *J Econ Dyn Control* 26:1359–1377.
21. Bhattacharya R, Majumdar M (2007) *Random Dynamical Systems* (Cambridge Univ Press, New York).
22. Mirman L, Zilcha I (1975) *J Econ Theory* 11:329–339.
23. Bertocchi G, Spagat M (1998) *J Econ Dyn Control* 23:209–231.
24. Luce R, Raiffa H (1957) *Games and Decisions* (Wiley, New York).
25. Vuong QH (1989) *Econometrica* 57:307–333.
26. Sin CY, White H (1996) *J Econometrics* 71:207–225.
27. Dechert W, O'Donnell S, Brock WA (2007) *J Diff Eq Appl* 13:121–133.
28. Sargent T (1999) *The Conquest of American Inflation* (Princeton Univ Press, Princeton).
29. Miller D (1993) *Acad Manage Rev* 18:116–138.
30. Scott JC (1998) *Seeing Like a State* (Yale Univ Press, New Haven, CT).
31. Gunderson LH, Holling CS (2002) *Panarchy* (Island, Washington, DC).
32. Peterson GD, Carpenter SR, Brock WA (2003) *Ecology* 84:1403–1411.
33. Carpenter SR, Brock WA (2006) *Ecol Lett* 9:311–318.
34. Brock WA, Carpenter SR (2006) *Ecol Soc* 11(2):10. Available at www.ecologyandsociety.org/vol11/iss2/art10/.
35. De Lima P (1998) *J Bus Econ Stat* 16:227–236.
36. Brock WA (2006) *Econ J* 116:F427–F440.
37. Brock W, Evans D (1985) *RAND J Econ* Autumn:398–409.
38. Magee S, Brock WA, Young L (1989) *Black Hole Tariffs and Endogenous Policy Theory* (Cambridge Univ Press, Cambridge, UK)
39. National Research Council (2002) *The Drama of the Commons* (Committee on the Human Dimensions of Global Change, Natl Acad Press, Washington, DC)
40. Imhof LA, Fudenberg D, Nowak MA (2005) *Proc Natl Acad Sci USA* 102:10797–10800.
41. Colander D (1984) *Neoclassical Political Economy* (Ballinger, Cambridge, MA).
42. Scheffer M, Westley F, Brock WA (2000) *Ecosystems* 3:451–471.
43. Carpenter SR, Folke C (2006) *Trends Ecol Evol* 21:309–315.