

Environmental foresight and structural change

M.B. Beck

*Warnell School of Forest Resources, University of Georgia, Athens, GA 30602-2152, USA
(mbbeck@uga.edu)*

Abstract: Policy-makers and the public, it has famously been said [Brooks, 1986], are more interested in the possibility of non-linear dislocations and surprises in the behavior of the environment than in smooth extrapolations of current trends. How indeed should we design our models to generate environmental foresight, to detect, in particular, threats to our environment lying “just beyond the horizon”? In facing this prospect of potentially profound dislocations in behavior, the problem is that the number of state variables in the model, whether they interact, how they interact, and the form of their interactions, may be evolving over time. What may have appeared to have been an insignificant mode of behavior in the past — buried within the uncertainty of the model and the historical data — may come to dominate behavior in the future. Technically, we may call this a change of structure. The concern of the paper is to address the challenge of constructing and employing models to generate environmental foresight in the presence of structural change. A number of case histories, ranging across lake eutrophication, urban ozone levels, the restoration of ecosystems, the circulation of waters in the North Atlantic, and the invasion of exotic species, are used to construct a much more immediate sense of the nature of structural change and, therefore, the character of the challenge of generating environmental foresight. Some mathematical and logical formalities are then introduced, both to define the issues more sharply and to open up the means with which to address them. This provides an opportunity to take stock of three rather different programs of model-building used, over the decades, to generate environmental foresight. We close by illustrating a set of possible responses to the essential challenge through a number of contemporary case studies: in assessing, *inter alia*, the reachability of the lay community’s hopes and fears for the future of their cherished piece of the environment; in apprehending and diagnosing the possibility of imminent structural change; and in examining the record of the past for emergence of the seeds of any such structural change.

Keywords: Adaptive control and management; analysis of uncertainty; cultural theory; reachable futures; recursive estimation; stakeholder futures; surprise; sensitivity analysis; watershed management

1. INTRODUCTION

Thinking about the future, would one really imagine the science base as remaining invariant over the decades now commonly within the reach of our forecasting horizons? The answer, of course, is “No, one would not”. Our knowledge is continually evolving. Yet we must necessarily invest the structures of the models we use with invariance. For how, otherwise, could we design a model in which the “birth”, “death”, or “extinction” of a state variable, for example, is to be mimicked [Allen, 1990; Kauffman, 1995]? How could one discover the rules by which systems rearrange themselves (the way in which their state variables interact with each other) and then use these to make projections into the future? For the time being, we are going to have to cope with the

continually evolving knowledge base employing models remaining squarely in the domain of the state-space representations of classical mechanics. They will contain an *invariant* number of states (\mathbf{x}) and parameters ($\boldsymbol{\alpha}$) with formal interactions among them that are just as invariant. In this domain, then, how should we accommodate the uncertain, changing nature of the knowledge base, assuming we believe that the environmental problems we now face are sensitive to such uncertainty and change?

This is a rather general, rough hewn block of a question, however; and we shall not attempt to answer it — at least, not in its entirety. In order to make any progress, fragments will have to be chipped from the block. We first present some case histories in dealing with the concepts of

trajectories, equilibria, excursions, dislocations and surprises in the behavior of the environment. Our purpose is to provide motivation and to develop a figurative sense of structural change. Some formalities of a mathematical and logical nature are then introduced. They are necessary for specifying more tightly those sub-problems and sub-questions, to which we may have some chance of responding through formal computational analysis. We shall close by illustrating our responses to date with a number of case studies.

2. CASE HISTORIES: WHAT EXACTLY IS THE PROBLEM?

Let us suppose the global system is evolving along a trajectory. In the past few decades we have become aware of the way in which the activities of humans have been sufficient — over the last couple of centuries or so — to cause the system to deviate from a trajectory along which it might otherwise have traveled, had it not been for, say, the industrial revolution.¹ This, then, is our predicament. As we stand, as we always do, on the threshold of the future, what is of concern to us about the environment of our global system? Because what is of concern to us — how we conceive of the future in relation to the trajectory discerned from the past — shapes the character of the problems to be explored computationally for the purposes of generating environmental foresight. It will shape too the direction of our search into the record of the past.

2.1 Equilibria and Deviations

That our perception is of a trajectory, as opposed to a temporary deviation away from an equilibrium, is in itself a conceptual advance [Holling, 1996]. For the moment, however, let us think of the trajectory, since it is evolving relatively slowly (over centuries), as a sequence of quasi-equilibria spanning relatively short periods of time (decades).

¹ The words we use here are fraught with the difficulties of “loaded” interpretations. For example, would the global system have experienced a natural trajectory of evolution in the absence of man, since man is part of nature? Without having to put quotation marks around almost every word, the following is written without presuming the evolution of the global system to be necessarily natural in the absence of man, or unnatural because of his presence and involvement.

In the classical domains of water and air pollution, the concern has been to reverse the imposition of a stress manifesting itself in a fully developed, palpable, local (or regional) environmental strain. In the case of Lake Erie in North America, for example, cultural eutrophication brought about by the excessive accumulation of nutrients could be reversed, in principle, by rectifying the anthropogenic diversion of these materials from their natural cycles into the aquatic environment [Schertzer and Lam, 2002]. From a knowledge of what Erie had once been like (in living memory), together with a knowledge of the behavior of nearby, similar lakes — not similarly stressed, however — the levers of policy to be pulled and the desired position to which the system ought then to return were predictable, by and large. Likewise, given a knowledge of how ozone is generated in the atmosphere from emissions of gaseous oxides of nitrogen (NO_x) and volatile organic carbon (VOC) compounds, the tangible environmental strain of injurious episodes of high concentrations of ozone in the air of urban areas could be predictably reversed [Dennis, 2002].

The record of environmental protection, as it happens, is *not* littered with the failures of any “classical” strategies of prediction and policy implementation. But there are some case histories bearing important lessons to be learned, such as those, for instance, of Schertzer and Lam [2002], Hornberger [2002] and Dennis [2002]. In particular, formal, computationally encoded models have now been used for a sufficiently long period for this horizon to include significant evolution in the science base — or, at least, evolution in those parts of it deemed sufficiently problem-significant to have been included in the model. Indeed, let us label such problem-relevant knowledge as the {presumed known}. By this very choice of words, it is obvious this {presumed known} will have a complement, i.e., the something unknown, or the something judged not to be problem significant. We shall call this the {acknowledged unknown}. Clearly, what may have been consigned to the {acknowledged unknown} at the time — and therefore omitted from the model used for prediction and the formulation of policy (deliberately, unwittingly, or in ignorance) — may subsequently come into play in thwarting the complete success of the actions taken to restore the system to the presumed trajectory from which it had earlier departed.

Thus, in the eutrophication of Lake Erie the {presumed known} of the 1960s and 1970s amounted essentially to the biochemistry of nutrient assimilation in an ecosystem, not the {acknowledged unknown} of the hydrodynamics

governing nutrient movement. In the 1980s and 1990s, a mixture of the flaws in the {presumed known}, which seemed less secure and overly crude in retrospect, and the growing perception of the significance of the {acknowledged unknown}, was brought to bear on explaining why the earlier regulatory actions had failed to relieve the strains of extended bouts of anoxia in the lake. In the case of urban ozone episodes, the {acknowledge unknown} of natural, biogenic emissions of VOCs in the rural surrounds of the metropolitan centres came to undermine the success of a regulatory policy locked firmly onto curtailing emissions of unnatural, anthropogenic VOCs from urban traffic — as though wilfully, almost perversely, the *only* allowable {presumed known} [Dennis, 2002].

Both of these case histories deal with deviations from a trajectory unfolding (and being refolded) over a matter of a few decades — a relatively short time-scale, within living memory; deviations furthermore, of broadly predictable properties, all the foregoing quibbling notwithstanding. Both too are instances where, for the purposes of implementing urgent regulatory action at a given point in time, the formulation of policy had inevitably to be crystallized around a particular, crisp image (computational encoding) of the science base, i.e., the {presumed known}. Such crystallization is not without its difficulties, however. The science base, or at least the formal predictions derived from it, can be in a state of considerable flux, even over the span of just a few years, as evidenced by Schneider and Thompson [1985] in their work on projections of changes in stratospheric ozone concentrations [see also Beck, 2002]. It would have been easy to have locked the formation of policy on to the wrong image of the moment. Yet policy and control there must be, based on what is deemed the best current image of the system's behavior. It would be naïve of us to suggest this could have been otherwise and unsurprising for the subsequent auditing of the actions taken to reveal the rise to significance of other elements in the evolving knowledge base. Yet the ethos of model-building in the 1970s and 1980s was to proceed as though the {presumed known} could be encapsulated in a largely unchanging model of what would in due course be revealed as a largely unchanging science base.

2.2 Dislocations and Structural Change

It occurs to us, however, that stresses can induce something of a more plastic, enduring character, beyond merely a reversible, elastic strain. The system, if pushed too much, may not return to its original “equilibrium”, or, better put, may not proceed to the position on the trajectory it would

have reached, had there never been any anthropogenically induced excursion. It may instead migrate towards patterns of movement about another equilibrium [Holling, 1978, 1996], that is, embark on an altogether different trajectory. So in looking to the future, we can become concerned about this: that the behavior of our environment in the future may come to be radically different from what we have known in the past, albeit the relatively recent past. And we are really rather creative in imagining what these threats to our environment might be [Leggett, 1996]. Concerns of this kind, we call herein the modern problems, without suggesting that the classical problems have somehow been conquered or rendered irrelevant. They are modern, for the purposes of our discussion here, in the sense that they have attributes requiring us to conceive of schemes of computational exploration beyond the classical forms of problem-solving in the cases of eutrophication and ozone just described. Two examples, both from oceanography and the marine sciences, illustrate the point — in their own slightly different ways.

First, there is the aspiration to restore the ecological systems of the oceans and seas to something of a quasi-pristine condition. The difficult question here, of course, is that of what constitutes this condition, a question subject to vigorous debate [Jackson, 2001]. To see why, let us first bring to mind the image of a model as comprising nodes (the state variables, x) connected by branches (causal influences underlain by parameters, α) — in literal terms here, a web of predator-prey interactions among the various biological species. To those familiar with molecular graphics, the web can be thought of as the three-dimensional image of a complex compound. With this in mind, today's dominant interactions amongst the state variables, with one subset of emboldened branches, say at time t^0 , shows a significantly different (emboldened) structure from that of times past, in living memory (t), let us say.² Indeed, if the distinction between emboldened (dominant) and ordinary (minor) branches in the web is sharply contrasted — so that only the emboldened branches remain apparent — we could visualize in the limit a *change of structure* in the system, between t and t^0 . This *per se* is significant. Jackson's point, however, is that the historical record (t^{\dots}), reaching further back than mere living memory, would show yet a quite different structure. In other words, our perception of these ecological systems, in the Caribbean and North Atlantic, for example,

² As in Figures 1 and (especially) 2, to be discussed in more detail later.

is that they have undergone a kind of smooth evolution over the centuries — arguably, an excursion from their original trajectory — which can appear to us as a structural change. The quasi-equilibrium of t is (arguably) *not* the point to which the system should be restored; it *might* (better) be that of t^{---} .

Strictly speaking, if no species has become extinct in these marine ecosystems, the number of nodes in the web has remained unchanged, although the nature of the interactions between the species may have changed, and certainly the set of dominant interactions has changed. We shall refer to this physical realization of change as an *apparent* structural change [Beck, 2002], to distinguish it thus from the true structural change of evolution envisaged in Allen [1990]. From the perspective of using models to generate environmental foresight — and this is the key — were we able to capture a description of the entire foodweb, one and the same structure of a model would suffice for describing its behavior from t^{---} (historical past) through t (living memory) to t^0 (the present) and on into the near-term future (t^+) and well beyond (t^{+++}). In this, in an all-encompassing, *singular* {presumed known}, movement along the trajectory of the system would be revealed as a continual structural undulation, of the rising to significance, and falling into insignificance, of the various modes of behavior attaching to the interactions (α) amongst all the known state variables (x) — apparent structural change of a kind. If such a singularity cannot be captured, we would have the apparent structural change of the sequence of structurally different images of merely the emboldened interactions at each stage in these blocks of time, from the past into the future.

Second, there are the ramifications reverberating from the conceptual insight of the ocean conveyor belt, as applied to circulation of the waters of the North Atlantic [Weaver, 1995]. The {presumed known} of geophysical theory and physical oceanography in this instance is extensive; so extensive, in fact, that it was not until the early 1990s that enough of it could be assembled into a model for practical solution on a computational platform, in order thus eventually to liberate the insight of the conveyor belt itself. An ominous “point of no return” in our potential futures has thereby been revealed [Weaver, 1995]: a location on a trajectory (again, perhaps a disturbed trajectory), which, should the system be pushed beyond it (by the accumulation of yet further deviant disturbances), would cause the conveyor belt eventually to be switched off, for the first time in living memory, with consequences of potentially profound significance for the climate of

north-west Europe. If the insight — the highly aggregated, high-level conceptual description of past behavior, extracted from the massive bulk of the {presumed known} — is not substantially flawed, it suggests the possibility of radically different, immensely “hard-to-predict” behavior arising in the future. The potential environmental strain is merely incipient; the strain may well be plastic, far from elastic; and the system, essentially of global proportions, has no nearby replicates to serve as reliable benchmarks against which to gauge the progress of restoration.

We have a sense here, then, of the kind of nonlinear dislocation of such great concern to us [Brooks, 1986]. It is change of a different order indeed to the smoother extrapolations of the eutrophication and ozone problems, even of the marine ecological systems of the Caribbean Sea and North Atlantic Ocean. *This*, then, is the problem.

2.3 Surprise

We know that these case histories of eutrophication and ozone, as related here, map neatly onto Holling’s concept of engineering resilience: policy is formed to maintain the system at a single, fixed equilibrium [Holling, 1996]; and use of the word “engineering” evokes a sense of manufactured security and completeness of the knowledge base (encoded in a model). In turn, Thompson [1997] has mapped the corresponding myth of nature, of “Nature benign” in this instance [Holling, 1986], into one of Cultural Theory’s four forms of social solidarity (or perspectives on the man-environment relationship), namely the Individualists. Theirs is an essentially peril-free outlook in which the environment is thought capable of absorbing all insult and injury before returning assuredly to a pre-disturbance equilibrium position. Restoration of the marine ecosystems of the Caribbean and North Atlantic *might* map across Holling’s myth of “Nature perverse but tolerant” into Thompson’s solidarity of the Hierarchists: restoration to the pre-disturbance equilibrium will be the norm, for as long as the environment is not placed under excessive stress, in which case excursions into the largely unfamiliar and undesirable will ensue. But the “point of no return” for the ocean conveyor belt seems like a phrase conjured up by the Egalitarians of Cultural Theory, just as much as it captures the essence of why herein we seek other means of computational exploration of the future. “Nature ephemeral”, the myth to which the Egalitarians are supposed to subscribe, says that any disturbance, no matter how small, may plunge the system into the wholly undesirable and

unfamiliar. Ecological resilience, for which Holling would have us design [Holling, 1996], would facilitate navigation around and through the worst of this “undesirable and unfamiliar”, with a continually receding horizon of foresight.

In classical terms, reducing the uncertainties — in preparing ourselves to cope with the impacts of climate change and other structural shifts and dislocations — has been equated with making a model so comprehensive ($\{\text{presumed known}\} \rightarrow \infty$) as to render the probability of significant structural change arising out of our ignorance vanishingly small ($\{\text{acknowledge unknown}\} \rightarrow 0$). Yet for all the computational power at our disposal, we are still surprised by the behavior of our environment. That is to say [Price and Thompson, 1997]:

A myth of nature provides its holder with a way of seeing the world [$\{\text{presumed known}\}$] and with a way of not seeing it [$\{\text{acknowledged unknown}\}$]. This means that if the world happens not to be the way the myth-holder is convinced it is, he or she will notice this discrepancy straight away. Enlightenment, therefore, is always time-lagged and, since it results in the enlightened one being tipped out of one quadrant of the Cultural Theory scheme and into one of the other three, it comes as something of a shock: a surprise.³

Surprise, in other words, is always relative, which explains why, whenever something unexpected befalls us, there is always someone who ‘saw it coming’!

The theory of surprise [Thompson and Tayler, 1986; Thompson et al., 1990] is built on this relativistic, but far from unconstrained, foundation:

an event is never surprising in itself;
it is potentially surprising only in relation to a particular set of convictions about how the world is;
it is actually surprising only if it is noticed by the holder of that particular set of convictions.

Someone must have seen the invasion of zebra mussels coming to Lake Erie and the other Great Lakes; certainly, these exotic creatures have acted as a policy-confounding feature [Schertzer and Lam, 2002]. Their abundance, and the consequent highly effective removal of phosphorus from these

waters, has wrought a distortion in the foodweb now deemed excessive; so excessive, in fact, as to give credence to the idea of augmenting the inflow of nutrients to the lake. That would overturn three decades of policy; and this would be surprising, to many of us. It is as though a new atom (an exotic species) were inserted into a recognizably stable and coherent molecule (ecosystem). Thereafter, the molecular graphic (the visualization of our computational model) would be animated, as the collection of nodes bounce around within the resulting looseness of their elastic attachments in response to the invader, just as they have done in several aquatic ecosystems in the northern United States [for example, Strayer et al, 1999; Matthews et al, 2002] — a physical manifestation of structural change. The system never quite returns to the location it might otherwise have attained on any pre-ordained trajectory.

3. MATHEMATICAL AND LOGICAL FORMALITIES

Our essential question is this: how should we go about generating environmental foresight while acknowledging the evolving nature of the knowledge base of primary science, manifest as apparent structural change in the computational encoding of a model? We must now introduce some formalities of models, thus to break this general question into a number of more specific, albeit no more markedly tractable, questions.

3.1 Constructing the Space Between the Model and the (Unknowable) Truth

Many models of the behavior of environmental systems can be defined according to the following (lumped-parameter) representation of the state variable dynamics of classical mechanics,

$$dx(t)/dt = f\{x, u, \alpha; t\} + \zeta(t) \quad (1a)$$

with observed outputs being defined as follows,

$$y(t) = h\{x, \alpha; t\} + \eta(t) \quad (1b)$$

in which f and h are vectors of nonlinear functions, u , x , and y are the input, state, and output vectors, respectively, α is a vector of model parameters, ζ and η are notional representations respectively of those attributes of behavior and output observation that are not to be included in the model in specific form, and t is continuous time. Should it be necessary, spatial variability of the system’s state can be assumed to be accounted for by, for example, the use of several state variables of the same attribute of interest at the several defined locations.

³ “Quadrant” implies four solidarities, the fourth being the Fatalists, who believe the myth of “Nature capricious”, i.e., that nature will do what it chooses irrespective of man’s behavior.

For any system, the choices of $[u, y]$ determine the (observable) *external description* of its behavior, to be labeled as B . Those aspects of the science base mobilized into the computational encoding of the model — the hypothetical mechanisms considered significant to the manner in which input, causative disturbances (u) are transcribed into output effects (y) — are signaled by the choices of $[f, h; x, a]$. In short, the *structure* of model is most succinctly conveyed in terms of $[f, h]$, which denote the logical inter-connections among u , x , and y , while a signifies parameterization of the particular mathematical expressions of all the hypothetical mechanisms believed to underpin these interactions. We may call $[x, a]$ the *internal description* of the system's behavior, as the complement of $[u, y]$.

In “truth”, then, the structure of the system's behavior can be supposed to be of (almost) infinitely high order, but invariant.⁴ Let us denote this as $[f^\infty, h^\infty]$. We, with our models in the realm of the finite, $[f^0, h^0]$ say, work on a much more macroscopic plane. Our models have a crude resolving power, even for those of a very high order ($+n$), with structure $[f^{+n}, h^{+n}]$. As the behavior of the system moves along a trajectory through time, certain of its myriad features may dominate affairs for a while, while others may be dominant over other periods. It may appear to us as though the structure of the system's behavior is varying with time, i.e., that for a while we would characterize this as $[f^0, h^0; t^-]$ and over another period as $[f^0, h^0; t^+]$, or even as $[f^i, h^i]$ during t^- but $[f^j, h^j]$ during t^+ (as seemingly manifest in the behavior of an ecological system before and after invasion by exotic species). All the fine-grained flutter and drift are lost imperceptibly — for a while — in the gap between what is included in the model, $[f^0, h^0]$, and the truth of the matter, $[f^\infty, h^\infty]$. With progress over the decades in the knowledge bases of the primary sciences, the tendency is in general for $[f^0, h^0]$ to grow incrementally towards $[f^{+n}, h^{+n}]$, *arguably* with an accompanying sense of closing in on the truth.⁵

⁴ In the absence of evolution that is, in particular, absent the birth of a truly novel species, which would constitute a state variable emerging from “somewhere outside” the vector x and requiring logically thereafter to be inserted into it.

⁵ There is a difference, of course, between what we can encode (hypothetically) on the computer and the actual nature of things; and not in all subjects is there inexorable growth in the order of the models being employed in the everyday practice of science (the case of understanding surface water acidification is one such example; Beck [2002]).

What exactly, however, should we suppose is the content of the *gap* between $[f^0, h^0]$ and $[f^\infty, h^\infty]$, the structural error in the model, that is? Put simply, this inadequacy, or approximation, may enter into equation 1 through a , ξ , and η , although these points of entry differ in their interpretation and significance. The principal distinction is between a , embedded within the choices for $[x, a, f, h]$, which signify that which we presume (or wish) to know of the system's behavior, relative to the purpose of the model, and $[\xi, \eta]$, which acknowledge in some form that which falls outside the scope of what we believe we know. The difference between the two is as the difference between what we have called above, respectively, the {presumed known} and the {acknowledged unknown}. Much, of course, must be subsumed under the latter, that is, under the definitions of ξ and η . We may have chosen to exclude from the model some of that which was known beforehand, but which was judged not to be significant; there may be features for which there are no clear hypotheses (and therefore no clear mathematical expressions), other than that these may in part be stochastic processes with presumably quantifiable statistical characteristics; there may be yet other features of conceivable relevance, but of which we are simply ignorant; and, as is most familiar, there may be factors affecting the processes of observation such that we are unable to have uncorrupted, perfect access to knowledge of the values of the inputs, states, or outputs.

Essentially, the model is all that we have to work with to cope with the gap between $[f^0, h^0]$ and $[f^\infty, h^\infty]$, where this gap will constitute the whole of the {acknowledged unknown} and of the {presumed known} being wrongly presumed known. It is all we have to apprehend something of significance, to our understanding and actions, within the gap. In particular, in this process of apprehension, the model — being the vessel containing all the relevant hypothetical knowledge from the science base — is to be pitted against all the relevant experience of observed past and, most importantly (as we shall see), *imagined* future behavior, collectively symbolized as $B(t^{---})$, $B(t^-)$, $B(t^+)$, $B(t^{+++})$, for example. In short, our original question can be restated in a more formal, but interim, manner as:

Given a model of fixed structure $[f^0, h^0]$, and given descriptions of behavior over time, $B(t^{---})$, $B(t^-)$, $B(t^+)$, $B(t^{+++})$, how should we use the model to cope with, even apprehend, the very high likelihood of changes in time of significance to our understanding and actions in

the space between $[\mathbf{f}^0, \mathbf{h}^0]$ and the truth of the matter, namely $[\mathbf{f}^\infty, \mathbf{h}^\infty]$?

3.2 Probing the Constructed Space

For several years, handling an apparent change of structure — from $[\mathbf{f}^0, \mathbf{h}^0; \hat{t}^-]$ to $[\mathbf{f}^0, \mathbf{h}^0; \hat{t}^+]$ — was accommodated through an immutable structure $[\mathbf{f}^0, \mathbf{h}^0]$ populated with parameters that were permitted to change with time, i.e., through the device of $\alpha(\hat{t}^-)$ migrating to $\alpha(\hat{t}^+)$ [see, for example, Beck, 1987]. But this is not entirely satisfactory as a description of the nature of the gap, as we now appreciate. It addresses merely the issue of the {presumed known} being in error, using the elements of the parameter vector α as tags. In other words the {presumed known} can be parameterized as {presumed known (α)} in order to identify more specifically the inadequacies of the model's *constituent* hypotheses, as revealed through those reconstructed estimates of the parameters found to change (significantly) with time. This does not cater for the need to detect changes of *significance* – not of pure chance – in the {acknowledged unknown}. Yet ξ and η are fictions, mere labels for the divergence between the presumed known content of the model and the unknowable truth. They are not computable from the givens of our question above. We need therefore to find a tractable means of probing the nature of the {acknowledged unknown}, no matter how approximate this might be.

To this end, as a first approximation, consider how we actually put a model to work. We compose the model ($[\mathbf{f}^0, \mathbf{h}^0]$), attempt to reconcile it with past observed behavior ($B(\hat{t}^-)$), use it to predict future behavior, and, when that behavior becomes manifest (as $B(\hat{t}^+)$), reconsider the model's formulation and proceed to the next round of prediction. Whether known as adaptive environmental assessment and management [Holling, 1978], or more narrowly as adaptive control in engineering systems [Åström and Wittenmark, 1989], these principles constitute what we may call a recursive, predictive working environment: with thoughts trained on a future with an ever receding forecasting horizon as we pass through time. We can place our model of equation 1 formally into this framework, putting it in fact into the naturally recursive format of what is called an innovations representation of the system's behavior, as follows:

$$d\mathbf{x}(t|t_{k-1})/dt = \mathbf{f}\{\mathbf{x}(t|t_{k-1}), \mathbf{u}(t), \alpha\} + \mathbf{K}\boldsymbol{\varepsilon}(t|t_{k-1}) \quad (2a)$$

$$\mathbf{y}(t_k) = \mathbf{h}\{\mathbf{x}(t_k|t_{k-1}), \alpha\} + \boldsymbol{\varepsilon}(t_k|t_{k-1}) \quad (2b)$$

Here attention has been restricted merely to the conventional intervals of time passing from one observing instant t_{k-1} to the next, t_k . In spite of the formalities, what is of importance will prove to be of conceptual, as much as algorithmic, significance. Thus, formally, the argument $(t|t_{k-1})$ signals a predicted value of the associated quantity at some (future) time t utilizing the model and all observed information, in particular, in respect of the observed output \mathbf{y} , up to and including that available at the most recent sampling instant, t_{k-1} . $\boldsymbol{\varepsilon}(t_k|t_{k-1})$ is the innovation, i.e., the mismatch between the predicted and observed values of the output at the next sampling instant in discrete time, t_k , in equation 2(b); $\boldsymbol{\varepsilon}(t|t_{k-1})$ in equation 2(a) is the value of this quantity at times not coincident with the sampling instant. \mathbf{K} is a weighting matrix and can be thought of as a device for distributing the impacts of the innovations among the constituent representations of the various state variable dynamics, i.e., the representations $\mathbf{f}_i\{\cdot\}$ for each state \mathbf{x}_i . \mathbf{K} is central to the conceptual argument we now present.

First of all, comparing equations 1 and 2, it is evident that the cleavage in the one, between the {presumed known} and the {acknowledged unknown}, is as that between $[\mathbf{f}, \mathbf{h}]$ and $[\xi, \eta]$ (in equation 1), while in the other (equation 2) it is as that between $[\mathbf{f}, \mathbf{h}]$ and $[\mathbf{K}\boldsymbol{\varepsilon}, \boldsymbol{\varepsilon}]$. Second, unlike ξ and η , $\boldsymbol{\varepsilon}(t_k|t_{k-1})$ is a computable quantity, being the mismatch (in equation 2(b)) between the forecast value of the output and the observed output — albeit not the truth of the matter (hence some of the necessary approximation in our argument).⁶ In this way, $\boldsymbol{\varepsilon}$ is a kind of gauge of the foregoing gap between $[\mathbf{f}^0, \mathbf{h}^0]$ and $[\mathbf{f}^\infty, \mathbf{h}^\infty]$. Third, just as we are familiar with the notion of reconciling the model's behavior with that observed of the real thing, in order to adjust and estimate the values of the model's *conventional* parameters (α), so this same process of reconciliation can be employed to reconstruct values for the elements (κ , say) of the matrix \mathbf{K} . Indeed, the original motivation for the algorithmic form of equation 2 was precisely this: to reconstruct κ , instead of setting their values by prior assumption, in order to reconstruct estimates of α [Ljung, 1979]. If then the estimates of κ remain essentially the same as their prior, presumed values of 0.0, none of the empirical mismatches between the model and the data — in effect the innovations $\boldsymbol{\varepsilon}$ — are fed back into the

⁶ Further, assumptions will have to be made as to how the computed value of $\boldsymbol{\varepsilon}$ from equation 2(b) is technically to be treated within equation 2(a), but this is literally a technical matter of no great significance for the conceptual argument being made.

predictions made of future behavior. Our predictive instrument is operating essentially on the basis of the {presumed known} alone. Should the elements of κ come to be reconstructed in the course of events as substantially non-zero, our predictive instrument is beginning to rely on the {acknowledged unknown}, and perhaps exclusively so.

Armed with these three conceptual interpretations of the formalities of equation 2, we can proceed to our vital insight into the role of \mathbf{K} . Given the association of $[\mathbf{K}\boldsymbol{\varepsilon},\boldsymbol{\varepsilon}]$ with the {acknowledged unknown}, κ can be attached to the parameterization of this entity in the same manner as α has been the device for parameterizing the {presumed known}. We have thus the {presumed known (α)} and {acknowledged unknown (κ)}, where now, given the computability of $\boldsymbol{\varepsilon}$, we have an empirical means of both identifying the inadequacies of what has been included in the model and apprehending something of significance in what has been excluded from it. Further, as with all the individual elements of α , what transpires in reconstructing the individual elements of κ can provide pointers to the specific consequences of this “something of significance” — something of substance in guiding the search for the reasons underlying the apparent structural change. Alternatively, think of it as follows. Randomness in the gap between $[\mathbf{f}^0, \mathbf{h}^0]$ and $[\mathbf{f}^\infty, \mathbf{h}^\infty]$ should cause flutter in $\boldsymbol{\varepsilon}$, possibly even of high amplitude (for example, from the spurious corruption of observing errors). Persistent mismatches of significance in $\boldsymbol{\varepsilon}$ should eventually cause adaptation and change within α and κ , the one pointing to structural errors in the expression of the {presumed known}, the other to something of significance being apprehended in the {acknowledged unknown}.

3.3 Metaphors and the Essential Challenge

Figuratively (and approximately), the structure of the model has been parameterized by the branches of the network of Figure 1 — an essential visual metaphor [Beck, 2002]. The $[\alpha_{ij}]$ are included within the basic rectangular frame connecting the states (\mathbf{x}) with each other, while the $[\kappa_{ij}]$ attach to the frame but point outwards symbolically into the space surrounding the structure of the frame. Figuratively, oscillation and/or deformation of the branches of Figure 1 should alert us to something being amiss in our understanding. The template of the model structure ($[\mathbf{f}^0, \mathbf{h}^0]$) has caught on something of significance in the space of all the possibilities around it, as it is being navigated through the given behavior (B) of the real system. The indications from engagement between the two,

what is more, should direct our attention into specific avenues for discovery of the source of the anomalies, through the tagging devices of α and κ .




Figure 1. Archetypal branch-node network structure of a model, as a metaphor for problem construction and problem-solving: nodes represent state variables (\mathbf{x}); branches connecting states with states are associated with the model’s parameters (α); branches extending into the space around the structure are associated with the elements (κ) of the matrix \mathbf{K} in equation 2.

To conclude, then, our original question, and its earlier, more formal, interim expression, can be restated more precisely as follows:

Given a model of fixed structure $[\mathbf{f}^0, \mathbf{h}^0]$, given the parameterizations of the {presumed known (α)} and {acknowledged unknown (κ)}, and given descriptions of behavior over time, $B(t^{---}), B(t), B(t^+), B(t^{+++})$, how should we use the model to cope with the very high likelihood of changes in time of significance to our understanding and actions in the space between $[\mathbf{f}^0, \mathbf{h}^0]$ and the truth of the matter, namely $[\mathbf{f}^\infty, \mathbf{h}^\infty]$?

One could re-phrase this through yet another metaphor:

As we navigate over the unfolding terrain of perceived actual behavior, along a particular heading for the vehicle of Figure 1, what combination of radar and vehicle must we design in order to apprehend (detect) threats to our environment lying beyond the horizon?

To some extent this echoes our original metaphor, of the global system evolving along a trajectory, and we, with our models, are interested in identifying that trajectory — especially its future course, beyond our direct line of sight, and while the ground on which we are standing (the science base) is changing under our feet. But that is to offer too many metaphors, perhaps, hence to

confuse rather than clarify. We must also issue a disclaimer. The value of these metaphors — and indeed the formal, algorithmic arguments from which they have been extracted — resides in their capacity for unraveling problems of some considerable complexity in order (ourselves) to comprehend these problems. All metaphors will be limited in their usefulness in a literal sense. Moreover, we are not literally envisaging application of the associated recursive algorithms in order to resolve the more specific problems thereby revealed.

4. SOME CHALLENGES: TO WHICH WE HAVE A POTENTIAL TO RESPOND

The whole of this concern is not trivial. It remains both abstract and recalcitrant — all the metaphors notwithstanding — too monolithic to attack, as a whole. We might have a greater chance of responding to the following, then, as fragments chipped away from the whole:

- (i) Given the lay person's apprehension at the prospect of surprises and qualitative dislocations in future scenarios of environmental change — in particular, apprehension amongst those belonging to the richly creative Egalitarian social solidarity of Cultural Theory [Thompson, 1997] — how can the science base be mobilized in order to explore the plausibility of these hopes and fears for the future?
- (ii) In the presence of gross uncertainty, what are the key constituent hypotheses about the system's behavior, i.e., what are the key model mechanisms, on which the reachability of these surprises might turn — for we had better start purchasing more science in these particular domains as a priority?
- (iii) Is the system in imminent danger of a major dislocation which would lead to a feared pattern of future behavior?
- (iv) Is there any evidence in the narrow window of the empirical record to suggest that the system may already have embarked upon a path of collapse into a feared future; can we identify and diagnose the seeds of any imminent change in this record of the past?
- (v) Can we design our models with the express purpose of discovering our ignorance, and at the earliest possible moment?

(vi) Can we make useful forecasts of possible behavior patterns in the face of these extrapolated dislocations?

(vii) And how should we design regulatory policies for minimizing the reachability of feared future states while maximizing the reachability of desired states?

Yet even these are still not trivial questions; and they will need further sculpting before we can set them up formally as tractable computational problems (subsequently in section 5).

4.1 Programs of Exploration and Enquiry

This issue (of coping with structural change) is more subtle than merely waiting for larger computational platforms on which to mount ever larger models, or of merely accounting for more uncertainty in all its myriad dimensions. For consider this, as a caricature of what has been, over the decades, the primary program for generating environmental foresight:

- (i) **Program (1):** Include in the model everything of conceivable relevance; in particular, maximize use of the Laws of Physics in this composition; maximize thus the number of constants (parameters, coefficients) that are either universally known or measurable, independently, in the field or laboratory; minimize thereby the freedom arbitrarily to adjust the parameterization of the model in the light of the observed behavior of the prototype; juxtapose the output of the model with this empirical evidence; and make a singular (deterministic) prediction.

Driven in pursuit of the goal of achieving the {acknowledged unknown (κ)} $\rightarrow 0$, with the strongly implied concomitant of the {presumed known (α)} $\rightarrow \infty$, the accompanying objective is that of the search for immutable *invariance* in the parameters (α) — in part through the successive elimination of any need for κ . This seems an illusion; but it is one to which we are all drawn in an unending quest. This is not a program to be abandoned. Yet it is not a program without flaws [Shackley et al, 1998], not the least of which is that of mounting an empirical program to secure the unimaginably extensive B against which to reconcile gargantuan mobilizations of theory — without ambiguity [Beck, 2002].⁷

⁷ Not to mention the work of Oreskes [1998], in which it is argued that “predictive power is itself a fallible judge of scientific knowledge”.

Thus caricatured, Program (1) does not formally admit uncertainty into its framework. Such admission is defining for Program (2):

- (ii) **Program (2):** Acknowledge the uncertainty attaching to the set of prior model concepts and the empirical record of observed behavior; presume therein that the parameters of the model are constants, but not known precisely; in the process of reconciling observed and estimated system responses, employ the past observations to adjust and quantify the uncertainty attaching to the model, i.e., to its constituent parameters; and make an ensemble of predictions that are therefore intrinsically uncertain.

Everything is uncertain, to some degree. Arguably, only in the limit of reaching the truth of the matter, i.e., $[f^\infty, h^\infty]$, could uncertainty be said to have been eliminated. Much of the intent of this second program, however, has always been more pragmatic, very mindful of the metaphysics of this limit. Models are employed in the service of making decisions. If the preferred course of action remains so, in spite of all the uncertainty and all the analyses of the sensitivity of the decision to this uncertainty, what does it matter if the model's predictions are uncertain? What does it matter, then, if these uncertain predictions have been generated from an uncertain model, reflecting an ineradicable ambiguity in the capacity of the science base to explain past behavior? Stated thus, of course, it could mistakenly be argued that the two programs — in their entirety — have different objectives. Their common cause is this: the unending quest of Program (1) becomes the pursuit of eliminating uncertainty in the parameterization (α) of the {presumed known} in Program (2). The uncertainty in α becomes the logical link between the ever limited capacity to characterize past behavior — the fingerprint of any and all the distortions wrought in the structure of the model as it is reconciled with this observed behavior — and the making of uncertain projections therefrom into the future [Beck and Halfon, 1991]. The program, in fact, entertains the notion of multiple candidate parameterizations (j), i.e., α_j^i , set within multiple candidate model structures (i), i.e., $[f^i, h^i]$ [Beven, 2002]. Each of these parameterizations of α , nevertheless, is almost universally presumed to be a realization of a random variable that remains fundamentally *invariant* with time.

Caricatures are what they are. They do not encapsulate everything about the subject. They exaggerate certain features, here for the purposes of setting out the heritage of something we may

now define as a third program of modeling the behavior of the environment:

- (iii) **Program (3):** Derive qualitative statements about possible future patterns of behavior from systematic organization and manipulation of current (non-quantitative) beliefs; acknowledge that the model's parameters are the focal points in a map of the scientific partial knowns and unknowns about the behavior of the system; presume that these parameters will, in general, change with time; assess the candidate parameterizations and parametric changes enabling the given futures to be reached; and direct interpretation of the observations of past behavior in the light of performing this predictive task.

Put another way, the goal of Program (3) is the development of an approach and a set of methods that will enable us to examine the record of the past in a manner guided expressly by prior contemplation of the reachability of certain feared patterns of future behavior. The intent is to contribute to the process of generating environmental foresight [Science Advisory Board, 1995], albeit in literally a round-about manner, using a rear-view mirror to complement Schellnhuber's [1999] forward-view mirror, which he defines as "contemplation of the future by reflection on the past".

A model may be constructed for a variety of purposes: (i) as a succinctly encoded archive of contemporary knowledge; (ii) as an instrument of prediction (in support of making a decision or formulating a policy); (iii) as an exploratory vehicle for discovery of our ignorance; and (iv) as a device for communicating scientific notions to a scientifically lay audience. These do not map neatly one-to-one onto the three programs of modeling. But we can see the evolution of purpose through (i) through (iii) as we pass over the three programs. Above all, the third program is about models as exploratory vehicles for the discovery of our ignorance. Traveling along a parallel path, then, we can look back to see how characterization of the parameters (α) has shifted: from an invariant, precisely known constant, bestowing thereby absolute rigidity on the structure of the model; through parameters as random variables, allowing thus a degree of random vibrational play to enter into a structure, which yet manifestly retains a recognizably constant overall form; and on to parameters as stochastic processes, where now the structure may undergo plastic deformations into almost unrecognizably different forms and shapes (as we have imagined earlier as a

result of the invasions of exotic species into an ecosystem).

We have no intention herein of getting caught on the horns of the hoary tension between dogmatic preferences for the large and comprehensive — the higher-order models (HOMs), $[f^{+n}, h^{+n}]$ — and the small and “statistical”, the lower-order models (LOMs), $[f^{-n}, h^{-n}]$, that is. Addressing our concern needs all the richness of variety of perspective that can be brought to bear on it, part of which admittedly requires that models of such vastly different scope can ultimately be compared on much the same plane, which in turn is technically not a simple matter to achieve [Young and Parkinson, 2002]. There are bigger battles to be fought.

Last, withdrawal of the present (t^0) from of our grand concerns — as the dividing line between past and future behavior, collectively expressed above as $B(t^{---})$, $B(t)$, $B(t^+)$, $B(t^{+++})$ — is also significant, as we shall now see in the case studies to follow.

5. CASE STUDIES

5.1 Reachable Futures

Our first case study concerns a man-made impoundment, Lake Lanier, on the Chattahoochee River just to the north of Atlanta, Georgia, in the south-eastern United States [Beck et al, 2002c]. The lake is located in a rapidly urbanizing area, serves several purposes (power generation, water supply, recreation, wastewater assimilation), and its quality — or ecological integrity — is a matter of hotly debated interest to many stakeholders from a variety of backgrounds. The intensity of development in this particular conurbation can be gauged by the fact that a wastewater treatment plant, whose price might ordinarily have been some \$60M, has been constructed at a cost of \$260M — because one of the local counties not only draws its water supply from Lanier, but is facing the prospect of returning its reclaimed water back to Lanier. In short, since its creation (in 1958) the lake has received substantial amounts of nutrients from its surrounding watershed, most of which (it is believed) are “locked” in the bed of the lake’s sediments. The lake is seemingly about to be engulfed by peri-urban development. There is a palpable fear for the lake’s future [Cowie, 2001]; that things cherished today may not obtain 25 years hence.

Here then is the challenge. Ordinary people are concerned about the future, even over the longer term. Some of them — perhaps most interestingly

here members of the Egalitarian social solidarity of Cultural Theory [Thompson, 1997] — will have richly creative imaginations, regarding what may go wrong with the cherished piece of the environment (Lake Lanier, in our case). Assuming we have vehicles for encouraging expression *inter alia* of these barely plausible, possibly bizarre, surprising futures (both feared and hoped for) in as varied and variegated a formal statement of $B(t^{+++})$ as possible [Beck et al, 2002c; Varis, 2002], does the relevant science base endorse or undermine their plausibility? More formally, the following may be hewn from the rough block of our primal, essential, motivating question:

Given a high-order model (HOM) of fixed structure $[f^{+n}, h^{+n}]$, given the parameterization of essentially just the {presumed known (α)} alone, and given (lay) stakeholder-authored descriptions of behavior over future time $B(t^{+++})$, establish the plausibility of such behavior coming to pass, and identify the subset of key parameters (α^K) on which the reachability of $B(t^{+++})$ appears likely to turn.

Two futures were cultivated for Lanier, in fact: a desired, $B_d(t^{+++})$, and a feared future, $B_f(t^{+++})$ [Cowie, 2001; Osidele, 2001]. The science base of the lake’s foodweb was encoded into a model with some 13 states (x) and 100 or so parameters (α) [Osidele, 2001; Osidele and Beck, 2002],⁸ and the two — stakeholder imagination and scientific theory — reconciled through an algorithmic procedure combining regionalized sensitivity analysis [RSA; Spear and Hornberger, 1980], tree-structured density estimation [TSDE; Spear et al, 1994], and uniform covering by probabilistic rejection [UCPR; Klepper and Hendrix, 1994]. In familiar terms, suffice it to say there are therein various mechanizations of Monte Carlo simulation, i.e., random sampling of the parameter space, with the means to discriminate those parameters deemed redundant (α^R) from those deemed to be key (α^K) to determining whether or not behavior ($B(t^{+++})$) is thereby obtained.

Establishing the plausibility of the stakeholder-authored futures can be approached through two measures: (i) the simple ratio of the number of randomly generated candidate parameterizations giving the behavior, relative to the total number of samples; and (ii) the fraction of the total volume of the sampled parameter space containing the most dense clustering of behavior-giving candidate parameterizations. The former, it will be noted, is

⁸ We acknowledge this is hardly a HOM by many people’s standards; it is, nevertheless, of a high order relative to what is discussed below.

cast in terms of the external description of the system, while the latter attaches to its internal description. Our findings, on both accounts, are that the desired future, $B_d(t^{++++})$, appears some two to four times more likely to come to pass than the feared future, $B_f(t^{++++})$, all the gross uncertainty surrounding the computational analysis notwithstanding [Osidele, 2001; Osidele and Beck, 2002]. The ratio of total phosphorus to total suspended solids in the lake's tributary inflow, and the maximum intrinsic growth rate and recruitment rate constants for fish, are key parameters with respect to reachability of $B_f(t^{++++})$. The tributary phosphorus/particulate ratio (again), fish recruitment rate (again), and the coefficient for phosphorus diffusion across the sediment-water interface in the lake, are key to $B_d(t^{++++})$ being attained.

5.2 Identifying Priorities for Purchasing More Science

Our second case study, of the much smaller Lake Oglethorpe in Georgia, was conducted as a nursery prototype of the foregoing Lanier study [Beck et al, 2002a]. The environmental and social settings of the problem are not radically different from those of Lanier, but our challenge will be posed somewhat differently:

Given an HOM of fixed structure $[f^{+n}, h^{+n}]$, given the parameterization of essentially just the {presumed known (α)} alone, and given descriptions of behavior over time $B(t)$ and $B(t^{++++})$, identify and reflect upon the significance of change over time in the subsets of key parameters, $\alpha^K(t)$ and $\alpha^K(t^{++++})$, on which the reachability of $B(t)$ and $B(t^{++++})$, respectively, appears likely to turn.

Figure 2 shows a summary of the outcomes of our computational articulation of this problem [Beck et al, 2002b]. In this instance, target future behavior $B(t^{++++})$ is a fairly straightforward expression of what should be expected of the lake from a regulatory perspective (arguably, therefore, from the perspective of the Hierarchist social solidarity of Cultural Theory). The ecosystem model contains nine state variables, as denoted by the nine nodes in the branch-node network diagrams of Figure 2. Within that structure, dominance — the emboldened branches — can be seen to shift from the nutrient-autotroph-herbivore sub-web, $\alpha^K(t)$ (Figure 2(a)), to the sediment-nutrient-autotroph sub-web, $\alpha^K(t^{++++})$ (Figure 2(b)). Over time, between past and future, the significance of interactions with the herbivores has declined, while that of the interactions with the sediment has risen, from insignificance to significance (all the

gross uncertainty of the problem context notwithstanding, once again). Put simply, if meeting the regulatory behavior were urgent, more science had better be purchased in the domain of sediment-water interactions, as a priority. For there are not the funds to push forward on all the scientific fronts associated with the many other constituent hypotheses parameterised through $\alpha^R(t^{++++})$.




Figure 2. A form of apparent structural change (Lake Oglethorpe, Georgia): (a) key parameters ($\{\alpha^K(t)\}$; emboldened branches) associated with past behavior attach primarily to the nutrient-autotroph-herbivore sub-web; (b) key parameters ($\{\alpha^K(t^{++++})\}$; emboldened branches) associated with future behavior attach primarily to the sediment-nutrient-autotroph sub-web.

If we had been unaware of the complete web of interactions, by not engaging it in the structure of the model, the shift from one to the other of the networks of merely emboldened branches in Figure 2 would have *appeared* to us an evolution, or dislocation, in the structure of the system's behavior (just as in the changes in the Caribbean Sea and North Atlantic chronicled by Jackson [2001]). Figure 2 demands a rather more subtle interpretation, however. In our formal compositions of the challenges of both the Lanier and Oglethorpe case studies, no mention has been made of the {acknowledged unknown}. The approach to addressing these challenges presumes the analyst will, in general, seek to encode as much of the science base as possible in the model —

what we believe to be the knowns, the partial knowns, perhaps even speculation about the unknowns [Beck et al, 2002a]. The distinction between what is crisply and boldly hypothesized as “known” and what is speculatively the form of the “unknown”, is somewhat indistinct. This has more to do with the mechanization of the problem-solving approach, however, than any display of the presumption that as the {presumed known} $\rightarrow\infty$ so the {acknowledged unknown} $\rightarrow 0$. It may be helpful, therefore, to conceive of the structural shift in Figure 2 as follows. There are always microscopic flutter and drift in the behavior of the system, below the resolving power of the model (or outside its scope). Whereas such flutter and drift are always present, macroscopic behavior, nevertheless, might not previously (t) have been influenced by it to any significant extent; the point is, it may well become so, over t^{+++} . Dislocation is perceived as change over time in the extent to which behavior is sensitive to the various elements in the model’s assembly of constituent hypotheses [Chen and Beck, 2002].

5.3 Parametric Change as the Agent of Control

As we stand in the present, on the threshold of the future, it is of great interest to us to know whether something is about to turn upwards or downwards. Formally, we have this challenge to face:

Given a model of fixed structure [$\mathbf{f}^0, \mathbf{h}^0$], given the parameterization of essentially just the {presumed known (α)} alone, and given descriptions of behavior over time $B(t)$ and $B(t^+)$, determine what changes with time in α , i.e., $\alpha(t)$ over the interval of time [$t \leq t \leq t^+$] — in particular, more narrowly focused in the vicinity of the present t^0 — are required to transfer the state of system from display of behavior $B(t)$ to that of $B(t^+)$.

There is no case study with which to respond to this, as yet. But we have a method in mind, now to be outlined, and can fashion a worthy computational problem from the Lanier and Oglethorpe studies.

A generic description of a system’s behavior (such as equation 1) has inputs \mathbf{u} , states and parameters [\mathbf{x}, α], and output responses, \mathbf{y} . For simplicity, let us assume there is essentially no difference between \mathbf{x} and \mathbf{y} ; target future behaviour ($B(t^+)$) can be specified in terms of either, i.e., as $\bar{\mathbf{x}}_f(t^+)$. In general, some of the inputs may be considered disturbances simply impinging upon the system and beyond our manipulation, while others are

quantities whose values we can choose at will, which we call controls. Control theory concerns itself with answering the question: what pattern of controls (\mathbf{u}) over some span of time will transfer the state from its present given value to another desired value, subject to appropriate assumptions about the nature of the model’s parameters (α). Method upon method is available for solving this problem; and solutions may look like the trajectory in Figure 3 [Taylor, 1993], for the (input) fossil fuel flux required to stabilize atmospheric carbon dioxide concentration at some target level $\bar{\mathbf{x}}_f(t^+)$ by the year 2300. Our challenge transposes certain features within the archetypal question of control theory, to ask: first, what pattern of variations of α over time will bring about some specified change in the state of the system, subject to appropriate assumptions about the nature of the inputs (\mathbf{u}); and, then, what methods from control theory might be available to deliver an answer [see Keesman, 2002]? Parametric change has thus become, in principle, the agent of control.

Within the HOM of the Lanier case study it turns out that a knowledge of the coefficient for phosphorus diffusion across the sediment-water interface (α), for example, is crucial to reaching the desired domain of behavior in the future, but nothing more specific is known about the behavior of this parameter over time. Suppose, for the purposes of this thought experiment, that the striped bass population is to be maintained at some level, $\bar{\mathbf{x}}_p(t^+)$, that we can assume an array of watershed development scenarios and policy controls (\mathbf{u}), i.e., that they are in some sense “known”, and that the structure of the model is of but modest scope, as conveyed by [$\mathbf{f}^0, \mathbf{h}^0$]. We imagine, in Figure 4, it is not beyond the scope of the existing methods of control theory — for a modestly sized model — to generate the sequence of parametric changes, $\dots, \alpha(t_{k-1}), \alpha(t_k), \alpha(t_{k+1}), \dots$, required to transfer the status of the striped bass population over time from $\bar{\mathbf{x}}_p(t)$ to $\bar{\mathbf{x}}_p(t^+)$. In particular, once given the anchoring device of a future destination for behavior, some indication of the nature of the changes in α to be sought in the vicinity of the present (t^0), including the empirical evidence of the recent past, is thereby revealed. In other words, pointers as to the *specific* nature of the seeds of any imminent structural change can be extracted. We would thus have been provoked into looking into the record of the past from a specific perspective, conditioned upon having contemplated what it might take to attain a particular form of behavior in the *future*.

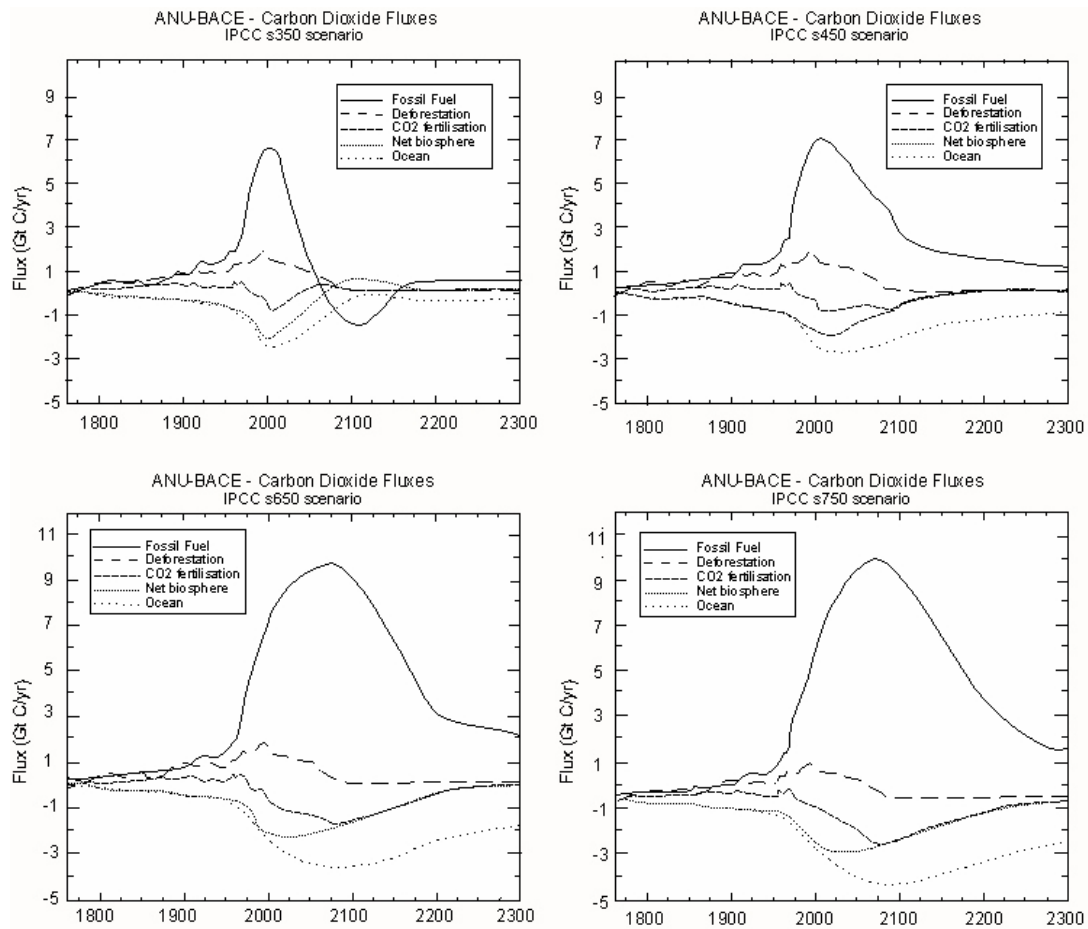


Figure 3. Estimated CO₂ fluxes from various sources required to follow four IPCC scenarios for atmospheric CO₂ concentration; results computed from the ANU-BACE model (reproduced with permission from Taylor [1993]).

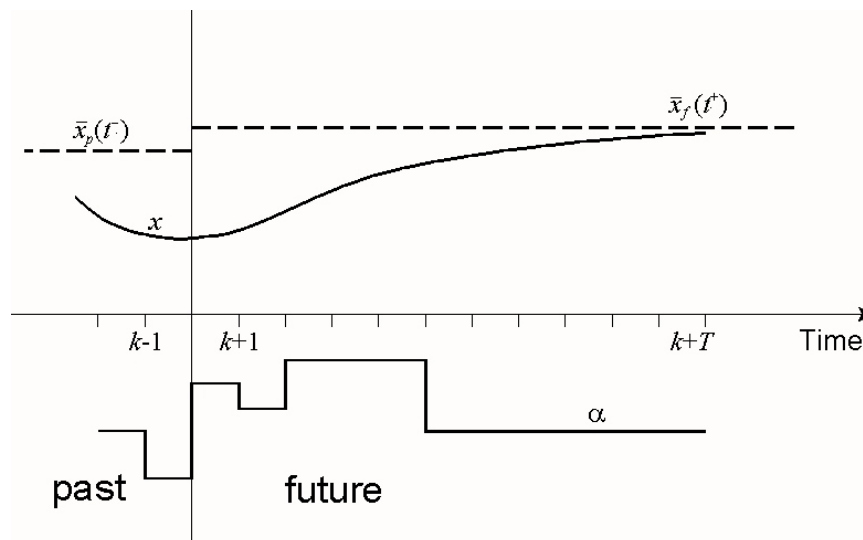


Figure 4. Putting a so-called principle of model-based predictive control [see Keesman, 2002] to work (conceptually) in solving the problem of computing those changes of the parameter values (α) required to transfer the state of the system (denoted as x) from its near-past/present position $\bar{x}_p(t)$ to some desired future position $\bar{x}_f(t^+)$.

5.4 Detecting the Seeds of Change: Elasto-plastic Deformation of Structure

Our fourth case study examines interference from the {acknowledge unknown} in the workings of the {presumed known} — from the putative existence of a population of phytoplankton (small organisms at the base of the aquatic foodweb) capable of intervening in the processes whereby easily degradable, waste organic matter is assimilated in a river, with or without damage to the health of that river's ecology. This will be our only challenge containing consideration merely of that which is retrospective, without any specific outlook on the future:

Given a model of fixed structure $[f^0, h^0]$, given the parameterizations of the {presumed known} (α) and {acknowledged unknown} (κ), and given a description of behavior over time $B(t)$, where the block of time t is defined by $[t_0 \leq t \leq t_N]$, reconstruct estimates of changes $\alpha(t)$ and $\kappa(t)$ in order to identify pointers to the nature of our ignorance in the space between $[f^0, h^0]$ and the truth of the matter, namely $[f^\infty, h^\infty]$.

Consider once again Figure 1. This, figuratively, is what is being attempted here: as the template of the model's structure is projected through the space of empirical experience of the real system, which constituent members of its frame — within what we think we know and within what we know we may not know — suffer distortion as a consequence of inconsistencies between hypothetical and empirical knowledge? For this we must mobilize the ideas prized open from the innovations representation of equation 2, implementing them through what is known as a Recursive Prediction Error (RPE) algorithm [Stigter, 1997; Stigter and Beck, 2002]. Our purpose is to test the worth not only of this solution procedure, but also of the above challenge, in this instance against a set of data from the River Cam in eastern England [Beck et al, 2002b].

Figure 5 is the outcome, where what is shown are the reconstructed changes with time, ..., $\hat{\alpha}^0(t_{k-1})$, $\hat{\alpha}^0(t_k)$, $\hat{\alpha}^0(t_{k+1})$, ..., and ..., $\hat{\kappa}^0(t_{k-1})$, $\hat{\kappa}^0(t_k)$, $\hat{\kappa}^0(t_{k+1})$, ..., in elements of *both* α and κ . The truth of the matter cannot be known in this instance, yet it is obvious from Figure 5(c) that something in the {acknowledged unknown} becomes significant at around t_{45} , as κ_1 abruptly assumes a non-zero value. Or it might be said that at this time the basis for the model's projection of immediate future behavior (one sampling interval ahead) veers from

a reliance on the {presumed known} to a reliance on the {acknowledged unknown}. That all three reconstructed estimates in Figure 5 remain relatively invariant after the deflection at about t_{45} gives us the basis for labeling this an inelastic, plastic deformation of the model's structure (although one has to be very careful in extracting such a conclusion from the confounding effects of the numerical implementation of the RPE algorithm; Beck et al, [2002b]). Think of this metaphor as follows. As the frame of the model's structure is subjected to the stresses induced, over time, by the loads arising from the mismatches between the hypothetical and empirical experiences of the real system, strains develop in the constituent members (hypotheses) of the model. Depending upon the material properties of these members (the levels of confidence attaching to the hypotheses), collapse may occur. Testing the model to the point of such explicit failure in its particular members is entirely the purpose of this challenge. We do *not* want to be left simply with the crude insight that the model — as a whole — is inadequate.




Figure 5. Reconstructed recursive estimates of three parameters in a candidate model for the assimilation of easily degradable organic matter in the River Cam: (a) parameter α_2^0 (day^{-1}); (b) parameter α_3^0 ($\text{gm}^{-3}\text{day}^{-1}$); and (c) gain matrix element κ_1^0 (dimensionless).




Figure 6. Pre-emptive projections into the future for the Cam case study: (a) using a model structure and candidate parameterisations (α_i^0) reconciled with (identified from) the most recent set of (past) data for the period t_{24} to t_{35} ; and (b) using a more refined structure (parameterized as α_i^1) that seeks to characterize the path by which the system came to be in its present state at t_{35} . Heavy, continuous lines represent projections based on a “best” candidate parameterization; thin, continuous lines are the envelopes of minimum and maximum projected values; and observations are indicated as +.

5.5 Design for Discovery of Ignorance

But what would we do, if we imagined ourselves in a present defined by something like t_{35} in Figure 5, with all the dislocation, if not surprise, in the behavior of the system still to come? Would we have the audacity (let alone the method), to ask the following. Is there any possibility whatsoever of detecting the onset of apparent structural change – in the domain of a model’s parametric space – before it manifests itself more palpably at the macroscopic level of the conventionally observed state variables? And if such change were detectable, could we entertain any hope of exploring the forecast consequences of propagating these seeds of change and dislocation into the future? Let us first state these challenges more formally as this:

Given a model of fixed structure $[f^0, h^0]$, given the parameterizations of the {presumed known α }, and given descriptions of behavior over time $B(t)$, reconstruct the path of variations in $\alpha(t)$, in order to identify the seeds of structural change in $B(t)$, thus to explore the

consequences of their further growth within $B(t^+)$.

What we need, from the past up to the present, is a mode of exploration — a combination of the vehicle and radar in the metaphor used earlier to explain Figure 1 — deliberately designed for discovery of our ignorance, and at the earliest possible moment. In other words, the manner in which the hypothetical science base is to be assembled into $[f^0, h^0]$, is done to that end.

We have not embarked on any such design as yet. We have, however, brought to bear all our *ad hoc* forensic methods of searching for the seeds of structural change [Chen and Beck, 2002] on the foregoing case study of the River Cam, with this result. Standing at about t_{35} in Figure 5, we have found that the behavior of water quality in the River Cam is subject to the workings of an unknown process, of potentially growing significance; that this something must generate both oxygen and easily degradable organic matter; and that its conjectured workings are in some way correlated with a declining streamflow and significant excursions of water temperature about

some nominal level (as possible input stimulae). Henceforth, the futures of Figure 6 are computable, under two subtly, but significantly different, conditions (and some strong assumptions).⁹ In this paler setting of our rather audacious challenge, nevertheless, there is significance. The projections of Figure 6(a) over t^+ , i.e., for $t \geq t_{36}$, which are not successful, have been generated using the most recent image of the candidate parameterization of the model at t_{35} , i.e., $\hat{\alpha}^0(t_{35})$. Those of Figure 6(b), which are more successful, are drawn from a revised, more refined, structure of the model $[\mathbf{f}^1, \mathbf{h}^1]$, based on a diagnosis and synthesis of the prior, accumulating forensic evidence of the “something untoward” being apprehended. They are therefore conditional upon the attaching more refined candidate parameterization, $\alpha^1(t_{35})$, as though this α^1 is a (hopefully) invariant parameterization of the parametric change discovered in $\alpha^0(t)$ (up to t_{35}) [Jakeman et al, 1994; Kokkonen and Jakeman, 2002]. Structure $([\mathbf{f}^1, \mathbf{h}^1])$, in short, embraces a hypothetical explanation of the *path* by which behavior — as perceived through the less refined (prior) structure $([\mathbf{f}^0, \mathbf{h}^0])$ — has evolved from the past (at t_0) up to the present (t_{35}). The difference is as between knowledge of the instant (the present) and knowledge of how we came to be in that present.

5.6 Inclined to Survive

This last challenge has to do with the following:

Given a low-order model (LOM) of fixed structure $[\mathbf{f}^n, \mathbf{h}^n]$, and given access to a description of behavior over a short window of time $B(t)$, determine the extent to which access to $B(t)$ changes the probability of reaching a terminal pattern of behavior $B(t^{+++})$.

Because the associated case study deals with the fate of populations of certain small rodents in the vicinity of Chernobyl, in the Ukraine, the intent of the challenge can be confused with determining the probability of the extinction of a species in a particular locale. As before, the words we use — “extinction”, “survival”, “collapse” [Kryazhimskii and Beck, 2002] — can cause difficulties. What we are about is this. Given a binary split of behavior in the future, for example, as between behavior being “essentially similar to” and

“radically different from” that of the past,¹⁰ are either of those future target domains more reachable, given the empirical record of the past relative to the maximally uninformed position of having no access whatsoever to such records? Put another way, given the tiny windows of empirical evidence typically available to us, in contemplating the unfolding of environmental problems with time constants significantly longer than the span of the window, can something meaningful be said of the entire mosaic, having observed the nature of just a single tile or two? Given the supposed, hypothetical trajectory of the system through time, does any observed inclination in the vicinity of the present, cause us to conclude that survival (or collapse, or whatever) is more or less likely than would otherwise be the case? If the likelihood is changed significantly as a result of access to the small window of empirical evidence, all the gross uncertainty notwithstanding, then the seeds of embarkation upon a path to the imagined future may be buried, and therefore deliberately to be sought after, somewhere in this evidence.

The method employed in response to this challenge takes its lead from the binary classification of behavior, compresses all the hypothetical knowledge of the science rather drastically into a (binary-valued) path-dependent stochastic process (the model $[\mathbf{f}^n, \mathbf{h}^n]$), and acknowledges substantial uncertainty in the parameterization of the model. In other words, it is accepted that the probability (r) of the next state transition being deemed “negative” (in the sense of an untoward current change in population status) and the threshold (β) of quantitative change in population size over any interval, above which the transition is so deemed, are parameters subject to wide ranges of possibilities. In the spaces of Figure 7, vertical axis w is the ratio of the conditional probability of reaching the “collapse” domain (C) in the future, i.e., $p\{C(t^{+++})|B(t)\}$, divided by the unconditional probability, i.e., $p\{C(t^{+++})\}$, given no access to the small window of empirical evidence. The surface represents the values of w resulting from all the various combinations of candidate parameterizations $[r, \beta]$, so that where that surface rises above the plane of $w = 1.0$, some inclination to collapse is signaled.

Figure 7 juxtaposes two outcomes of the analysis [Kryazhimskii and Beck, 2002]. For Figure 7(a), the window of evidence spans just the two years of

⁹ We must assume foreknowledge of the ensuing, observed, future, input stimulae, $\mathbf{u}(t^+)$; the foresight for which we struggle is always dogged by some such limitation.

¹⁰ Which is clearly a concept underlying all of the conceptual development (from Hornberger and Spear, 1980) culminating in many of the foregoing challenges.

annual observations for 1981/2 and 1982/3, while that for Figure 7(b) spans the four years of 1981/2, 1982/3, 1983/4, 1984/5. It is apparent how access to the broader window enhances the probability of coming to the view that the system (of rodent populations) was destined to enter a collapse state in the longer term (beyond 1986) prior to the Chernobyl accident (of 1986). The surface of Figure 7(b) is subject to a more widespread upheaval above the plane of $w = 1.0$ than that of Figure 7(a). Having contemplated thus the approach of the system to a distant future state, should we thereby be provoked into re-examining the record of the past in a different light?

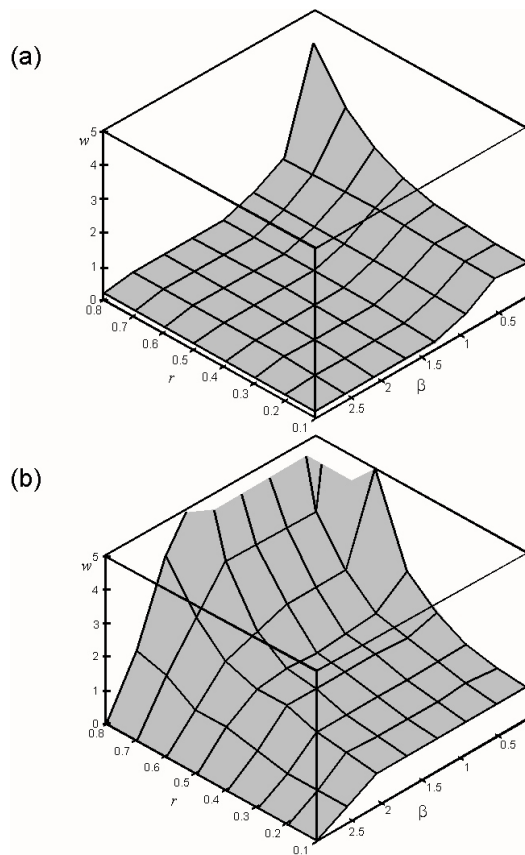


Figure 7. “Collapse indicator” surfaces for projections of future outcomes for a rodent population in the vicinity of Chernobyl: (a) given access to observed population levels for 1981/2 and 1982/3; and (b) given access to observed population levels for 1981/2, 1982/3, 1983/4, and 1984/5.

6. CONCLUSIONS

It will not have escaped notice that this paper has been written as a synopsis of a recently published book [Beck, 2002]. Conspicuous by its absence from that book was a reasonably all-encompassing

case study, which that described above of Lake Lanier in Georgia, USA, has become (and not especially by design). All this “foresight” should be “foresight for action”. Does that mean we have fashioned a better framework for action? Not really, for the recursive, predictive framework at the core of the formal procedures of this paper (epitomized in equation 2) is merely another expression of adaptive management and control [Holling, 1978; Åström and Wittenmark, 1989]. And it is hard — very hard in fact — to better these concepts, although we have tentatively suggested embedding them in a somewhat broader framework we have called “adaptive community learning”, as a result of the experience of the Lanier case study [Beck et al, 2002c].

When the Science Advisory Board of the US Environmental Protection Agency called for the means to generate environmental foresight [Science Advisory Board, 1995], it expected largely the following kind of response: an area is to be identified in which scientific data are sparse and/or in conflict and the scientist conducting the enquiry is to submit — to the process of scientific peer review — an opinion on the interpretation of the data as portending some threat to the environment. In other words, the extant historical record gathered within the paradigm of (normal) scientific enquiry is to be examined and interpreted by a practising scientist whose opinion will be judged by other practising scientists. Put cryptically, such a foresight-generating framework taps into the combination of {scientific empirical observations & scientific opinion}. Worthy and necessary though this is, it is not the only thing that could be done. Cast in like terms, our procedure of adaptive community learning — as now defined and as encompassing an alternative line of enquiry [Beck et al, 2002c] — draws upon a combination of {scientific models & stakeholder imagination}. It differs in both elements from that of {scientific empirical observations & scientific opinion}, offering thus a wider search for the possibility of surprises and being undoubtedly eclectic, if unconventional, in the sources of information into which it taps. It has about it a whiff of the public directing upon which issues the torchlight of scientific enquiry is to be shone. In this scheme of things, someone — and someone *essentially other* than those mobilizing the science base into the computational model — is imagining the future, not predicting it. In the process of exploring the reachability of this imagined target future, and the implications thereof, assurance of quality is needed in the dimension of judging whether the model has been well or ill designed with respect to the task, in particular, of

discovering our ignorance at the earliest possible moment.

So it is, therefore, that in a paper about “foresight” almost no mention has been made of “prediction”. For a paper addressing so obviously the future, fundamentally within the framework of computational models, it is perhaps just as surprising that the word “validation” — or evaluation [Oreskes, 1998], or quality assurance [Beck and Chen, 2000], or quality assistance [Risbey et al, 2002] — has barely surfaced, if at all. Directing a model at the goal of “design for discovery of ignorance” has only just arisen, as far as we are aware; judging the quality of the process of going about this purpose remains, therefore, an open question. This notion of mobilizing models as vehicles of exploration, furthermore, shifts attention away from the arguably more familiar idea of piling on empirical assessment after empirical assessment in order (eventually) to converge upon the invariance of the (unknowable) truth in the structure of the model.

In many ways, what has been set out in this paper, as a program for coping with structural change while generating environmental foresight, can be cast comfortably within the “participatory, inverse approach” to problem-solving that would constitute a Sustainability Science [Kates et al, 2001]. Unquestionably the origins of any advances that have been made across the paper have been steeped in the discussion of what has been called Post Normal Science [Funtowicz and Ravetz, 1990; 1993]. These advances have their reflections in the works of Robinson [1988] and Kieken [2002], for example.

The focus was once on analyzing uncertainties presumed to be reducible (in the end) in a world view of “constancy” [Beck, 1987]. Attention herein has been channeled into an outlook of a more fluid kind: onto a problematique in which the science base is itself evolving as much as are the hopes and the fears of the lay community for their cherished pieces of the environment, on which our close scientific scrutiny is to be trained.

7. ACKNOWLEDGMENTS

I am indebted to all those who participated in the International Task Force on Forecasting Environmental Change, who wrote the book of that Task Force (*Environmental Foresight and Models: A Manifesto*), from which the account of this paper has been composed. I am especially grateful to three former graduate students: Jining Chen, Hans Stigter, and Femi Osidele. Part of the research on

which this paper is based has been funded by the Water and Watersheds Program of the US Environmental Protection Agency, Grant # R825758. I am deeply indebted to this support. I also wish to acknowledge my current appointment as Visiting Professor in the Department of Civil and Environmental Engineering at the Imperial College of Science, Technology, and Medicine in London.

8. REFERENCES

- Allen, P.M., Evolution, innovation, and economics, in G. Dosi, C. Freeman, R. Nelson, G. Silverberg and L. Soete (eds.), *Technical Change and Economic Theory*, 95-119, Pinter, London, 1990.
- Åström, K.J. and B. Wittenmark, *Adaptive Control*, Addison Wesley, 1989.
- Beck, M.B., Water quality modeling: a review of the analysis of uncertainty, *Water Resources Research*, 23(8), 1393-1442, 1987.
- Beck, M.B., *Environmental Foresight and Models: A Manifesto*, Elsevier, Oxford, 2002.
- Beck, M.B. and J. Chen, Assuring the quality of models designed for predictive tasks, in A. Saltelli, K. Chan, and E.M. Scott (eds.), *Sensitivity Analysis*, 402-420, Wiley, Chichester, 2000.
- Beck, M.B. and E. Halfon, Uncertainty, identifiability and the propagation of prediction errors: a case study of Lake Ontario, *J. Forecasting*, 10(1/2), 135-161, 1991.
- Beck, M.B., J. Chen and O.O. Osidele, Random search and the reachability of target futures, in M.B. Beck (ed.), *Environmental Foresight and Models: A Manifesto*, 207-226, Elsevier, Oxford, 2002a.
- Beck, M.B., J.D. Stigter and D. Lloyd Smith, Elasto-plastic deformation of the structure, in M.B. Beck (ed.), *Environmental Foresight and Models: A Manifesto*, 323-350, Elsevier, Oxford, 2002b.
- Beck, M.B., B.D. Fath, A.K. Parker, O.O. Osidele, G.M. Cowie, T.C. Rasmussen, B.G. Norton, A. Steinemann, B.C. Patten, S.R. Borrett, D. Cox, M. Mayhew, X-Q. Zeng and W. Zeng Developing a concept of adaptive community learning: case study of a rapidly urbanizing watershed, *Integrated Assessment*, (submitted), 2002c.
- Beven, K.J., Uncertainty and the detection of structural change in models of environmental systems, in M.B. Beck (ed.), *Environmental Foresight and Models: A Manifesto*, 227-250, Elsevier, Oxford, 2002.

- Brooks, H., The typology of surprises in technologies, institutions, and development, in W.C. Clark and R.E. Munn (eds.), *Sustainable Development of the Biosphere*, 325-348, Cambridge University Press, Cambridge, 1986.
- Chen, J. and M.B. Beck, Detecting and forecasting growth in the seeds of change, in M.B. Beck (ed.), *Environmental Foresight and Models: A Manifesto*, 351-373, Elsevier, Oxford, 2002.
- Cowie, G.M., Foresight for Lanier: Summary Results of a Workshop, Warnell School of Forest Resources, University of Georgia, Athens, Georgia, 2001.
- Dennis, R.L., The ozone problem, in M.B. Beck (ed.), *Environmental Foresight and Models: A Manifesto*, 147-165, Elsevier, Oxford, 2002.
- Funtowicz, S.O. and J.R. Ravetz, *Uncertainty and Quality in Science for Policy*, Kluwer, Dordrecht, 1990.
- Funtowicz, S.O. and J.R. Ravetz, Science for the post normal age, *Futures*, 25(7), 739-755, 1993.
- Holling, C.S., *Adaptive Environmental Assessment and Management*, Wiley, Chichester, 1978.
- Holling, C.S., The resilience of terrestrial ecosystems: local surprise and global change, in W.C. Clark and R.E. Munn (eds.), *Sustainable Development of the Biosphere*, 292-320, Cambridge University Press, Cambridge, 1986.
- Holling, C.S., Engineering resilience versus ecological resilience, in P.C. Schulze (ed.), *Engineering Within Ecological Constraints*, 31-43, National Academy of Engineering, Washington DC, 1996.
- Hornberger, G.M., Impacts of acidic atmospheric deposition on the chemical composition of stream water and soil water, in M.B. Beck (ed.), *Environmental Foresight and Models: A Manifesto*, 131-145, Elsevier, Oxford, 2002.
- Hornberger, G.M. and R.C. Spear, Eutrophication in Peel Inlet, I, problem-defining behaviour and a mathematical model for the phosphorus scenario, *Water Research*, 14, 29-42, 1980.
- Jackson, J.B.C., What was natural in the coastal oceans?, *Proceedings National Academy of Sciences*, 98(10), 5411-5418 (8 May), 2001.
- Jakeman, A.J., D.A. Post, and M.B. Beck, From data and theory to environmental model: the case of rainfall runoff, *Environmetrics*, 5(3), 297-314, 1994.
- Kates, R.W., W.C. Clark, R. Corell, J.M. Hall, C.C. Jaeger, I. Lowe, J.J. McCarthy, H.J. Schellnhuber, B. Bolin, N.M. Dickson, S. Faucheux, G.C. Gallopin, A. Grübler, B. Huntley, J. Jäger, N.S. Jodha, R.E. Kasperson, A. Mabogunje, P. Matson, H. Mooney, B. Moore, T. O'Riordan and U. Svedin, Sustainability science, *Science*, 292, 641-642 (27 April), 2001.
- Kauffman, S., *At Home in the Universe*, Oxford University Press, Oxford, 1995.
- Keesman, K.J., Parametric change as the agent of control, in M.B. Beck (ed.), *Environmental Foresight and Models: A Manifesto*, 415-424, Elsevier, Oxford, 2002.
- Kieken, H., Integrating structural changes in future research and modelling on the Seine River basin, in Integrated Assessment and Decision Support, Proceedings First Biennial Meeting of the International Environmental Modelling and Software Society, Lugano, Italy, 2002.
- Klepper, O. and E.M.T. Hendrix, A method for robust calibration of ecological models under different types of uncertainty, *Ecological Modelling*, 74, 161-182, 1994.
- Kokkonen, T.S. and A.J. Jakeman, Structural effects of landscape and land use on streamflow response, in M.B. Beck (ed.), *Environmental Foresight and Models: A Manifesto*, 303-321, Elsevier, Oxford, 2002.
- Kryazhinskii, A.V. and M.B. Beck, Identifying the inclination of a system towards a terminal state from current observations, in M.B. Beck (ed.), *Environmental Foresight and Models: A Manifesto*, 425-451, Elsevier, Oxford, 2002.
- Leggett, J.K., The threats of climate change: a dozen reasons for concern, in J.K. Leggett (ed.), *Climate Change and the Financial Sector: The Emerging Threat, The Solar Solution*, 27-57, Gerling Akademie Verlag, Munich, 1996.
- Ljung, L., Asymptotic behaviour of the Extended Kalman Filter as a parameter estimator, *IEEE Transactions on Automatic Control*, 24, 36-50, 1979.
- Matthews, D.A., S.W. Effler, C.M. Brooks Matthews, C.A. Siegfried and M.E. Spada, Responses of Onondaga Lake, New York, to early stages of rehabilitation: unanticipated ecosystem feedbacks, *Water Environment Research*, 73(6), 691-703, 2002.
- Oreskes, N., Evaluation (not validation) of quantitative models, *Environmental Health Perspectives*, 106(Supplement 6), 1453-1460, 1998.
- Osidele, O.O., Reachable futures, structural change, and the practical credibility of environmental simulation models, PhD Dissertation, Warnell School of Forest

- Resources, University of Georgia, Athens, Georgia, 2001.
- Osidle, O.O. and M.B. Beck, Integrating stakeholder imagination with scientific theory: a case study of Lake Lanier, USA, in *Integrated Assessment and Decision Support*, Proceedings First Biennial Meeting of the International Environmental Modelling and Software Society, Lugano, Switzerland, 2002.
- Price, M.F. and M. Thompson, The complex life: human land uses in mountain ecosystems, *Global Ecology and Biogeography Letters*, 6, 77-90, 1997.
- Risbey, J., J. van der Sluijs, P. Klopogge, J.R. Ravetz, S.O. Funtowicz and S. Corral Quintana, Application of a checklist for quality assistance in environmental modelling to an energy model, in *Integrated Assessment and Decision Support*, Proceedings First Biennial Meeting of the International Environmental Modelling and Software Society, Switzerland, Italy, 2002.
- Robinson, J.B., Unlearning and backcasting: rethinking some of the questions we ask about the future, *Technological Forecasting and Social Change*, 33, 325-338, 1988.
- Schellnhuber, H.J., Earth system analysis and the second Copernican Revolution, *Nature*, 402(Supplement), C19-C23, 2 December, 1999.
- Schertzer, W.M. and D.C.L. Lam, Lake Erie and the evolving issues of the quality of its water, in M.B. Beck (ed.), *Environmental Foresight and Models: A Manifesto*, 105-130, Elsevier, Oxford, 2002.
- Schneider, S.H. and S.L. Thompson, Future changes in the atmosphere, in R. Repetto (ed.), *The Global Possible*, Yale University Press, New Haven, Connecticut, pp 397-430, 1985.
- Science Advisory Board, Beyond the horizon: using foresight to protect the environmental future", Report EPA-SAB-EC-95-007, Science Advisory Board, US Environmental Protection Agency, Washington, DC, 1995.
- Shackley, S., P.C. Young, S. Parkinson and B. Wynne, Uncertainty, complexity and concepts of good science in climate change modelling: are GCMs the best tools?, *Climatic Change*, 38, 159-205, 1998.
- Spear, R.C., and G.M. Hornberger, Eutrophication in Peel Inlet, II, identification of critical uncertainties via generalised sensitivity analysis, *Water Research*, 14(1), 43-49, 1980.
- Spear, R.C., T.M. Grieb, and N. Shang, Parameter uncertainty and interaction in complex environmental models, *Water Resources Research*, 30(11), 3159-3169, 1994.
- Stigter, J.D., The development of a continuous-discrete recursive prediction error algorithm in environmental systems analysis, PhD dissertation, University of Georgia, Athens, Georgia, 1997.
- Stigter, J.D. and M.B. Beck, Application of a continuous-discrete recursive prediction error algorithm as a model evaluation tool, *Mathematical Biosciences* (submitted), 2002.
- Strayer, D.L., N.F. Caraco, J.J. Cole, S. Findlay and M.L. Pace, Transformation of freshwater ecosystems by bivalves: a case study of zebra mussels in the Hudson River, *BioScience*, 49(1), 19-27, 1999.
- Taylor, J.A., Fossil fuel emissions required to achieve atmospheric CO₂ stabilisation using ANU-BACE: a box diffusion carbon cycle model, in M.J. McAleer and A.J. Jakeman (eds.), *Proceedings International Congress on Modelling and Simulation*, 2, 765-770, 1993.
- Thompson, M., Cultural theory and integrated assessment, *Environmental Modeling and Assessment*, 2, 139-150, 1997.
- Thompson, M. and P. Tayler, The surprise game: an exploration of constrained relativism, *Warwick Papers in Management*, 1, School of Industrial and Business Studies, University of Warwick, Warwick, UK, 1986.
- Thompson, M., R. Ellis and A. Wildavsky, *Cultural Theory*, Westview, Boulder, Colorado, 1990.
- Varis, O., Belief networks: generating the feared dislocations, in M.B. Beck (ed.), *Environmental Foresight and Models: A Manifesto*, 169-205, Elsevier, Oxford, 2002.
- Weaver, A.J., Driving the ocean conveyor, *Nature*, 378, 9 November, 135-136, 1995.
- Young, P.C. and S. Parkinson, Simplicity out of complexity, in M.B. Beck (ed.), *Environmental Foresight and Models: A Manifesto*, 251-301, Elsevier, Oxford, 2002.