Support

The conference has been financially supported by the following institutes and research programmes:

Commission of the European Communities, Directorate for Science, Research and Development, Climatology Programme

Dutch Ministry of Housing, Physical Planning and Environment, National Research Programme "Global Air Pollution and Climate Change" (NOP)

Netherlands Organization for Scientific Research, Priority Program "Nonlinear Systems"

Royal Netherlands Academy of Sciences

Wageningen Agricultural University, as part of its programme to commemorate its 75th Anniversary.

Sponsors

Shell Nederland bv

Unilever

Predictability and Nonlinear Modelling in Natural Sciences and Economics

edited by

J. Grasman

Department of Mathematics,
Agricultural University,
Wageningen, The Netherlands

and

G. van Straten

Department of Agricultural Engineering and Physics,
Agricultural University,
Wageningen, The Netherlands

KLUWER ACADEMIC PUBLISHERS
DORDRECHT / BOSTON / LONDON

UNDERSTANDING UNCERTAIN ENVIRONMENTAL SYSTEMS

M. B. BECK

Warnell School of Forest Resources, University of Georgia
Athens, Georgia 30602-2152, USA

Summary

Developments over the past two decades in the identification of models of environmental systems are reviewed, with special reference to the quality and pollution of surface freshwaters. As in so many fields, the early 1970s were a time of great expectations: it would not be long, we believed, before the admittedly less well defined problems of environmental systems analysis would nevertheless yield to the already vast array of methods available from applied mathematics and control theory (which had been so successful in their application, for example, to the analysis of aerospace systems). Such a yielding has still to come to pass, at least for multivariable models of more than, say, five or six state variables. In the past decade, because of the seemingly insuperable difficulties of model identifiability, we have promoted the pragmatic view that what really matters is the ability to generate "robust" predictions that are maximally insensitive to a lack of identifiability. Such pragmatism, coupled with a continuing dearth of successful techniques of system identification, does not bode well. The digital computing technology on which we are able to realise our "set of concepts" (our models) continues to expand rapidly. A similar expansion, although less dramatically so, is apparent in the technology of instrumentation and remote sensing, through which our "given data" are acquired in ever greater volumes. No such expansion is evident in the capacity of the brain to juggle with disparate facts and figures until the ever more comprehensive, given data can be reconciled with the increasingly massive sets of concepts. Whither, then, is environmental system identification bound in the next decade? A modest attempt to answer this question will be made, by way of conclusion.

Keywords Kalman filter, system identification, identifiability, predictability, uncertainty, water pollution

1. Introduction

In a recent article -- on interactive computing as a teaching aid - MacFarlane (1990) has presented a three-element characterisation of knowledge. According to the American philosopher Lewis these three elements are (as reported by MacFarlane):

(i) the given data;

(ii) a set of concepts; and

(iii) acts which interpret data in terms of concepts.

It is readily apparent that the problem of system identification (the derivation of a model whose behaviour bears the closest possible resemblance to the observed behaviour of the actual system) is covered exactly by the third of the three elements. What then are the prospects for success in these "acts which interpret data in terms of concepts", that is, for success in reconciling the candidate model with the given data? In other words, what are the prospects for progress in understanding uncertain environmental systems?

From contemporary experience we know that the scope and resolution of both the "given data" and the "set of concepts" necessary for understanding the behaviour of environmental systems are expanding at an increasing rate. We know too that this rate of expansion, if anything, is greater in respect of the set of concepts (the model) . The General Circulation Models (GCMS) of climatology and meteorology are of suitably massive proportions, with typically seven or more state variables (wind velocities, air temperature, and so on; Folland *et al.*, 1991) to be accounted for at over 10^5 spatial locations -- and doubtless soon to be still more. The computing capacity now available for realising models of the behaviour of a system offers us a truly staggering, expanding universe of possibilities. Indeed, in the popular scientific press this potential is neatly captured in headlines such as "Is It Real, Or Is It A Cray?" (Pool, 1989) or "Speculating In Precious Computronium" (Amato, 1991).

Technical support for manipulating the logical consequences of our "set of concepts" is thus assured. Technical support for acquiring the "given data" is likewise assured, although it might always be argued to be (relatively) inadequate. For example, in the area of modelling ocean circulation patterns it has been said that the rate of expansion in computing power (for realising ever more refined models) will bring about a need for a ten-fold increase every six years in data acquisition (for defining boundary and initial conditions). Given that US$2 billion per annum are reported to be required to service the current data-retrieval systems for monitoring ocean circulation, it is almost inconceivable that access to data will ever expand at a rate faster than the access to computing power.

No such assurances as these exist for development of the technical support necessary for engaging the model in a meaningful interpretation of the data. Indeed, how does such "interpretation" come about? It is a result of juggling with, and sifting through, a unique assortment of disparate facts and figures assembled by the individual, upon which some kind of order is eventually imposed. It is a subjective mental process. News of advances in computational capacity is abundant; news of advances in the technology of instrumentation and remote sensing is commonplace; news of the *increasing* capacity of the brain to juggle with disparate facts and concepts is non-existent.

Furthermore, what form of technical support would be desirable for promoting, provoking, or stimulating acts that interpret the data in terms of a set of concepts? This review of the tentative attempts to answer such questions is organised on a simple chronological basis, beginning with the hope of enlightenment in the 1970s,, passing through the clouds of uncertainty gathering during the 1980s and now, in the 1990s, looking forward to the prospect of rekindled hopes of further enlightenment.

2. Enlightenment: a better sense of the problems

Few of us -- working in the early 1970s -- could have guessed at the richness of "paradigms" now available for description and computerbased realisation of our theories about the behaviour of environmental systems. There are, for example, the following options: of classical differential calculus; of qualitative simulation (and the calculus of fuzzy logic); of cellular automata; and of pictorial simulation. All but the first of these have either been enabled or profoundly influenced by developments of just the past decade in the hardware and software of electronic computing. All are applicable, in principle, to the characterisation of problems associated with the contamination of surface water systems (Camara *et al.*, 1987, 1990; Castro *et al.*, 1993).

But such choice was not available to us two decades ago. Systems were conceived of as assemblies of mechanisms, characterised by state variables (x) subjected to input load disturbances (u) that generate various forms of output response (y). It was almost beyond question that description of their behaviour should take the form of a set of algebraic and/or differential equations. In terms of Lewis's three-element characterisation of knowledge, albeit perhaps with some licence:

(i) the input-output data $\{u,y\}$ constituted the external description of the system's behaviour and were the "given data";

(ii) the states and parameters (α) of the model, i.e., $\{x,\alpha\}$, constituted the internal description of the system's behaviour and were therefore a formal realisation of the "set of concepts"; and

(iii) calibration of the model was the "act which interpreted the data in terms of the set of concepts".

In respect of this last, it was generally assumed that all of the appropriate constituent hypotheses of the model would already have been assembled and correctly expressed; all that remained was to tune the parameters of the model, as an instrument requiring calibration for subsequent prediction. In retrospect, of course, one can ask whether the fine tuning of an already well structured instrument is the essence of progress in understanding, especially in the presence of gross uncertainty (attaching both to one's prior theories and to the observations of behaviour).

This restatement of Lewis's characterisation of knowledge presumes an important distinction between the input-output space $\{u,y\}$ and the state-parameter space $\{x,\alpha\}$.

Such a distinction is important on two accounts. First, the recourse to a state-parameter space description of the system's behaviour suggests that the objective is indeed to reconcile an assembly of constituent hypotheses (the model, or the set of concepts) with a set of observations. This does not suggest that a model cast in the input-output space is devoid of, or unrelated to, a set of concepts -- as we shall see later -- but that it may not ultimately be the most useful vehicle for interpretation of the field data. The inputoutput models of time-series analysis are rather primitive vehicles for such interpretation. They might best be used as the means of preparing the data for subsequent interpretation through some other form of model (Beck, 1991). Second, for reasons of academic discipline, a state-parameter space description is by far the more popular form of expression of the models of environmental systems analysis.

Choosing thus a middle course, between the partial differential equations of what might frequently be referred to as a "physicsbased" model and the algebraic, discrete-time equations of the "black-box" models of control theory and time-series analysis, we have the following form of "conceptual" model

$$\dot{x}(t) = f\{x,u,\alpha;t\} + \xi(t) \qquad (1a)$$

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

in which the dot notation in $\dot{x}(t)$ denotes differentiation with respect to time t, ξ is a vector of unknown disturbances (or errors of model structure, or errors of observation of u) associated with the state vector dynamics and η is a vector of errors associated with observations of y, assumed pragmatically to have been made at discrete instants in time t_k.

2.1 Filtering theory as the conceptual framework

One can look at the problem of calibration as a matter of signal processing. Such a view is entirely sympathetic to the notion of models as vehicles ' for the interpretation of data. And the classic solution to the problem of reconstructing information about $\{x,\alpha\}$ from information about $\{u,y\}$ in the presence of uncertainty (or noise), as all control theorists will know, is the filtering theory of Kalman (Kalman, 1960; Kalman and Bucy, 1961). Quintessentially, the filter reconciles a prediction from the model with an observation of the system through a process of feedback (to the model!), in which the account taken of the mismatch between theory and observation is modulated according to the balance of uncertainties attaching to these two "elements of knowledge".

Over the years filtering theory has come to reflect something of a universal framework for exploring and formally defining -- yet not necessarily solving (as will become readily apparent) -- many of the most interesting sub-problems of uncertainty, identifiability and predictability. This is, however, a very personal, perhaps idiosyncratic, view (Beck, 1987).

Within this framework, and with reference to the model of equation (1), the problem of calibration can formally be defined as the problem of combined state-parameter estimation:

Problem #1: State-parameter estimation

Given the observations and the model structure, i.e., given $\{u,y;f,h\}$, and given assumptions about the various sources of uncertainty, i.e., $\{P_{xx}(t_0), P_{\alpha\alpha}(t_0), S(t_k), Q(t_k), R(t_k)\}$, determine the best estimates of $\{x,\alpha\}$ and the uncertainty attaching to these estimates, i.e., $\{P_{xx}(t_N), P_{\alpha\alpha}(t_N)\}$.

Here, P_{xx} is defined as the variance-covariance matrix of the state estimation errors, $P_{\alpha\alpha}$ is the variance-covariance matrix of the parameter estimation errors, S and R are respectively the variance-covariance matrices of the input and output observations, and Q is the variance-covariance matrix of the process (excluding the errors now accounted for explicitly under S). t_0 and t_N represent respectively the discrete instants of time at the beginning and end of the observation period, with k = 0, 1, 2, ..., N sampling instants. In fact, to be pedantic, prior estimates of $\{x,\alpha\}$ at t_0 must also be assumed in the above problem definition.

The obvious flaw in this statement of calibration -- as the *first* act of reconciling the model with the data -- is that the model structure, as reflected in (f,h) (and also in the choices of the elements in $\{x,\alpha\}$), is not given a priori. *Before* determination of a best $\{x,\alpha\}$ the problem of model structure identification, defined formally as follows, must be addressed:

Problem #2: Model Structure Identification

Given the observations $[u,y]$ and given assumptions about the various sources of uncertainty, i.e., $\{P_{xx}(t_0), P_{\alpha\alpha}(t_0), S(t_k), Q(t_k), R(t_k)$ determine $[f,h;x,\alpha]$ and the accompanying $\{P_{xx}(t_N), P_{\alpha\alpha}(t_N)\}$.

Between these two statements (of state-parameter estimation and model structure identification), the degree of belief in the likely success of our prior theories in describing in general the behaviour of an environmental system has diminished. The burden of specifying $[f,h]$ correctly has shifted, from an almost complete reliance on prior theory and conjecture, to some engagement of the field observations in this process (which would in fact be more consistent with Lewis's three-element characterisation of knowledge). This change, and elucidation of the problem of model structure identification, did not arise from any philosophical consideration, but rather from a case study of organic waste degradation in a stretch of lowland river in eastern England (Beck and Young, 1976).

In short, the instrument of prediction may need more than just finetuning; it may need substantial re-design. The questions of interest to reconciliation of the model with the data are: which design is closest to the "truth"; how can we approach this "truth" at the fastest possible rate from some starting point; and what is a useful starting point for the prior model structure?

Exactly how one might go about answering these questions -- within the framework of filtering theory -- is summarised in Beck (1986, 1987) . The fact that one can obtain recursive estimates $[\{\hat{x}(t_k), \hat{\alpha}(t_k)\}]$ across the period of the observed record from t_0 to t_N has been crucial, however, to insights about the nature of the problem of model structure identification. And this capacity to provide temporally varying parameter estimates is in turn the distinctive feature of a filtering-like algorithm, which thereby sets it apart from any other approach to a solution of this problem. But like any other approach, the filter cannot directly identify the "true" structure $[f,h]$, since this implies, inter alia, a means of estimating integer values for the numbers of differential and algebraic equations in the model. Its distinction lies in revealing unreasonable fluctuations in the recursive estimates of $\hat{\alpha}(t_k)$ that result from significant discrepancies between the structure underlying the observed field data $[f,h]$ and the structure of the candidate model, let us say $[f',h']$. In fact, we would hope these fluctuations are only superficially "unreasonable" and that behind them lies a plausible explanation.

Since individual parameters relate to constituent model hypotheses we have, in principle, a means to establish the "success" or failure of these *individual* hypotheses (as opposed to the more customary assessment of whether the model *as a whole* succeeds or fails) . Furthermore, after working on a number of case study problems (Beck and Young, 1976; Beck, 1982, 1985) it was possible to distil out a more systematic organising principle for the procedure of model structure identification, which comprised the following elementary questions of (Beck, 1986):

(i) how to expose the *failure* (inadequacy) of the constituent hypotheses of a model structure;

(ii) how to *infer* the form of an improved model structure from diagnosis of the failure of an inadequate structure?

Preparing tests of the model structure in order to answer these questions can again be realised within the framework of filtering theory. Respectively:

(i) It can be assumed that the constituent parameter is unknown but invariant with time, i.e., α, with the expectation that a variable parameter estimate, i.e., $\hat{\alpha}(t_k)$, will deny that prior assumption; or

(ii) It can be assumed that the parameter is variable, but varying in an unknown (random-walk) fashion, i.e., $\alpha(t_k)$, with the expectation of the posterior result that a more useful model of the parameter variations can be postulated through interpretation of $\hat{\alpha}(t_k)$.

Access to these tests is gained via the use of an assumed intensity of random perturbation of the parameter dynamics, by analogy with the state vector dynamics. Thus, distinguishing now between Q_{xx}, as the variance-covariance matrix of the unknown state perturbations, and $Q_{\alpha\alpha}$, as the variance-covariance matrix of a corresponding set of unknown parameter perturbations, the specification of the latter can be equated with a quantitative characterisation of the degree of confidence attaching to each constituent model hypothesis.

Moreover, setting $Q_{\alpha\alpha} = 0$ is the most dramatic way of formulating the test of (i) above; it ought to have the greatest possibility of exposing unambiguously the failure of a constituent model hypothesis. The ability of a physical engineering structure to resist deformation when placed under a test load is, by analogy, dependent upon the mechanical properties of its structural members; and these can be likened to the degree of confidence attached to each constituent hypothesis in an abstract model. The more confidently, or the more boldly the hypotheses are assumed to be stated, so the model structure is less flexible, more rigid, more brittle, and the more demonstrative should be the failure of the test structure.

Quite the opposite, however, is needed for the test of (ii) above. In this, speculation about a possibly improved specification of the model structure is the objective. The test draws its strength from the inherent flexibility of the model structure, which can be easily moulded to the patterns in the data and which can, therefore, be suggestive of ways in which to modify hypotheses.

All this, of course, is fine in principle, but not in practice.

2.2 Towards the limits

It is "fine in principle" because of the philosophical underpinnings that can be attached to what is in effect a Popperian programme of falsifying boldly stated, constituent hypotheses. And such association has uniquely been enabled as a result of using filtering theory as the conceptual framework for grappling with the problems of identification. In terms of understanding uncertain environmental systems and the evolution of knowledge MacFarlane (1990) equally so establishes a strong association between the work of Lewis and Popper:

> Popper's and Lewis's approaches to knowledge are essentially the same, with different emphases. Both split concepts from interpretations, and both emphasise the distinct role of the individual mind in generating and using knowledge. Both regard the acquisition of

knowledge as an iterative feedback process. Popper emphasises the objectivity of concepts, and Lewis emphasises the pragmatic role of the individual mind.

It is poor in practice because, first, too many prior assumptions - on $\{P_{xx}(t_0), P_{\alpha\alpha}(t_0), S(t_k), Q(t_k), R(t_k)\}$ not to mention the prior estimatexs of $[x,\alpha]$ at t_0 -- must be made in order to implement the filter.

Second, it is equally poor in practice because of the difficulty of absorbing and interpreting the sheer volume of diagnostic information yielded from the filter (i.e., in the temporal variations of $\{P_{xx}(t_k), P_{\alpha\alpha}(t_k)\}$ and the recursive estimates $[\hat{\alpha}(t_k), \hat{\alpha}(t_k)]$ themselves). This was especially true in the more ambitious exercises in model structure identification attempted at the time, principally in respect of higher-order, state-space descriptions of the degradation of organic material and the proliferation of algal populations in the Bedford Ouse River (Beck, 1982; 1983). In anything but the smallest of models it is difficult to determine unambiguously where the constituent model hypotheses can be said to have failed. A lack of variation with time of a parameter estimate can result from two causes: the associated hypothesis is crucial and "correct"; or it is simply redundant (not identifiable), and no relevant information has been transferred from the external description of the system to this particular constituent of its internal description. One might then be able to resolve this ambiguity of interpretation through inspection of the variations with time of the diagonal, and then off-diagonal, elements of $P_{\alpha\alpha}(t_k)$ (although these possibilities were not greatly exploited at the time). But in short, the self-same sophistication of the questions that could be asked through the framework of filtering theory had become the barrier to progress in unscrambling the answers so generated. Precisely the same advantage that had led to insights into what the problems actually were, had become the cause of downfall in progress towards their practical solution.

Third, for a programme of research that had shunned the use of larger-scale, arguably "physics-based", models as vehicles for the interpretation of field data, not least because of the presumed difficulty of discriminating between key and redundant hypotheses, limits to the notion of the alternative "small being beautiful" had become apparent. From where, for example, would one pick a (posterior) hypothesis for replacing a demonstrably failed hypothesis in a prior model of insufficient content? This is no easier to answer than the alternative of identifying and casting out a redundant hypothesis from a prior model with surplus content. Unlike the input-output models of time-series analysis, there are no systematic rules for extension or reduction of the number of terms (hypotheses) in the structure of these conceptual, state-space models.

Fourth, the power of the classical experiments of laboratory science lay presumably in promoting the possibility of "acts which interpret data in terms of concepts" by reducing the "set of concepts" under scrutiny to as small a set as possible and by maximising the scope for acquiring a large volume of the "given data". In principle, the possibility of progress in the identification of a model should be enhanced when the order of $[u,y]$ is very much greater than the order of $[x,\alpha)$ or, more succinctly, $O[u,y] >> O[x,\alpha]$ (where order (O) increases with the number of elements in the respective vectors and, for the external description, is an increasing function of the density of temporal sampling). Achieving the condition of $O[u,y] >> O[x,\alpha]$ can rarely be the case in the analysis of environmental systems. But where it is, as in the paucity (*not* simplicity) of the set of concepts required to describe pollutant transport and dispersion in a river, in combination with the facility of implementing repeatable dye-tracing experiments with high-frequency sampling, progress can be dramatic (Beer and Young, 1983; Young and Wallis, 1986). It may even prompt a shift in paradigm of the set of concepts (Young and Lees, 1993), though notably through the formal identification of input-output models cast resolutely in terms of the external description of the system's behaviour $[u,y]$. Such models do not directly permit reconciliation of the data with the set of concepts, in the sense intended here. They enable a translation from a "raw", $[u,y]$, to a "refined", $[\hat{u}, \hat{y}]$, external description of the system's observed behaviour, where the latter may provoke insights into the possible nature of $[x,\alpha]$ that are different from those apparent from the former (as, for example, in understanding what may lie beneath the now celebrated time-series of carbon dioxide variations in the upper atmosphere; Young *et al.*, 1991).

Such, however, is not the norm, the epitome of which is quite the opposite, i.e., $O[u,y] << O[x,\alpha]$. Above all, it was this that rendered the conceptual framework of filtering theory "poor in practice" or, rather more correctly, poor in *widespread* practice.

We had come to understand better what the problems were, in the process of assisting these "acts which interpret data in terms of concepts"; some principles for their more systematic resolution had emerged; yet even in the analysis of the relatively small-scale systems upon which the development of these principles had been based, it was extremely difficult to make them work well.

3. Uncertainty: the escape to pragmatism

By the mid-1980s the search for a method of resolving the problem of model structure identification, as posed above, had had to be put to one side. So too, though less consciously and with less speed, had the notion that there might be some "true" structure to be revealed 'through systematic reconciliation of the model with the observed patterns of behaviour.

In retrospect, our view of what was possible had been drifting: away from these more philosophical and "absolutist" ideas; towards a more pragmatic and "relativistic" position. This was not without purpose (nor without success) . After all, in the light of the above, it was natural to ask just what -- at bottom -- would be achievable? And still more so was this questioning needed, given that even for hydrological systems with copious volumes of data, it had long been found that it was not possible to recover a uniquely best set of parameter estimates allowing a match between the model and the observations (Ibbitt and O'Donnell, 1971; Johnston and Pilgrim, 1976; Sorooshian and Gupta, 1983). Given a candidate specification of the model structure, $[x,\alpha]$, the imperative was to establish what choice and form of observation $[u,y]$ would make this internal description identifiable? Or, put the other way round, what $[x,\alpha]$ could be recovered unambiguously from a given $[u,y]$? In fact, could we ever recover the "true" values of the parameters? And would it matter if we could not?

3.1 Identifiability: theoretical bounds on the possible

These questions, though seemingly theoretical, have a certain practical significance. In seeking to understand the acidification of surface waters, correct identification of the paths along which water flows from its impact with the ground and subsequent entry into a stream is crucial. These flow paths determine the soils and minerals with which the water has contact, the nature of the chemical interactions experienced and the duration of these interactions. Their significance is reflected in the values assumed by the various parameters in a conceptual state-space model. Hence there is a distinct need to know what form of observational data $[u,y]$ -- on water flows, on which natural chemical "tracers", at what sampling frequency, in what sequence, and with what degree of confidence -- will enable *in theory* the determination of a set of unambiguously estimated values for these parameters.

In other words, by a suitable rearrangement of the sets of "knowns" and "unknowns" in the previous definitions of state-parameter estimation and model structure identification, it is possible to develop a definition of a priori identifiability as follows:

Problem #3: A Priori Identifiability

Given a candidate model structure $[f,h]$ and given assumptions about $\{P_{xx}(t_0), P_{\alpha\alpha}(t_0), Q(t_i)\}$ determine which candidate set of observations $[u,y]$, with what degree of uncertainty $\{S(t_k), R(t_k)\}$, will allow the estimation of $[x,\alpha]$ with an acceptable degree of uncertainty $\{P_{xx}(t_N), P_{\alpha\alpha}(t_N)\}$.

This clearly goes beyond the notion of a priori identifiability in the deterministic sense of Bellman and Åström (1970), Pohjanpalo (1978), or Godfrey *et al.* (1982). Here "a priori" connotes simply the act of establishing what is possible *before* implementation of the progamme of observations in the field, thus giving the problem a title not strictly in conformity with usage elsewhere (as in Walter and Pronzanto, 1990).

In welcome contradistinction to what has been said earlier, Problem #3 is a problem to which the framework of filtering theory allows a successful solution (Beck *et al.*, 1990; Kleissen *et al.*, 1990). This success is bought at some expense, however, principally in that: (i) the analysis must be conducted for specific parameterisations $[\hat{\alpha}]$, and may therefore be but localised; (ii) $[\hat{\alpha}]$ must be prevented (within the filter) from varying with time, if there is to be any clarity in interpreting the results of the analysis (this is accomplished by assuming the hypothetical observations $(u,y]$ available to the filter are identical with those generated from the nominal reference trajectory of the model's solution); and (iii) the identifiability of a parameter must be equated with expansion and contraction over time of the elements of the estimation error variance-covariance matrix $P_{\alpha\alpha}(t_k)$.

The test of Problem #3 seeks not to determine whether convergent estimates of the model's parameters can as such be obtained but is instead designed to explore the way in which uncertainty is propagated through the model and, in particular, how information (as the reciprocal of uncertainty) is transferred from the external description of the system $[u,y]$ to the component parts of its internal description $[x,\alpha]$. In contrast to the problems of model calibration and model structure identification, whose focus is on the nature of the estimates of $[x,\alpha]$, the focus in this test is on the properties of $P_{\alpha\alpha}(t_k)$. Whether, and to what extent, the elements of this matrix contract or expand will be indicative of whether the constituent model parameters are more or less identifiable. The information-transcribing mechanism, uniquely associated with the filtering algorithm, is itself revealing of how access to which observations, and which combinations of field conditions and naturally perturbing events, can enhance (or corrupt) the confidence attaching to the various parts of the model.

Yet such analysis, in the context of needing to understand the behaviour of uncertain environmental systems, is in the end only shadow-boxing. It does not facilitate acts which interpret data in terms of concepts. Rather, since it deals with hypothetical data from contemplated experiments, it merely enables bounds to be set on what is possible for such reconciliation in the best of all theoretical worlds. This is not, however, without practical

implications, as, for example, in the design of better pumping tests for the identification of aquifer parameters in the characterisation of groundwater fields.

In the real world, as well as in this theoretical world, we now know that the bounds on what is possible are close by. At most just a handful of parameters α can be recovered unambiguously from a single input-output pair, monitored with a high frequency (Hornberger *et al.* 1985; Beven, 1989; Jakeman *et al.*, 1990; Kleissen *et al.*, 1990). This has the ring of relativism about it, in the sense that we may expect it to hold irrespective of the scale of resolution of the observations. Yet it may not be a linear property, in the sense that two input-output pairs will probably not permit as many as two handfuls of model parameters to be recovered. And there is evidence of a theoretical nature suggesting that the presence of nonlinearities in the candidate model will degrade still further the "degree" of identifiability of the parameters (Kleissen *et al.*, 1990).

A number -- and it is notably small -- has been placed on what is achievable.

3.2 Sparseness of the data: progress under substantial uncertainty

We might draw at least three conclusions from this experience, that:

(i) most, if not virtually all, of the conceptual state-space models used to describe the behaviour of environmental systems are not identifiable and are therefore incapable of unambiguous reconciliation with the given data (which is not to say that they are not useful models);

(ii) the entire concept of identifiability is, as a consequence of (i), simply not useful (or that our quantitative analysis of it is defective); or

(iii) from the pragmatic point of view of making predictions, this inability to eliminate ambiguity is not material (what matters is that the ambiguity is apparent and its implications quantifiable).

This last is an area in which progress has been made, although it might in fact be seen as something of a retreat: from seeking the ideal of acquiring a uniquely *optimal* set of parameter estimates; through a search for a cluster of merely *relatively good* candidate parameterisations of the model; to being content with just an *acceptable* set of such candidate parameterisations.

This retreat was not actually a matter of conscious pursuit. Its origins were born of other motivations. One of these was the recognition (in the late 1970s) that $O[u,y]$ may be so small as to amount to little more than merely a subjective, expert, qualitative appreciation, ["*u*","*y*"] say, of the external description of the system's behaviour. So common was (and still is) this the case -- of "rich" sets of concepts confronted with "impoverished" sets of sparse data -- that the absence until then of systematic methods enabling some kind of progress to be made under these conditions, is a curiosity (Hornberger and Spear, 1981). Having thus shifted from a quantitative to a qualitative description of observed behaviour, expectations of what is possible from the acts which interpret the field data in terms of a set of concepts must similarly so be drawn back. One can at most investigate which, among the many constituent mechanisms (hypotheses) in this rich set of concepts, are key (and which redundant) to discriminating between the model's capacity to generate what is defined as acceptable behaviour and its complement (not-the-behaviour). Crucially, it would then be to a better understanding of the nature of these key mechanisms so identified that the inevitably limited capacity for implementing field experiments should gainfully be directed (Hornberger and Spear, 1980; Spear and Hornberger, 1980).

In the face of gross uncertainty the distinction sought in this analysis, between key and redundant mechanisms, must necessarily be based on a sufficiently large sample of randomly generated candidate model parameterisations. Each parameterisation classified as acceptable constitutes an equally probable interpretation of past behaviour so that, with a twist to the original intention of the analysis, the ensemble of such acceptable parameterisations may be used for computing an ensemble of predictions (Fedra *et al.*, 1981). Any ambiguity, distortion, or uncertainty residing in the model following this justifiably passing attempt at its reconciliation with the data, is thereby apparent and quantitatively reflected in its predictions of future behaviour.

3.3 Identifiability, predictability and pragmatism

We can thus escape to the pragmatism of making predictions; and there have been many subsequent examples of precisely this (Keesman and van Straten, 1991; Klepper *et al.*, 1991; Beck and Halfon, 1991; and Beven and Binley, 1992). Again, in the case of surface water acidification, all the attraction of so doing is readily evident. For a situation of $O[u,y] >> O[x,\alpha]$ even the best attempts at identification lead to ambiguous, if not contradictory, interpretations of past behaviour: $[x,\alpha]^1$ and $[x,\alpha]^2$. In effect, it is found that the water has contact either strictly with an upper soil horizon alone or strictly with both an upper and lower soil horizons (Beck *et al.*, 1990). Either could be used for predictive purposes.

The interesting question -- called herein a question of "predictability" (albeit not in the terms discussed by Wegman (1989), for example) -- is whether this makes any material difference. More formally, and once more rearranging the sets of "knowns" and "unknowns" in the preceding problem definitions, we have:

Problem #4: Predictability

Given $[f,h;u]$, together with assumptions about $\{Q(t_k), S(t_k), R(t_k)\}$, and two (or more) interpretations of past behaviour crystallised through $[x,\alpha]^{1,2}$ and $\{P_{xx}(t_N), P_{\alpha\alpha}(t_N)\}^{1,2}$, determine whether $[y]$ differs significantly from $[y]^2$ in the light of $\{T(t_k)\}^1$ and $\{T(t_k)\}^2$.

Here $[y]^{1,2}$ are understood to be output responses of interest, upon which decisions may be based, and which are generated from $[x,\alpha]^{1,2}$ respectively; $\{T(t_k)\}^{1,2}$ are the respective error variance-covariance matrices associated with $[y]^{1,2}$. It should also be noted that t_k now refers to discrete instants in time over a forecasting horizon starting at t_N, for which a future pattern of input disturbances u is assumed to be known, albeit with a degree of uncertainty (as expressed by $S(t_k)$).

It will come as no surprise that this problem can be approximately solved within the context of filtering theory, simply by assuming that the next sampling instant for observing the system's behaviour is infinitely far into the future (Beck, 1983; Beck and Halfon, 1991).

4. Contemporary scene

If there is a "true" structure believed to govern the behaviour of the system -- which is essentially the assumption that motivated the developments of the 1970s (Section 2 above) -- our goal may not necessarily be to have this "truth" revealed by systematic and protracted attempts at reconciliation of the model with the data. Our outlook on the way in which we seek to understand the world around us and, perhaps more pragmatically, to utilise this understanding, has changed. Progress throughout the 1980s (as set out in Section 3) has brought to us to a position in which we recognise that each successive attempt at reconciling the model with the data will result in a distorted interpretation of what is observed; this distortion can be quantified (not least within the conceptual framework of filtering theory); its effects can be faithfully reflected in the predictions generated from the distorted model structure; and the task is then to choose between alternative courses of policy and regulatory action in a manner that is maximally insensitive to this ambiguity and distortion. In a sense too, the field has drawn back from the search for optimality in the performance of its models.

4.1 Behaviour under novel conditions: reachable futures

There is a paradox. The greater the degree of extrapolation from past conditions, so the greater must be the reliance on a model as the instrument of prediction; hence, the greater the desirability of being able to quantify the reliability of the model, yet the greater is the degree of difficulty in doing just this. What is more, the contemporary problems of environmental protection are increasingly of such a form where prediction of behaviour under quite novel conditions is called for. For example, there is a need to predict the fate of utterly novel chemicals in the environment before they are released into that environment. Or alternatively, in the case of surface water acidification, there is a need to extrapolate from small-scale catchments observed over a matter of years to behaviour over entire continental regions and over decades into the future.

We are strongly accustomed to the idea of behaviour being specified in terms of time-series of observations of the system's inputs and outputs $[u,y]$. This is not only rarely the case, as already acknowledged in the work of Hornberger, Spear and Young (Young *et al.*, 1978; Hornberger and Spear, 1980; and Spear and Hornberger, 1980), it is also very restrictive. The notion that $[u,y]$ can be replaced by more qualitative, linguistic descriptors ["u" , "y"] may come to have a profoundly liberating influence on the subject of understanding uncertain environmental systems.

The analyst has immense freedom to be creative in defining the task or purpose, of a model. Fitting the historical data as closely possible has been a traditional such purpose, although this was not an end in itself, merely a means to a better understanding of the system's behaviour. In their seminal work Young *et al.* (1978) were concerned to locate a sample of randomly generated values for the model's parameters that enabled the model outputs to satisfy certain crude constraints, i.e., ["u", "y"], on what is defined (not actually observed) to be an acceptable statement of past behaviour. Yet if behaviour can be so defined for the past, so too can it be for the future (Beck, 1987, 1991), such that the task shifts to that of locating a sample of randomly generated values for the model's parameters that enable the model outputs to match certain crude constraints on what has been defined to be radically different behaviour of the system in the future.

The questions of interest become ones of whether and how a prespecified pattern of future behaviour is, as it were, "reachable" (Beck, 1991). They can formally be defined thus:

Problem #5: Reachable Futures

Given $\{f,h;``u"\}$ and $\{``y"\}$, a prespecified pattern of *future* output responses (that may be radically different from those of the past), determine from $\{``x", ``\alpha"\}$ which constituent parameters $\{\alpha\}^K$ are key to enabling the model to generate $\{``y"\}$ and which $\{\alpha\}^R$ are redundant.

Here $\{``u", ``y"\}$ has been assumed, loosely speaking, to subsume the previous use of $\{u,y\}$ and $\{S(t_k),R(t_k)\}$ and likewise $\{``x", ``\alpha"\}$ subsumes $\{x,\alpha\}$ and $\{P_{xx}(t_k), P_{\alpha\alpha}(t_k), Q(t_k)\}$. Though this may seem a novel problem at first sight, it is not. It is merely a modest rearrangement of the more familiar control problem of finding what input, regulatory action $\{u\}$ (as opposed to $\{\alpha\}$) will transfer $\{y\}$ to some desired performance level $\{y^d\}$. In practice, answers to such questions are of current interest in determining, for example, what aspects of a lake's biochemistry, possibly in combination with which changes in the lake's local climate, will lead to an expressly feared radical change of behaviour in the future.

But to turn matters entirely on their head, if such future responses $\{``y"\}$ are truly radically different from those of the past, then in theory the values of $\{\alpha\}^K$ thus identified ought strictly not to be identifiable (or at most barely identifiable) from the observed record of past behaviour. Now this, if a sensible assertion, would make a virtue out of a lack of model identifiability! Indeed, it provides pointers for where to search -- within the low amplitude, perhaps relatively low-frequency noise of the short records of past behaviour -- for the barely identifiable seeds (constituent mechanisms) of the radically different behaviour that is feared in the future. We have been preoccupied with identifying the major, dominant modes of behaviour captured unambiguously in the "signal". Yet it may well be the minor modes of behaviour, buried within the "noise" at the fringes of our understanding, that have the capacity of becoming the dominant modes of future behaviour.

4.2 Filtering theory: a renaissance

In the end, then, there is a distinct impression of things turning back on themselves, perhaps nowhere more so than in the revival of interest in filtering theory as a means of solving the problem of model structure identification. For we also need pointers, as observed in Section 2, for where to search for a (posterior) hypothesis for replacing a demonstrably failed hypothesis in a prior model of insufficient content.

In 1979 Ljung published a paper on a modified form of extended Kalman filter (let us say LEKF for short) that would improve the estimation of parameters in a conceptual, state-space model (Ljung, 1979). The attraction and potential of this algorithm for the purposes of identifying the model's structure were immediately obvious (Beck, 1987). It held out the possibility, among other advantages, of changing significantly the composition of two of the most important practical difficulties of solving this problem, because (with reference back to Section 2):

(i) It reduced the number of arbitrary prior assumptions that had been necessary to implement the EKF, specifically from the plethora of $\{P_{xx}(t_0), P_{\alpha\alpha}(t_0), S(t_k), Q(t_k), R(t_k)\}$ to just $\{P_{\alpha\alpha}(t_0), S(t_k), R(t_k)\}$;

(ii) It provided access to estimates not only of $[x(t_k), \alpha(t_k)]$ but also the elements of the gain matrix $K(t_k)$, the focal point of the feedback mechanism whereby the filter modulates the account taken of the mismatch between theory and observation according to the balance of the uncertainties attaching to the model and the data.

The balance between "prior assumptions" and "posterior performance", as gauged not just by $[\hat{x}(t_k), \hat{\alpha}(t_k)]$, nor merely with $\{P_{\alpha\alpha}(t_k)\}$ in addition, but now also with $K(t_k)$, has been markedly shifted towards the latter. Indeed, the very way in which the set of concepts are reconciled with the given data -- through use of the filter's gain matrix -- has been made more sensitive to the confrontation of the one "element of knowledge" with the other.

The potential of Ljung's algorithm seems to have remained largely unrealised since its publication, and undoubtedly so in the field of environmental systems analysis, at least until very recently, that is (Stigter, 1993). In fact, liberated thus from the confines of a conventional view of filtering theory, all manner of interesting questions may be opened up. For example:

(i) Does the LEKF have a useful directional property, in the sense of revealing through its "posterior performance" (rather than by "prior assumption") to which constituent hypotheses in the model the failure to match the given observations is due?

(ii) Should the gain matrix of a filter be chosen not -- as is conventionally the case -- in order to minimise the variance of the state-parameter estimation errors but rather to maximise sensitivity to the detection of a structural error or the "seeds" of a structural change?

(iii) Or indeed, should the gain matrix be specified through a neural net trained to detect structural anomalies?

These, however, are already in the realm of speculation for the distant future.

5. In conclusion

Progress can seem painfully slow at times. For many of the ideas reviewed towards the end of this paper are not new. It is really rather disquieting, for example, to have to draw attention to the gap of fourteen years between the publication of Ljung's algorithm and its subsequent successful implementation on a problem of interest to this paper. What is rather different from previous such reflections is the organisation of this review around Lewis's three-element characterisation of knowledge: (i) the given data; (ii) a set of concepts; and (iii) acts which interpret data in terms of concepts. It is this last that has been of principal interest, albeit within the context of making predictions of future behaviour, and inevitably under uncertainty.

At the heart of filtering theory -- as admittedly with any algorithm of system identification -- is a feedback mechanism of correction, adaptation and learning. Drawing a parallel with Lewis's third element of knowledge, what the filter achieves so elegantly is modulation (through its gain matrix) of the account taken of the mismatch between the data and the set of concepts according to the balance of whether the data are believed to be less uncertain than the set of concepts, or vice-versa. Moreover, in contrast to other algorithms, estimated values of the model's parameters can be generated as a function of time. These estimates are accompanied by estimates of the variance of the parameter estimation errors, which are also available as a function of time.

The conceptual framework of filtering theory facilitates insights of a general nature into the mechanics of reconciling the set of concepts with the given data. Fluctuations over time of the reconstructed model parameter estimates are indicative of the failure (or success) of the constituent hypotheses of the model. By analogy with a physical engineering structure, these fluctuations are conditioned upon the strength of the hypotheses and their ability to withstand the various loads imposed on the structure (as a consequence of the mismatch between the model and the data). Such an insight has a very strong association with Popper's view on the acquisition of knowledge. The fundamental difficulties with this, however, are that: (i) the principles of model structure identification so derived can barely be made to work successfully for even the simplest of models; (ii) much of any distortion of the model's structure may be governed by a plethora of notoriously arbitrary prior assumptions; (iii) the evidence gathered from the test is voluminous and not easily distilled into the essence of understanding how the model's structure might be improved; and (iv) the algorithms of filtering theory are being used for a purpose for which they were never intended.

Given this impasse it was easier to make progress on other fronts.

Thus first, decisions are made on the basis of predictions. All predictions are subject to uncertainty; this uncertainty derives in part from the residual uncertainty in the model; and the pattern of residual uncertainty in the model is a function of whatever distortion, or ambiguity, remains as a result of all the preceding attempts to recondile the set of concepts with the data. What matters is whether or not the same decision should be made in the light of these distortions and ambiguities. This is a line of enquiry we can successfully pursue, using, if nothing else, the ubiquitous Monte Carlo simulation. Pragmatically, it may not matter that the distortions are incapable of rectification and the ambiguities incapable of resolution. It is crucially important that they can be quantified, however, through measures of model uncertainty, and their consequences accounted for in the propagation of prediction errors.

Second, and entirely theoretically, if there were a need to resolve the ambiguities in understanding the past, what form of field observations with what accuracy would be needed? This too is a question for which there is a means of obtaining an answer. Contraction and expansion of the uncertainty attaching to the model's parameter estimates, as uniquely illuminated through the estimated variance-covariance matrices of a filtering algorithm, determine when information in the external description of the system (the data) can be productively and counter-productively brought to bear on the constituents of its internal description (the model).

Third, the obligation of forecasting the future under conditions substantially different from those of the past will be the downfall of most, if not all, models. There is a dilemma (Beck, 1983). The "large" model -- that will result from including everything of conceivable relevance to the problem at hand -- may indeed be capable of predicting "correctly" such radically different behaviour; but we would place little confidence in this prediction. The "small" model -- that will result from any honest analysis of the past data -- may quite "incorrectly" predict behaviour in the future essentially similar to that of the past; worse still, we might place great confidence in this erroneous prediction. There is, of

course, no way in which we could have fore-knowledge of either of these results. Instead, more fruitful progress can be made in identifying which constituent mechanisms in the model, and/or what degree of climate change, for instance, may be key to the reaching (or not) of some predefined radically different "target" behaviour of the given system in the future.

Last, and turning to matters of contemporary interest, progress may now be possible in two of the principal areas where the original insights of filtering theory failed nevertheless to enable the development of practical solutions to the problems of interpreting the given data in terms of a set of concepts. Certain modifications of the basic algorithms allow us to dispense with some of the arbitrary prior assumptions and to have access to a greater variety of fee ' dback, diagnostic information. It remains to be seen whether there are ways of manipulating all of this information for the purposes of accelerating a reconciliation of the model with the data.

Acknowledgements

The work reported on in this paper has been supported in part by the Visiting Scientists and Engineers Program of the United States Environmental Protection Agency (project title "The Analysis of Uncertainty in Environmental Simulation"). The paper is also a contribution to the International Task Force on Forecasting Environmental Change, which is currently receiving support from the National Water Research Institute of Environment Canada, the National Institute of Public Health and Environmental Protection (RIVM) of the Netherlands, and the International Institute for Applied Systems Analysis, Laxenburg, Austria.

The author is currently Visiting Professor in the Department of Civil Engineering at Imperial College, London.

References

Amato, I. (1991), 'Speculating in precious computroniuml', *Science*, **253**, 23 August, 856-857.

Beck, M. B. (1982), 'Identifying models of environmental systems behaviour', *Mathematical Modeling*, **3**, 467-480.

Beck, M.B. (1983), 'Uncertainty, system identification, and the prediction of water quality', in *Uncertainty and Forecasting of Water Quality* (M.B. Beck and G. van Straten, eds), Springer, Berlin, pp 3-68.

Beck, M.B. (1985), 'Lake eutrophication: identification of tributary nutrient loading and sediment resuspension dynamics', *Appl Math Comput*, **17**, 433-458.

Beck, M.B. (1986), 'The selection of structure in models of environmental systems', *The Statistician*, **35**, 151-161.

Beck, M.B. (1987), 'Water quality modeling: a review of the analysis of uncertainty', *Wat Res Res*, **23(8)**, 1393-1442.

Beck, M.B. (1991), 'Forecasting environmental change', *J Forecasting*, 10(1&2), 3-19.

Beck, M.B. and E. Halfon (1991), 'Uncertainty, identifiability and the propagation of prediction errors: a case study of Lake Ontario', *J Forecasting*, 10(1&2), 135-161.

Beck, M.B. and G. van Straten (eds) (1983), '*Uncertainty and Forecasting of Water Quality*, Springer, Berlin.

Beck, M.B. and P.C. Young (1976), 'Systematic identification of DO-BOD model structure", *Proc Am Soc Civil Eng, J Environ Eng Div*, **102(EE5)**, 902-927.

Beck, M.B., F.M. Kleissen and H.S. Wheater (1990), 'Identifying flow paths in models of surface water acidification', *Rev Geophys*, **28(2)**, 207-230.

Beer, T. and P.C. Young (1983), 'Longitudinal dispersion in natural streams" *J Environ Eng*, **109(5)**, 1049-1067.

Bellman, R. and K-J. Åström (1970), 'On structural identifiability', *Math Biosci*, **7**, 329-339.

Beven, K.J. (1989), 'Changing ideas in hydrology - The case of physically-based models', *J Hydrol*, **105**, 157-172.

Beven, K.J. and A.M. Binley (1992), 'The future of distributed models: model calibration and uncertainty prediction', *Hydrological Processes*, **6(3)**, 279-298.

Camara, A.S., F.C. Ferreira, D.P. Loucks and M.J.F. Seixas (1990), 'Multidimensional simulation applied to water resources management', *Wat Res Res*, **26(9)**, 1877-1886.

Camara, A.S., M. Pinheiro, M.P. Antunes and M.J.F. Seixas (1987), 'A new method for qualitative simulation of water resources systems. 1. Theory', *Wat Res Res*, **23(11)**, 2015-2018.

Castro, A.P., D.L. Gallagher and A.S. Câmara (1993), 'Dynamic water quality modeling using cellular automata', *Technical Report*, Virginia Polytechnic Institute, Blacksburg, Virginia.

Fedra, K., G. van Straten and M.B. Beck (1981), 'Uncertainty and arbitrariness in ecosystems modelling: a lake modelling example', *Ecol Modelling*, **13**, 87-110.

Folland, C., J. Owen, M.N. Ward and A. Colman (1991), 'Prediction of seasonal rainfall in the Sahel Region using empirical and dynamical methods', *J Forecasting*, **10(1&2)**,21-56.

Godfrey, K.R., R.P. Jones, R.F. Brown and J.P. Norton (1982), 'Factors affecting the identifiability of compartmental models', *Automatics*, **18(3)**, 285-293.

Hornberger, G.M. and R.C. Spear (1980), 'Eutrophication in Peel Inlet, I. Problem-defining behaviour and a mathematical model for the phosphorus scenario', *Wat Res*, **14**, 29-42.

Hornberger, G.M. and R.C. Spear (1981), 'An approach to the preliminary analysis of environmental systems', *J Env Management*, **12(1)**, 7-18.

Hornberger, G.M., K.J. Beven, B.J. Cosby and D.E. Sappington (1985), 'Shenandoah Watershed study: calibration of a topographybased, variable contributing area hydrological model to a small forested catchment', *Wat Res Res*, **21(12)**, 1841-1850.

Ibbitt, R.P. and T. O'Donnell (1971), 'Fitting methods for conceptual catchment models', *Proc Am Soc Civil Eng, J Hydr Div*, **97(HY9)**, 1331-1342.

Jakeman, A.J., I.G. Littlewood and P.G. Whitehead (1990), 'Computation of the instantaneous unit hydrograph and identifiable component flows with application to two small upland catchments', *J Hydrol.*, **117**, 275-300.

Johnston, P.R. and D.H. Pilgrim (1976), 'Parameter optimization for watershed models', *Wat Res Res*, **12(3)**, 477-486.

Kalman, R.E. (1960), 'A new approach to linear filtering and prediction problems', *Am Soc Mech Eng Trans, J Basic Eng*, **82**, 3545.

Kalman, R.E. and R.S. Bucy (1961), 'New results in linear filtering and prediction theory', *Am Soc Mech Eng Trans, J Basic Eng*, **83**, 95-108.

Kleissen, F.M., M.B. Beck and H.S. Wheater (1990), 'The identifiability of conceptual hydrochemical models', *Wat Res Res*, **26(12)**, 2979-2992.

Klepper, O., H. Scholten and J.P.G. van de Kamer (1991), 'Prediction uncertainty in an ecological model of the oosterschelde Estuary', *J Forecasting*, **10(1&2)**, 191-209.

Ljung, L. (1979), 'Asymptotic behaviour of the extended Kalman filter as a parameter estimator', *IEEE Trans Automat Contr*, **24**, 36-50.

MacFarlane, A.G.J. (1990), 'Interactive computing: a revolutionary medium for teaching and design', *Computing & Control Engineering Journal*, **1(4)**., 149-158.

Pohjanpalo, H. (1978), 'System identifiability based on the power series of the solution', *Math Biosci*, **41**, 21-33.

Pool, R. (1989), 'Is it real, or is it a Cray?', *Science*, **244**, 23 June, 1438-1440.

Sorooshian, S. and V.K. Gupta (1985), 'The analysis of structural identifiability: theory and application to conceptual rainfallrunoff models', *Wat Res Res*, **21(4)**, 478-495.

Spear, R.C. and G.M. Hornberger (1980), 'Eutrophication in Peel Inlet, II. Identification of critical uncertainties via generalised sensitivity analysis', *Wat Res*, **14**, 43-49.

Stigter, J.D. (1993), 'Application and further development of the extended Kalman filter in a river water pollution problem', *Graduation dissertation*, Department of Applied Mathematics, University of Twente, Enschede, The Netherlands.

van Straten, G. and K.J. Keesman (1991), 'Uncertainty propagation and speculation in projective forecasts of environmental change: a lake-eutrophication example', *J Forecasting*, **10(1&2)**, 163-190.

Walter, E. and L. Pronzanto (1990), 'Qualitative and quantitative experiment design for phenomenological models - A survey', *Automatics*, **26**, 195-213.

Wegman, E.J. (1988), 'On randomness, determinism and computability,, *J Statistical Planning and Inference*, **20**, 279-294.

Young, P.C., and M. Lees (1993), 'The active mixing volume: a new concept in modelling environmental systems', in *Statistics and the Environment* (V Barnett and R Turkman, eds) , Wiley, Chichester (to appear).

Young, P.C., and S.G. Wallis (1986), 'The aggregated dead zone (ADZ) model for dispersion in rivers', in *Water Duality Modelling in the Inland Natural Environment*, British Hydraulics Research Association, Cranfield, England, pp 421-433.

Young, P.C., G.M. Hornberger and R.C. Spear (1978), 'Modelling badly defined systems - some further thoughts', in *Proceedings SIMSIG Simulation Conference*, Australian National University, Canberra, pp 24-32.

Young, P.C., C.N. Ng, K. Lane and D. Parker (1991), 'Recursive forecasting, smoothing and seasonal adjustment of non-stationary environmental data', *J Forecasting*, **10(1&2)**, 57-89.