Bedford T, Wilson KJ, Daneshkhah A. **Assessing parameter uncertainty on coupled models using minimum information methods**. *Reliability Engineering and System Safety* 2014, 125, 3-12.

Copyright:

DOI link to article:

http://dx.doi.org/10.1016/j.ress.2013.05.011

Date deposited:

22/04/2016

Assessing parameter uncertainty on coupled models using minimum information methods

Tim Bedford[a], Kevin J Wilson[a,*], Alireza Daneshkhah[b]

[a]*Department of Management Science, University of Strathclyde, Glasgow, G1 1QE, UK*
[b]*Cranfield Water Science Institute, Cranfield University, Vincent Building, Cranfield, Bedford, MK43 0AL*

Abstract

Probabilistic inversion is used to take expert uncertainty assessments about observable model outputs and build from them a distribution on the model parameters that captures the uncertainty expressed by the experts. In this paper we look at ways to use minimum information methods to do this, focussing in particular on the problem of ensuring consistency between expert assessments about differing variables, either as outputs from a single model, or potentially as outputs along a chain of models. The paper shows how such a problem can be structured and then illustrates the method with two examples; one involving failure rates of equipment in series systems and the other atmospheric dispersion and deposition.

Keywords: Minimum information, coupled models, expert judgement, Probabilistic Risk Analysis, Gaussian plume.

1. Introduction

An important element of Probabilistic Risk Analysis is assessment of uncertainty in model outputs. Physical models do not perfectly represent the phenomena they are meant to describe for many reasons: lack of complete understanding of the physical phenomena, deliberate simplifications in the model (for example, because of the need to run the model quickly), inadequate choice of model parameters, and so on.

Bedford and Cooke [1] stress the importance of assessing uncertainties for observable quantities when using probability. Since it is typically model outputs that are observable quantities, and many model parameters are not directly observable (maybe having no direct physical interpretation) it is therefore necessary to consider ways of taking probability distributions that describe the uncertainty in model output quantities and "back-fitting" these to generate a distribution on the model parameters so that this matches the uncertainty specified for the output parameters in the following sense: If we randomly choose a set of model parameters and compute the model outputs then those

outputs are also random. The distribution of the outputs obtained is called the push-forward of the distribution on the model parameters and should match the uncertainty for the observable quantities.

The problem of Probabilistic Inversion (PI) is simply the problem of computing a distribution on the model parameters with the property that its push forward matches that specified for the distribution on the observable quantities. See Bedford and Kraan [2] and Kurowicka and Cooke [3] for different approaches to this problem. Having defined an uncertainty distribution on the the model parameters, this then allows us to make predictions, incorporating our uncertainty, on model outputs for any set of model inputs.

The probabilistic inversion problem is typically either under- or over-specified. This means that the constraints imposed by the output distributions are either not sufficiently strong to lead to a single solution to the PI problem, or they are mutually contradictory, and give rise to an infeasible problem.

Previous approaches have used minimally informative distributions to solve the problem of under-specification, and have either ignored the problem of infeasibility or used slightly ad-hoc approaches to deal with it. This paper develops the ideas proposed in Bedford [4] to use minimum information methods [5, 6] to provide guidance in the specification of constraints that are not infeasible. It also gives a method to map out the feasible region for the constraints. Minimum information methods use constraints on expected values rather than quantile information, though we discuss how quantile information can be incorporated.

The paper generalizes the context given above to one of "coupled" models, that is, the situation where we have several models which are sequentially linked in the sense that the outputs of the first are the inputs of the second, and so on. Hence specifications may be made of distributions of input and output parameters, and the distributions on the model parameters are supposed to push forward to these.

The main contribution of this paper is therefore to show how the minimum information approach can be used to generate solutions to the PI problem in the context of coupled models. The rest of the paper is structured as follows. In Section 2 we provide the mathematical setup to the problem, considering coupled models and minimum information modelling. In Section 3 we outline the solution to the continuous PI problem for coupled models and in Section 4 we give the equivalent solution to the discretized problem. Section 5 investigates the specification of feasible constraints for the coupled PI problem and considers the effect of fixing the constraints. In Sections 6 and 7 we provide two examples; the first considering the failure rates of machines located in series and the second based on the dispersion and deposition problem considered by Kraan in [7]. Finally, in Section 8, we give some conclusions and a discussion of areas for further work.

2. Mathematical set up

2.1. Coupled models

We consider models M_i taking input parameters y_i and giving output parameters y_{i+1}. So the output of model 1 is the input for model 2. The models are deterministic models, but are assumed to have parameters (which may be vectors) that we denote x_i. An
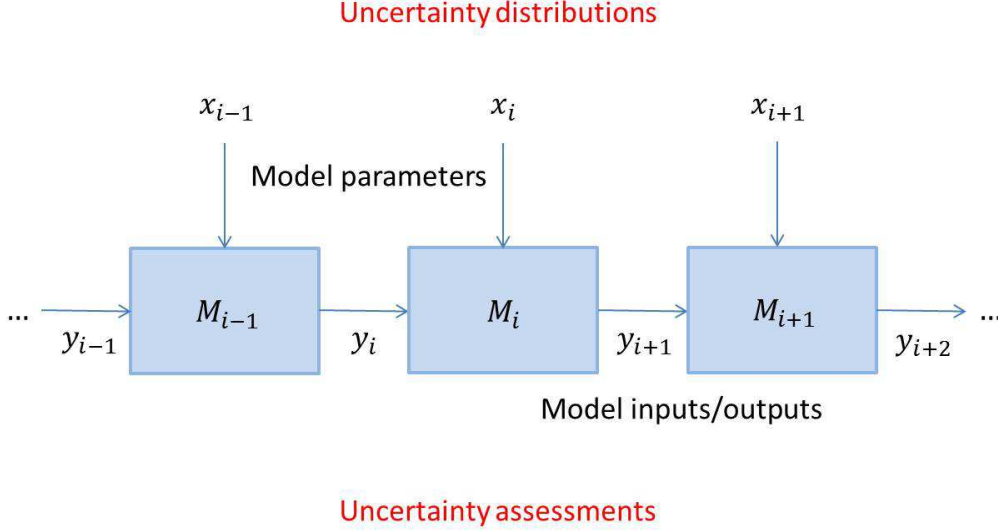
Figure 1: A diagram illustrating the form of coupling of models.

illustration of this situation is given in Figure 1.

Some of the parameters may be directly measurable. Some we might be able to choose. However, often the model parameters are not known directly and we have to infer them, or infer uncertainty distributions over them. The inputs and outputs of the models are observable quantities in principle, and we may therefore be able to use expert judgement to assess distributions or expected values. In probabilistic inversion we try to find a distribution for the parameters that matches (when propagated through the model) the distributions specified by the experts for the observable quantities.

As noted above, the PI problem is typically either over- or under-constrained. Over-constraint happens because the quantities to be assessed are typically functionally related through the model, so it easy for experts to provide information on observables that is mutually inconsistent (assuming the model is correct). Under-constraint happens because there is consistency but there is still not enough information to uniquely determine the distribution on the parameters. In our approach we use a minimum information property to solve the problem of under-constraint, and a sequential approach to expert elicitation to avoid the problem of over-specification.

2.2. Minimum Information Modelling

Suppose we have continuous random variable X, which could be multi-dimensional, and two densities, $f_1(\cdot)$ and $f_2(\cdot)$. Then the relative information of $f_1(\cdot)$ to $f_2(\cdot)$ is a measure of how similar the two distributions are. It is defined as

$$I(f_1 \mid f_2) = \int f_1(x) \log\left(\frac{f_1(x)}{f_2(x)}\right) dx.$$

If $f_1(\cdot) = f_2(\cdot)$ then the relative information is zero. It then increases as the deviation between the two distributions becomes greater.

In the problems considered in this paper we wish to find the distribution, $f_1(\cdot)$, which has minimum information, with respect to the background distribution $f_2(\cdot)$, subject to some real valued functions $h_1 \ldots, h_k$ taking expectations $\alpha_1, \ldots, \alpha_k$. These are known as the *constraints*. The choice of background distribution is an important consideration as this specifies what $f_1(\cdot)$ should look like in the absence of more specific information. Typical choices might be the uniform, log-uniform and Normal distributions (but this can be subject to sensitivity analysis as we discuss later).

The minimum information distribution is therefore, in some sense, the "simplest" distribution which satisfies the required criteria. If, relative to the background distribution, this distribution exists then it is unique and takes the form, Lanford [8],

$$f(x) = \frac{\exp\left\{\sum_{i=1}^{k} \lambda_i h_i(x)\right\}}{Z(\lambda)}, \tag{1}$$

for Lagrange multipliers $\lambda = (\lambda_1, \ldots, \lambda_k)$ (depending on $\alpha_1, \ldots, \alpha_k$), where $Z(\lambda)$ is the normalising constant. An important contribution of this paper is a procedure which provides a method of specifying the constraints using expert elicitation which ensures such a unique minimally informative solution exists. We can approximate this density arbitrarily well using discrete densities

$$p(x_j) \propto \exp\left\{\sum_{i=1}^{k} \lambda_i h_i(x_j)\right\}, \tag{2}$$

where x_j, $j = 1, \ldots, n$ is a suitable discretization of x. A derivation of these results is outlined in Bedford and Wilson [9].

We see that, for minimum information problems, constraints are naturally specified as expectations on functions of the parameters. We can incorporate quantile assessments as the expert judgement by specifying the h_i's as indicator functions in terms of the quantile functions of the parameters. An example is given in Bedford *et al* [10] and the technique is used in the failure rates example in Section 6.

We shall now consider how to evaluate the feasible values of constraints when the specifications of those constraints are made sequentially. Initially we consider the continuous problem.

3. The Continuous Problem

Consider a collection of parameters from a deterministic model or sequence of models, denoted X. In general X could have a large number of dimensions. Suppose that expert elicitation results in the specification of the expectations of some functions, $h_1, \ldots, h_p$, of these parameters, denoted $\alpha_1, \ldots, \alpha_p$. These are the constraints in the problem. We wish to find the distribution with minimum information which satisfies these expectations.

As we saw in the previous section this minimum information distribution has a density at x which is proportional to

$$\exp\left\{\sum_{i=1}^{p} \lambda_i h_i(x)\right\},$$

for some parameters $\lambda_1 \ldots, \lambda_p$. We can define the function $\phi : \mathbb{R}^p \to \mathbb{R}^p$ so that ϕ gives the expected values obtained from choosing the minimum information distribution with parameters $\lambda_1, \ldots, \lambda_p$. That is,

$$\phi(\lambda_1 \ldots, \lambda_p) = (\alpha_1, \ldots, \alpha_p).$$

In particular,

$$\phi(\lambda_1, \ldots, \lambda_p) = \left(\frac{\int h_1(x) \exp\{\sum_i \lambda_i h_i(x)\} dx}{\int \exp\{\sum_i \lambda_i h_i(x)\} dx}, \ldots, \frac{\int h_p(x) \exp\{\sum_i \lambda_i h_i(x)\} dx}{\int \exp\{\sum_i \lambda_i h_i(x)\} dx} \right).$$

The function ϕ is invertible and has good analytical properties.

We wish to specify α_1 and use this to explore the possible values α_2 can take. Having done this the next step is to find the possible values for α_3 having specified α_1, α_2. We can continue in this way, evaluating the possible specifications of each expectation consistent with previously specified values, until all of the required expectations have been set. Doing this stepwise allows us to guide the expert in making allowable choices for the constraints.

To see how this can be achieved, initially consider functions $\psi_1, \psi_2, \ldots, \psi_p$ which are defined, in terms of $\phi = (\phi_1, \ldots, \phi_p)$, as

$$\psi_1(\lambda_1) \quad = \quad \phi_1(\lambda_1, 0, \ldots),$$
$$\psi_2(\lambda_1, \lambda_2) \quad = \quad (\phi_1(\lambda_1, \lambda_2, 0, \ldots), \phi_2(\lambda_1, \lambda_2, 0, \ldots)),$$

and so on. Since each ϕ_i is monotone, $\psi_1, \psi_2, \ldots, \psi_p$ are monotone. Therefore $\psi_1, \psi_2, \ldots, \psi_p$ are invertible and so we can invert ψ_1 to obtain ψ_1^{-1}. This allows us to calculate λ_1 from a given α_1, for example using Newton's method. We can also use this relationship to explore the possible values of α_1.

Having specified $\alpha_1 = \alpha_1^*$ we would like to find the values of λ_1, λ_2 for a particular α_2. That is, we are interested in

$$\psi_2^{-1} \mid_{\alpha_1 = \alpha_1^*} (\alpha_2) = (\lambda_1, \lambda_2). \tag{3}$$

Thus we are mapping $\text{image}(\psi_2) \cap \{\alpha_1 = \alpha_1^*\} \to \mathbb{R}^2$. This is a map of one variable. Initially, having specified $\alpha_1 = \alpha_1^*$, we can identify a value $\tilde{\alpha}_2$ so that

$$\psi_2^{-1} \mid_{\alpha_1 = \alpha_1^*} (\tilde{\alpha}_2) = \left(\lambda_1^{(1)}, 0 \right),$$

by taking $\tilde{\alpha}_2 = \mathrm{E}[h_2(X)]$ under the distribution with parameters $(\lambda_1^{(1)}, 0)$. We can use this as a starting value for exploring the range of possible values for α_2. The range can then be explored using the inverse function in (3). Having found the set of feasible values we specify the value of $\alpha_2 = \alpha_2^*$ in this range and use ψ_2^{-1} as above to find $(\lambda_1^{(2)}, \lambda_2^{(2)})$.

The general step begins with the specification of $\alpha_1 = \alpha_1^*, \ldots, \alpha_{n-1} = \alpha_{n-1}^*$ having been made. The relevant conditional mapping is then

$$\psi_n^{-1} \mid_{\alpha_1 = \alpha_1^*, \ldots, \alpha_{n-1} = \alpha_{n-1}^*} (\tilde{\alpha}_n) = \left(\lambda_1^{(n-1)}, \ldots, \lambda_{n-1}^{(n-1)}, 0 \right).$$

This is used to explore the space of possible values for α_n. Having specified $\alpha_n = \alpha_n^*$ we use it to calculate $(\lambda_1^{(n)}, \ldots, \lambda_n^{(n)})$. We see that $\lambda_i^{(j)} = 0$ for $i > j$.

It is clear from the above formulation that the order the specifications are made is going to have an effect as we are considering expectations one at a time. This is discussed in Section 7.3. To solve such problems in practice we work with the discrete version of the problem. This is set out in the next section.

4. The discretized problem

Suppose we have discretized at points $x_1, \ldots, x_m$. In two dimensions these could be points on the square. As before we have constraints on some functions $h_1, \ldots, h_p$ in the form of expected values $\alpha_1, \ldots, \alpha_p$. The minimum information distribution has mass at point x_j which is proportional to $\exp(\sum_i \lambda_i h_i(x_j))$.

We replace integrals by sums in the definition of ϕ. That is,

$$\phi(\lambda_1, \ldots, \lambda_p) = \left(\frac{\sum_j h_1(x_j) \exp\left\{\sum_i \lambda_i h_i(x_j)\right\}}{\sum_j \exp\left\{\sum_i \lambda_i h_i(x_j)\right\}}, \ldots, \frac{\sum_j h_p(x_j) \exp\left\{\sum_i \lambda_i h_i(x_j)\right\}}{\sum_j \exp\left\{\sum_i \lambda_i h_i(x_j)\right\}} \right).$$

Now, since we have a simple, analytic formula for ϕ, we can determine $D\phi$, the matrix of derivatives with respect to the λ's. This is a $p \times p$ matrix with elements

$$\frac{\partial \phi_k}{\partial \lambda_l} = \frac{\sum_j h_k(x_j) h_l(x_j) A_j}{\sum_j A_j} - \frac{[\sum_j h_k(x_j) A_j] \times [\sum_j h_l(x_j) A_j]}{[\sum_j A_j]^2},$$

where $A_j = \exp\{\sum_i \lambda_i h_i(x_j)\}$. As long as $D\phi$ is of full rank then we can invert it, numerically, to obtain the matrix of derivatives with respect to the α_k's.

This enables us to work out, to first order, using Newton's method for example, how the Lagrange multipliers change when we keep $\alpha_1, \ldots, \alpha_{k-1}$ fixed, and change α_k slightly. When we approach the boundary of the convex set of possible α_k values, the Lagrange multipliers go to plus or minus infinity.

This enables us to work out which parameter vector gives a particular constraint vector. Hence we can systematically map out the ranges of allowed expectations as follows.

1. With the first function h_1 we take $\lambda_1 = \pm N$, where N is some large number, to obtain upper and lower bounds on α_1.

2. The expert chooses α_1. We calculate, numerically by inversion, what the corresponding λ_1 is, denoted l_1.

3. For the next function h_2 we take $\lambda_1 = l_1$ and $\lambda_2 = 0$. This means that our starting value for α_2 is whatever we had for the expectation value with respect to the minimum information distribution we used in Step 1.

4. Using this as a starting value for α_2 we increase α_2, keeping α_1 fixed, and work out the corresponding values of λ_1 and λ_2 using $(D\phi)^{-1}$. When λ_1, λ_2 become large we have reached the largest feasible α_2.

5. We repeat Step 4 decreasing α_2 from its initial value.

6. We repeat Steps 3-5 with the other functions.

We shall illustrate this methodology by applying it to two examples; the first a model of atmospheric dispersion and dispersion and the second a model assessing the failure rates of machines in series. However, first we consider the mapping out of the feasible regions for multiple constraints.

5. Investigation of the feasible region for the constraints

We can investigate the convex set of possible combinations of different constraints known as the feasible region. To do so consider the normalising constant from (1). It is

$$Z(\lambda) = \int \exp\left\{\sum_i \lambda_i h_i(x)\right\} dx,$$

for vector x and parameters $\lambda = (\lambda_1, \ldots, \lambda_p)$. Taking logs and differentiating with respect to Lagrange multiplier λ_i recovers the expectations defined previously,

$$\frac{d}{d\lambda_i} \log Z(\lambda) = \frac{\int h_i(x) \exp\{\sum_i \lambda_i h_i(x)\} dx}{\int \exp\{\sum_i \lambda_i h_i(x)\} dx} = \mathrm{E}_\lambda[h_i(x)].$$

If we differentiate a further time with respect to the same λ_i we obtain the variance of the constraint function associated with this expectation,

$$\begin{aligned}
\frac{d^2}{d\lambda_i^2} \log Z(\lambda) &= \frac{\int h_i^2(x) \exp\{\sum_i \lambda_i h_i(x)\} dx}{\int \exp\{\sum_i \lambda_i h_i(x)\} dx} - \left[\frac{\int h_i(x) \exp\{\sum_i \lambda_i h_i(x)\} dx}{\int \exp\{\sum_i \lambda_i h_i(x)\} dx}\right]^2 \\
&= \mathrm{Var}_\lambda(h_i(x)) \geq 0.
\end{aligned}$$

Let us define the function g to be the vector of the constraint functions, so that $g = (h_1(x), \ldots, h_p(x))$. Then the expectation of g is simply $\mathrm{E}_\lambda[g] = (\mathrm{E}_\lambda[h_1(x)], \ldots, \mathrm{E}_\lambda[h_p(x)])$.

If we make the assumption that g is bounded then $\mathrm{E}_\lambda[g]$ is also bounded. We know that there is a unique λ associated with any possible α in the convex hull of g. We also see that each expectation $\mathrm{E}_\lambda[h_i(x)]$ is increasing as λ increases. This will allow us to map out the feasible region of $\mathrm{E}_\lambda[g]$.

Consider $\lambda_1 \ldots, \lambda_p$ in the sphere of radius 1,

$$\sum_{i=1}^p \lambda_i^2 \leq 1. \tag{4}$$

If we take as our starting point the sphere with unit radius defined by (4) this will map out a region of values, corresponding to this λ, for $\mathrm{E}_\lambda[g]$. If we then multiply the vector $(\lambda_1, \ldots, \lambda_p)$ by some $r > 1$, so that we have $r \times (\lambda_1, \ldots, \lambda_p)$, we are now in the sphere defined by

$$\sum_{i=1}^p (r\lambda_i)^2 \leq r^2.$$

That is, we are moving outward in a radial direction. From this we obtain a new, larger, region for $\mathrm{E}_\lambda[g]$. As r becomes large this region will tend to the overall feasible region for $\mathrm{E}_\lambda[g]$.

We shall illustrate this with a simple example concerned with coin tossing.

Table 1: The possible outcomes, background distribution and contributions to the MI solution for the coin tossing example.

$\{X,Y\}$	0,0	0,1	1,0	1,1
Background dist.	1/4	1/4	1/4	1/4
$e^{\lambda_i h_i(x,y)}$	1	1	e^{λ_1}	$e^{\lambda_1+\lambda_2}$

5.1. Coin tossing example

Suppose we have two coins and that we suspect that each is biased. Let X, Y represent one realisation for each coin of a single toss. That is,

$$X, Y = \begin{cases} 1, & \text{if heads,} \\ 0, & \text{if tails,} \end{cases}$$

with probabilities p_X, p_Y. We also suspect that the tosses are not independent. Further suppose that the functions we are going to ask an expert to specify are $h_1 = x$ and $h_2 = xy$. We know the form of the minimum information distribution in this case. A sensible measure to take as the background distribution would appear to be $(1/4, 1/4, 1/4, 1/4)$ on the four different possible outcomes of the two tosses.

Table 1 sets out each of the outcomes, the associated background distribution and the minimum information terms found analytically. From this we can calculate $Z(\lambda)$. It is

$$Z(\lambda) = \frac{1}{4}(2 + e^{\lambda_1} + e^{\lambda_1+\lambda_2}).$$

If we take logarithms then differentiate we obtain the expectations of each of the functions. They are

$$\mathrm{E}_\lambda[h_1(x,y)] = \frac{d}{d\lambda_1} \log Z(\lambda) = \frac{e^{\lambda_1} + e^{\lambda_1+\lambda_2}}{2 + e^{\lambda_1} + e^{\lambda_1+\lambda_2}},$$

and

$$\mathrm{E}_\lambda[h_2(x,y)] = \frac{d}{d\lambda_2} \log Z(\lambda) = \frac{e^{\lambda_1+\lambda_2}}{2 + e^{\lambda_1} + e^{\lambda_1+\lambda_2}}.$$

We can use these equations to map out the feasible region for the expectation of $g = (h_1(x,y), h_2(x,y))$. We begin with $r = 1$ and then increase it gradually up to $r = 10$. The results are given in Figure 2 for $r = 1, 2, 3, 10$. We see that, as r increases, the region becomes larger until, when $r = 10$, it is covering the entire feasible region for g. Some pseudo-code illustrating how to explore the feasible region in this way is given in Algorithm 1 in the Appendix.

We can also carry out such an investigation in 3 dimensions. Suppose that X, Y are defined as above but now we are interested in the expectations associated with three functions, namely

$$h_1(x,y) = x, \quad h_2(x,y) = y, \quad h_3(x,y) = xy.$$

We can apply the methodology as above. The function $Z(\lambda)$ is now given by

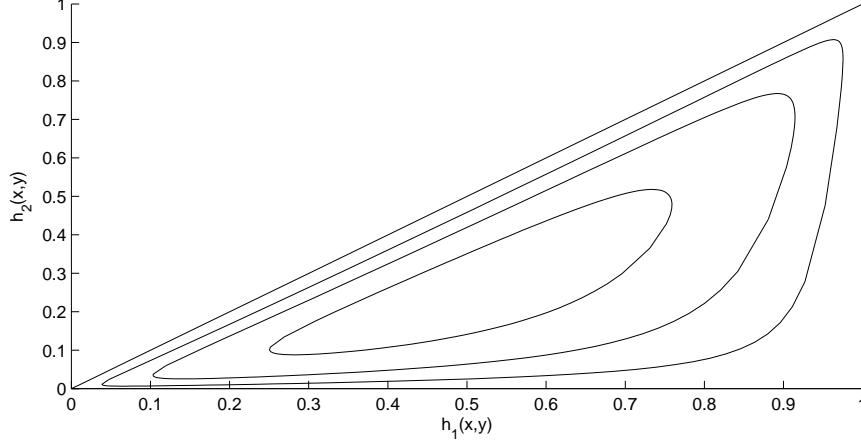$$Z(\lambda) = \frac{1}{4}(1 + e^{\lambda_1} + e^{\lambda_2} + e^{\lambda_1+\lambda_2+\lambda_3}).$$

Figure 2: The feasible region for the expectation of g given $r = 1, 2, 3, 10$.

Taking logarithms and differentiating we can find the expectations of the three functions of interest. They are

$$\mathrm{E}_\lambda[h_1(x, y)] = \frac{e^{\lambda_1} + e^{\lambda_1 + \lambda_2 + \lambda_3}}{1 + e^{\lambda_1} + e^{\lambda_2} + e^{\lambda_1 + \lambda_2 + \lambda_3}},$$

$$\mathrm{E}_\lambda[h_2(x, y)] = \frac{e^{\lambda_2} + e^{\lambda_1 + \lambda_2 + \lambda_3}}{1 + e^{\lambda_1} + e^{\lambda_2} + e^{\lambda_1 + \lambda_2 + \lambda_3}},$$

and

$$\mathrm{E}_\lambda[h_3(x, y)] = \frac{e^{\lambda_1 + \lambda_2 + \lambda_3}}{1 + e^{\lambda_1} + e^{\lambda_2} + e^{\lambda_1 + \lambda_2 + \lambda_3}}.$$

We can then use these three expressions to evaluate the convex region associated with a sphere in $(\lambda_1, \lambda_2, \lambda_3)$ of radius 1 and expand this, moving outwards along the radius, until we are satisfied we are evaluating the entire feasible region for the three constraints. A plot of this, for a radius $r = 10$, is given in Figure 3.

5.2. The discretized case

As, in practice, we usually solve the minimum information problem in a discretized version of the parameter space we now consider such a discretized space in the context of exploring the feasible region of $\mathrm{E}_\lambda[g]$. In this case the quantity $Z(\lambda)$ is

$$Z(\lambda) = \sum_j \exp \left\{ \sum_i \lambda_i h_i(x_j) \right\},$$

and the expectation of $h_i(x)$ is

$$\mathrm{E}_\lambda[h_i(x)] = \frac{d}{d\lambda_i} \log Z(\lambda) = \frac{\sum_j h_i(x_j) \exp \left\{ \sum_i \lambda_i h_i(x_j) \right\}}{\sum_j \exp \left\{ \sum_i \lambda_i h_i(x_j) \right\}}.$$
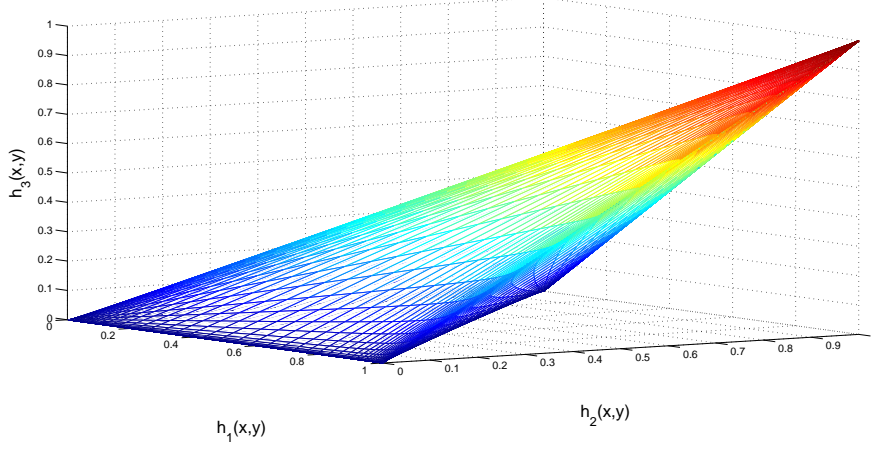
Figure 3: The feasible region for the expectation of g in the three dimensional example.

We then proceed, evaluating the feasible region beginning on a sphere of radius 1 and moving outwards in a radial direction, in the same manner as in the continuous case.

5.3. Fixed constraints

Suppose we have two constraints, α_1 and α_2. Let us further suppose that we fix α_2 for all time at some value $\tilde{\alpha}_2$. What effect does this have? Consider the mapping ϕ, in terms of $Z(\lambda_1, \lambda_2)$, as

$$\phi : \left(\frac{d}{d\lambda_1} \log Z(\lambda_1, \lambda_2), \frac{d}{d\lambda_2} \log Z(\lambda_1, \lambda_2) \right).$$

Fixing $\alpha_2 = \tilde{\alpha}_2$ is equivalent to fixing

$$\frac{d}{d\lambda_2} \log Z(\lambda_1, \lambda_2) = \tilde{\alpha}_2. \tag{5}$$

Now, when we consider α_1, which is given by

$$\frac{d}{d\lambda_1} \log Z(\lambda_1, \lambda_2) = \alpha_1,$$

it is a function of both λ_1 and λ_2. However, (5) has fixed the relationship between λ_1 and λ_2 and so λ_1 is now an implicit function of λ_2. Thus α_1 is now a function of one variable, and any value it takes will be "on the curve" between λ_1 and λ_2 implied by (5).

An example of such a curve for the dispersion and deposition example is given in Figure 10.

6. Example: Failure rates of machines in series

Suppose we have two machines, M_1, M_2, which are physically co-located and, within a larger process, are to be used in series. In this sense the two machines are then coupled as the physical outputs of one machine will become the physical inputs of the next machine. Suppose we are interested in the failure rates of the two machines, $\lambda_1(t), \lambda_2(t)$ respectively. In particular, we wish to evaluate the uncertainty distributions over the failure rates of the machines. It would seem reasonable that the failure rates, $\lambda_1(t), \lambda_2(t)$, would not be statistically independent as a result of the coupled nature of the machines and their physical co-location. For example, if it were not built into the modelling process elsewhere, an extreme weather event at the location of the machines would likely cause both to fail. This would result in the failure rates of the machines being correlated.

A simple model for the failure rate of machines which captures a constant failure rate due to random failures for a proportion of the lifetime of a machine or component and an increasing failure rate as a result of degradation beyond a certain age was proposed by van Gestel [11]. An illustration of the form of the model is given in Figure 4.
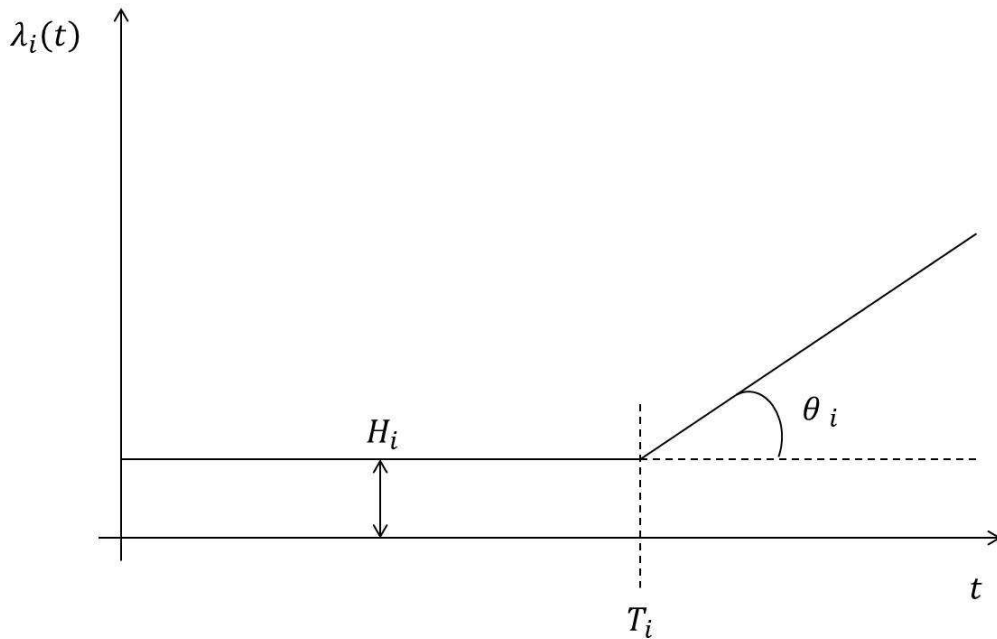

Figure 4: A simple model for the failure rate of machine i incorporating three parameters, (H_i, T_i, θ_i).

In Figure 4 we see that the model contains three parameters, (H_i, T_i, θ_i). An advantage of the model is that each of the parameters has an intuitive meaning; H_i is the failure rate of machine i under purely random failures, T_i is the time at which machine i moves from the random failures period into the wear out/degradation period and θ_i is the angle of the gradient in the failure rate during the wear out phase. It is possible to extend this model in order to incorporate early life failures by adding a linear section

with a negative gradient prior to the constant random failures section. This would add two additional parameters to the model, the changepoint time between early life and random failures and the angle of the gradient in early life. However, for this illustrative example, we shall consider the simpler three-parameter model.

If we return to the coupled machines we see that we have 6 parameters in total and so our parameter set is $x = (H_1, T_1, \theta_1, H_2, T_2, \theta_2)$. The changepoint time parameter of the first model, T_1, is an observable quantity and so an expectation on it can be elicited directly and will not affect the inference on the other parameters in the model. We could also ask experts about their expectation on the probability that the first failure of machine 1 was after T_1. This is given by the expectation of $\Pr(t > T_1) = \exp\{-H_1 T_1\}$. We discuss how to elicit such quantities in Section 6.1.

Thus, the first two constraints on machine 1 are given by

$$h_{1,1}(x) = T_1, \quad h_{1,2}(x) = \exp\{-H_1 T_1\}.$$

We would also like to elicit some information concerning θ_1 in order that we can evaluate the uncertainty on all of the machine 1 parameters. We can do so by asking the expert to assess the expected probability that machine 1 will have failed in the first year after it enters the ageing state. That is, we evaluate the expectation of $\Pr(t < T_1 + 1 \mid t \geq T_1) = [R(T_1) - R(T_1 + 1)]/R(T_1)$. This gives the third constraint as

$$
\begin{aligned}
h_{1,3}(x) \quad &= \quad \frac{\exp\{-H_1 T_1\} - \exp\left\{-H_1 T_1 - \frac{1}{2}(H_1 + \tan\theta_1)\right\}}{\exp\{-H_1 T_1\}} \\
&= \quad \left[1 - \exp\left\{-\frac{1}{2}(H_1 + \tan\theta_1)\right\}\right].
\end{aligned}
$$

It is clear that the specification of the first constraint will affect the possible values for the second and third constraints and so it will be important to perform the steps sequentially as in the previous example. We can also ask the expert to specify the same values for machine two giving three further constraints;

$$h_{2,1}(x) = T_2, \quad h_{2,2}(x) = \exp\{-H_2 T_2\},$$

$$h_{2,3}(x) = \exp\{-H_2 T_2\}\left[1 - \exp\left\{-\frac{1}{2}(H_2 + \tan\theta_2)\right\}\right].$$

As the failure rates of the machines will not, in general, be statistically independent, we are also interested in the joint behaviour of the different elements of the failure rates. To consider the correlations between the different elements we shall define a further constraint function relating parameters of the two machines. This is simply the probability that the difference between the times the two machines start ageing is less than one year. To define this probability we set the constraint equal to the indicator function for the event of interest. That is,

$$h_{(1,2),1}(x) = \begin{cases} 1, & \text{if } |T_1 - T_2| < 1, \\ 0, & \text{otherwise.} \end{cases}$$

6.1. Eliciting failure rate information from experts

The elicitation of expert judgements for probabilistic models should always be carried out on observable quantities [12, 1]. In the case of the failure rate model for machines considered here it is not reasonable to elicit information from experts about expectations on the failure rates directly. When he proposed the model, van Gestel [11] elicited information for a single machine from experts by asking them for the changepoint time between random failures and ageing T_i, the percentage of random failures per time period H_i and the mean time to failure.

Our minimum information approach differs as it is specified in terms of expected values. Our first proposed constraint for each machine is the expectation of the change-point time between random failures and the start of ageing. It is felt that engineers with experience of observing the failure behaviour of large numbers of similar machines over long time periods would be comfortable with this. The second constraint we have used for each machine is the expected probability that the first failure of the machine is after ageing starts. This is observable in the sense that the expert could be asked to specify this as the average proportion of machines which would not fail in this time over very large datasets. The third constraint for each machine, the expected probability that the machine will have failed a year after entering the ageing state given that it had not failed up to this time, can be elicited in a similar way.

The constraint which considers the relationship between failures in the two models is of a different form to the other constraints. Rather than asking the expert for an expected value directly, it instead asks the expert for a probability found as the expectation of an indicator variable. This probability is that ageing in the two machines begins less than a year apart. This is an observable quantity as the point at which a machine starts ageing is observable.

6.2. Simulating from the minimum information distribution

Making specifications on these 7 constraint functions sequentially will uniquely determine the minimum information distribution over the parameters of the failure rate functions for the two machines. Let us denote this distribution $g(x)$. We can then simulate from this distribution which gives us uncertainty distributions on the two failure rates of the machines. To do so we apply the following steps.

- Sample a uniform random variable $u = (u_1, \ldots, u_6)$, where $u_i \sim U(0,1)$.

- A sample from the MI distribution, $\tilde{x}$, can then be found. The first sampled parameter is $\tilde{x}_1 = G_1^{-1}(u_1)$, where $G_1^{-1}(\cdot)$ is the inverse cumulative distribution function of x_1. The rest of the parameters are then sampled sequentially as

$$\tilde{x}_i = G_i^{-1}(u_i \mid \tilde{x}_1, \ldots, \tilde{x}_{i-1}),$$

  where $G_k^{-1}(\cdot \mid \cdot)$ is the relevant one dimensional, conditional, cumulative distribution function.

- This gives one realisation of $\lambda_1(t)$ and $\lambda_2(t)$, the failure rates of the two machines.

- Repeat this a large number of times which gives uncertainty distributions on $\lambda_1(t)$ and $\lambda_2(t)$ representing the epistemic uncertainty of the expert.

Constraint function	Constraint	Lagrange multiplier	Value
$h_{1,1}(x)$	10.2000	$\lambda_{1,1}$	0.0991
$h_{1,2}(x)$	0.8160	$\lambda_{1,2}$	0.5996
$h_{1,3}(x)$	0.0119	$\lambda_{1,3}$	-73.2895
$h_{2,1}(x)$	7.9600	$\lambda_{2,1}$	-0.1346
$h_{2,2}(x)$	0.7674	$\lambda_{2,2}$	-5.1478
$h_{2,3}(x)$	0.0177	$\lambda_{2,3}$	-106.5014
$h_{(1,2),1}(x)$	0.2000	$\lambda_{(1,2),1}$	-10.5432

Table 2: The constraints on the parameters of the failure rates of the two machines.

Some pseudo-code to do this is given in Algorithm 3 in the Appendix. We would also like to calculate the predictive distribution of the failure rate for each machine. The failure rate of machine i at time t is given by

$$\lambda_i(t) = \begin{cases} H_i, & t < T_i, \\ H_i + \tan\theta_i(t - T_i), & t \geq T_i. \end{cases}$$

We can use this and the relationship that $R_i(t) = \exp\{-\int_0^t \lambda_i(s)ds\}$ to find the reliability of machine i at time t. This is

$$R_i(t) = \begin{cases} \exp\{-H_i t\}, & t < T_i, \\ \exp\left\{-H_i t - \dfrac{1}{2}\tan\theta_i(t - T_i)^2\right\}, & t \geq T_i. \end{cases}$$

The density of t can be calculated from the reliability as $f_i(t) = \partial/\partial t[1 - R_i(t)]$. In this case, again, we can find this analytically. It is

$$f_i(t) = \begin{cases} H_i \exp\{-H_i t\}, & t < T_i, \\ [H_i + \tan\theta_i(t - T_i)]\exp\left\{-H_i t - \dfrac{1}{2}\tan\theta_i(t - T_i)^2\right\}, & t \geq T_i. \end{cases}$$

The predictive distribution of the failure rate can now be calculated, from evaluations $j = 1, \ldots, J$ of the simulation at times $t_k, k = 1, \ldots, K$ resulting in reliability $R_{i,j}(t_k)$ and density $f_{i,j}(t_k)$, as

$$\lambda_i(t_k) = \frac{\sum_{j=1}^{J} f_{i,j}(t_k)}{\sum_{j=1}^{J} R_{i,j}(t_k)}. \tag{6}$$

6.3. Numerical Illustration

Suppose that we have elicited the expectations of the constraints given in the previous section for the two machines using the sequential method outlined in the paper and illustrated explicitly in the next example. The resulting constraints for the minimum information distribution are given in the left hand side of Table 2.

Using these constraints we can construct the discretized minimum information distribution of the form of (2). The discretization used in this example takes 50 points in

each dimension. All background distributions are taken to be uniform. The choice of background distribution will be investigated further later in this example. The resulting minimum information distribution has Lagrange multipliers as given in the right hand side of Table 2.

In the previous section we gave a method for sampling from the distribution on the failure rate over the three parameters for a single machine. A plot of 100 samples of (H_1, T_1, θ_1) taken from the minimum information distribution with the Lagrange multipliers from Table 2 is given in Figure 5.
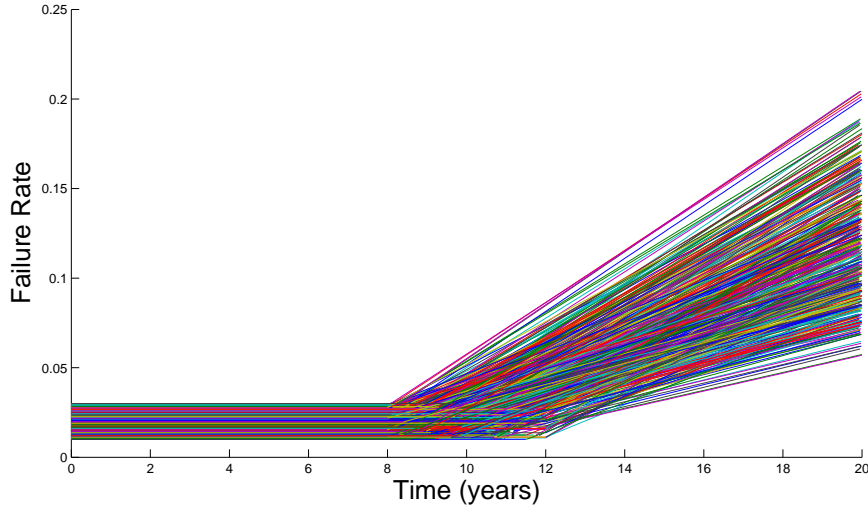

Figure 5: Realisations of the failure rate of machine 1 resulting from 100 samples from the minimum information distribution.

We see that there is uncertainty evident in all three parameters of the failure rate function. The height of the constant rate prior to the change point varies in the samples, though by a relatively small amount. There is also variation in the changepoint time between the constant and linear sections of the plot which represents the uncertainty on T_1. The greatest uncertainty is in the linear section of the plot. The uncertainty in T_1 and θ_1 both contribute to this.

It is important to check that the ranges used for the background distributions of the parameters are sufficiently large. We do not wish for the majority of the density for any parameter to be close to the boundaries of this background distribution. Doing so would indicate that the range of the background distribution should be increased. In general this can be checked by plotting the marginal densities of the parameters. In this case a simple heuristic is to consider the simulated parameter values in Figure 5. As long as there aren't a large number of simulated values close to the background distribution boundaries then the ranges of the distributions are sufficient. This is the case in this example.

We can also calculate the predictive distribution of the failure rates using the 100

samples from the minimum information distribution for this example. It is given for machine 1, for the first 15 years of life, in Figure 6.
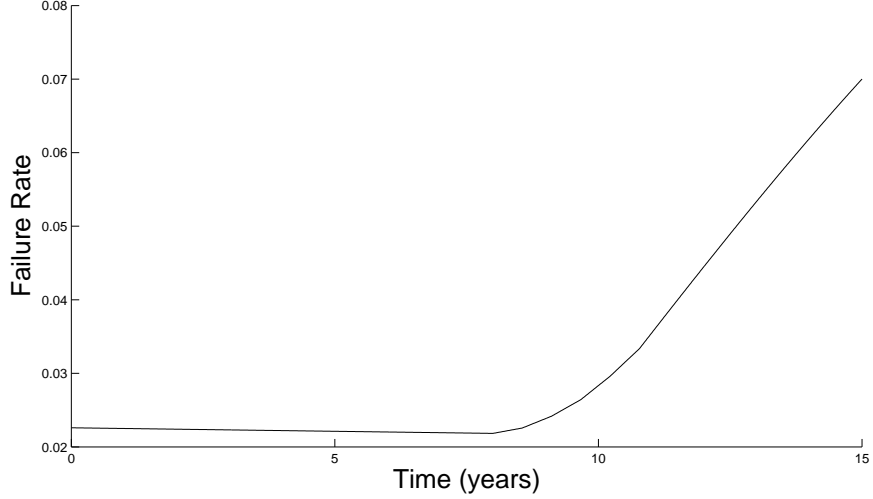

Figure 6: The predictive distribution of the failure rate for machine 1 calculated using the 100 samples.

We see that the predictive distribution is smoother than the individual realisations from the simulation. In particular the changepoint between random failures and ageing is fairly smooth This is due to the averaging of the simulations in (6). We also see a slight non-linear quality to both sections of the failure rate in the predictive distribution. This is again as a result of averaging over non-linear quantities in (6).

We can also use this example to examine the effect of the assumption of background distribution on the minimum information solution. In Figure 7 we see the predictive distribution for the failure rate of machine 1 under uniform, log-uniform and Normal background distributions for each of the parameters. All three choices of background distribution give a similar shape in the random failure period of the machine's life. The uniform and log-uniform distributions also give very similar solutions in the ageing section of the plot. The predictive distribution using Normal background distributions is further from these in this section. This illustrates the way in which a sensitivity assessment can be made of the degree to which the distribution depends on assumptions made about the background distribution in the minimum information method. In this case it shows that the background distribution makes little difference except in the tail.

7. Example: Atmospheric dispersion and deposition

7.1. One-stage dispersion model

Atmospheric dispersion is the process by which gases, which could be harmful or even radioactive, travel through the atmosphere from a point of origin such as a factory or
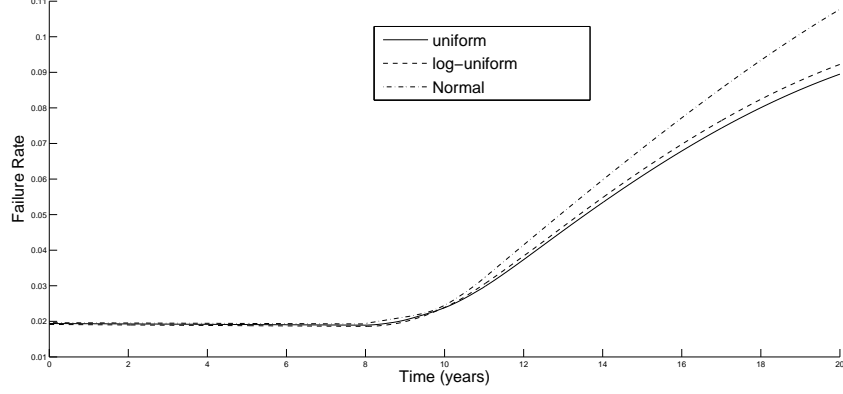
Figure 7: The predictive distribution of the failure rate for machine 1 calculated using the 100 samples under uniform, log-uniform and Normal background distributions.

power plant. One commonly used model to describe such a process is the Gaussian plume model. This model is based on simplifying assumptions such as a constant wind direction, no plume rise and the land over which the contaminant is travelling being fairly even, [7].

The model is defined in terms of the downwind distance x from the source, the crosswind displacement y and the vertical displacement z. It then assumes that the plume concentration is of Gaussian form,

$$C(x, y, z) = \frac{Q_0}{2\pi u \sigma_y(x) \sigma_z(x)} \exp\left\{ -\frac{y^2}{2\sigma_y^2(x)} \right\} \left[\exp\left\{ -\frac{(z - H)^2}{2\sigma_z^2(x)} \right\} + \exp\left\{ -\frac{(z + H)^2}{2\sigma_z^2(x)} \right\} \right],$$

where $C(x, y, z)$ is the time integrated concentration of the contaminant at (x, y, z), Q_0 is the initial quantity of contaminant released, u is the windspeed and H is the centreline of the plume. The two standard deviations, $\sigma_y(x)$ and $\sigma_z(x)$, are the lateral and vertical plume spreads at x respectively. They are commonly assumed to follow deterministic power law models,

$$\begin{aligned} \sigma_y(x) &= a_y x^{b_y}, \\ \sigma_z(x) &= a_z x^{b_z}, \end{aligned}$$

for unknown parameters a_y, b_y, a_z, b_z. That is, both lateral and vertical plume spreads are functions of the downwind distance.

Let us consider the lateral plume spread. The parameters of the power law model, a_y and b_y, do not have simple physical interpretations. They can be inferred from tracer experiments but this is unreliable. As they are unobservable we cannot elicit their values from experts. However, σ_y is observable. Having elicited a distribution for σ_y we can then use this to define probability distributions over a_y, b_y. This is probabilistic inversion.

By contrast σ_z is not observable. Harper *et al* [7] indicate, however, that the centreline

concentration ratio,

$$C(x, 0, H) = \frac{Q_0}{2\pi\sigma_y\sigma_z u},$$

which is the Gaussian plume model evaluated at $y = 0, z = H$, is. Taking expectations and rearranging, if we elicit $\mathrm{E}[C(x, 0, H)]$ from the expert, then

$$\frac{2\pi u}{Q_0}\mathrm{E}[C(x, 0, H)] = \mathrm{E}\left[\frac{1}{a_y a_z x^{b_y + b_z}}\right].$$

Since our notation is slightly different to that used up to now it is worth considering how this model fits into the framework we have adopted above. The model consists of three elements (C, σ_y, σ_z), so its outputs are the plume concentration and the lateral and vertical plume spreads. The model inputs are Q_0, u and H. The model parameters are (x, y, z), the place we are interested in plume concentration, which we choose, and a_y, b_y, a_z, b_z which are unknown, and whose distributions will be inferred through PI.

7.1.1. Numerical illustration

Suppose we are interested in two downwind distances, 1km and 2km from the source. We could ask the expert for the expected lateral plume spreads and the concentration ratios at each of these distances. That is, the expert specifies,

$$\mathrm{E}(\sigma_{y,1}) = \mathrm{E}(a_y) = \alpha_1, \quad \mathrm{E}(\sigma_{y,2}) = \mathrm{E}(a_y 2^{b_y}) = \alpha_2,$$

$$\mathrm{E}\left(\frac{1}{a_y a_z}\right) = \alpha_3, \quad \mathrm{E}\left(\frac{1}{a_y a_z 2^{b_y + b_z}}\right) = \alpha_4,$$

as described above for the plume spreads and ratios at 1km and 2km respectively. However, the range of possible values for α_2 will depend on the chosen α_1, the possible values for α_3 depend on α_1 and α_2 and so on. We shall find the minimum information distribution satisfying the constraints above using the methodology set out in Section 4. Some pseudo-code for defining a minimum information distribution in this way is given in Algorithm 2 in the Appendix.

The background distributions we use to find this minimum information distribution are taken to be uniform with respect to $\log(a_y)\log(a_z), b_y, b_z$. This follows work in Kraan and Bedford [2]. The discretizations take 50 points in each dimension with suitable ranges chosen for each parameter. The ranges used for the background distributions are

$$\log a_y, \log a_z \in (0, 4), \quad b_y, b_z \in (0, 2).$$

Initially, given these ranges for the parameters, the possible ranges the expectations can take are

$$\alpha_1 \in (1, 55), \quad \alpha_2 \in (1, 218),$$
$$\alpha_3 \in (0, 1), \quad \alpha_4 \in (0, 1).$$

All are found by calculating the expectation α_i for large positive and negative values of the relevant λ_i. Figure 8 indicates the range of possible values of the first constraint α_1 and how they relate to the value of λ_1. We see that for large positive and negative λ_1, α_1 tends to the limits of its range of possible values.
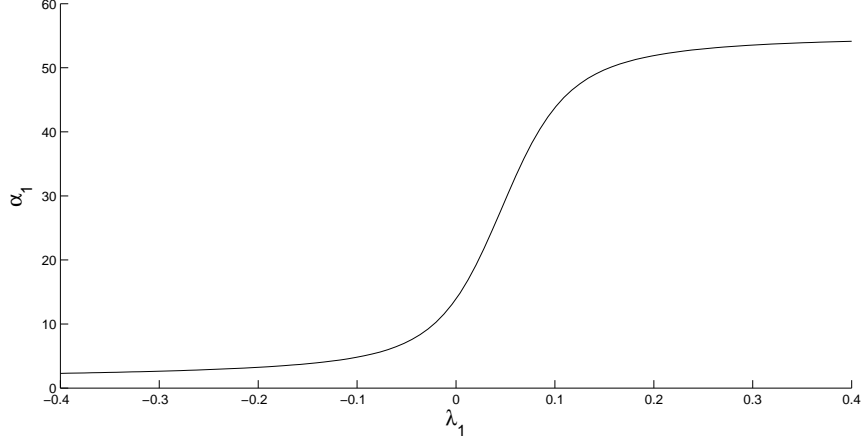
Figure 8: The possible range of α_1 and corresponding values of λ_1.

We wish to find an initial value for α_2 so we can calculate its range under this specification for α_1. Keeping $\lambda_1 = 0.116$ and setting $\lambda_2 = 0$ gives the expectations $\alpha_1 = 45, \alpha_2 = 97.691$. Using this as a starting point we can calculate the range of values α_2 can take given this α_1. This is

$$\alpha_2 \mid \alpha_1 = 45 \in (45, 180).$$

We can see that some values which were possible originally are now not.

We can calculate the range of possible values for α_2 given different specifications of α_1. Figure 9 shows these as points for 5 chosen values of α_1 and compares these to the feasible region mapped out using the methods described in Section 5. We see that the two different methods are in agreement. The possible range of α_2 is smallest for low values of α_1 and larger for high values. It is always reduced from the original possible range of $(1, 218)$.

Let us suppose that the expert, on receipt of this information, specifies $\alpha_2 = 80$. The parameters of the minimum information distribution become

$$\lambda_1 = 0.1401, \quad \lambda_2 = -0.0128.$$

Note that λ_1 has changed as α_1 now depends on both λ_1 and λ_2. Using these parameter values and $\lambda_3 = 0$ leads to an initial expectation of $\alpha_3 = 0.0082$. From this we calculate the range of possible values for α_3. They are

$$\alpha_3 \mid (\alpha_1 = 45, \alpha_2 = 80) \in (0.00336, 0.186).$$

Again this is reduced from the original range.

If the expert specifies $\alpha_3 = 0.0105$ then the parameters of the minimum information distribution are $\lambda_1 = 0.1425, \lambda_2 = -0.0128, \lambda_3 = 2.3025$. Repeating this process for the final constraint gives a feasible range of

$$\alpha_4 \mid (\alpha_1 = 45, \alpha_2 = 80, \alpha_3 = 0.0105) \in (0.00214, 0.0781).$$
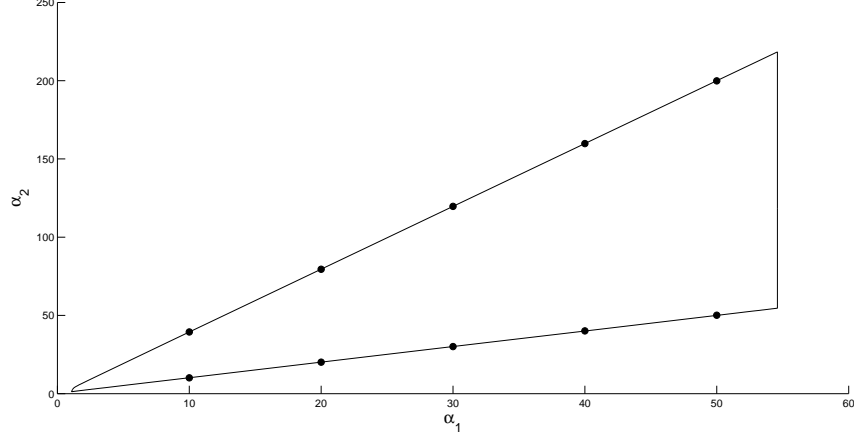
Figure 9: The feasible region of α_2 evaluated for different specified values of α_1.

This is significantly smaller than the initial possible range of $(0, 1)$. Therefore this process will be useful in guiding the expert to make a plausible choice for α_4. If the expert settles on $\alpha_4 = 0.0038$ then the final parameter values are $\lambda_1 = 0.1315, \lambda_2 = -0.0124, \lambda_3 = -0.4500, \lambda_4 = 4.8963$.

7.2. Two-stage coupled dispersion and deposition model

Another quantity of interest when a contaminant is released into the atmosphere is how much of it is deposited into the ground at different distances from the source. This is known as deposition. Two different types are considered generally depending on the absence or presence of rain; dry and wet deposition.

Dry deposition can be measured in terms of the time integrated surface contamination at $(x, y, 0)$. This measures how much contaminant is deposited into the ground from the plume and is given by

$$
\begin{aligned}
C_d(x, y, 0) &= v_d C(x, y, 0), \\
&= v_d \frac{Q_0}{\pi \sigma_y \sigma_z} \exp\left\{ -\left(\frac{y^2}{2\sigma_y^2} + \frac{H^2}{2\sigma_z^2} \right) \right\}
\end{aligned}
$$

where v_d is an unknown parameter called the dry deposition velocity. We can elicit expectations for $C_d(x, y, 0)$ at different downwind distances and crosswind displacements. We see that the dispersion and deposition models are coupled.

In this two-stage example we shall consider eliciting expectations for the lateral plume spread and the surface contamination at the two downwind distances considered previously and a crosswind displacement of 1km. That is, we elicit the means of

$$
\sigma_{y,1}, \sigma_{y,2}, C_d(1, 1, 0), C_d(2, 1, 0).
$$

Thus, the first two constraints, α_1 and α_2, are the same as in Section 7.1.1. We also have

$$
\begin{aligned}
\alpha_3 &= \mathrm{E}[C_d(1,1,0)], \\
\alpha_4 &= \mathrm{E}[C_d(2,1,0)].
\end{aligned}
$$

We then wish to use these elicited values to define minimum information distributions over the model parameters

$$
\Omega = (a_y, b_y, a_z, b_z, v_d).
$$

Background distributions will be as previously for a_y, b_y, a_z, b_z and uniform on the log-scale for v_d as it is multiplicative.

7.2.1. Numerical Illustration

The first two expectations, α_1 and α_2, will have the same initial possible ranges as previously. The corresponding initial ranges for α_3 and α_4 are

$$
\begin{aligned}
\alpha_3 &\in [1 \times 10^{-4}, 0.1931], \\
\alpha_4 &\in [1 \times 10^{-4}, 0.1931].
\end{aligned}
$$

We will set $\alpha_1 = 45$ and $\alpha_2 = 80$ as previously. However, first suppose that in order to do so it is decided that no matter what the values of other constraints α_2 must be 80. Therefore when we wish to specify α_1, this specification for α_2 defines an implicit relationship for λ_1, λ_2. A plot of the curve representing this relationship is given in Figure 10.
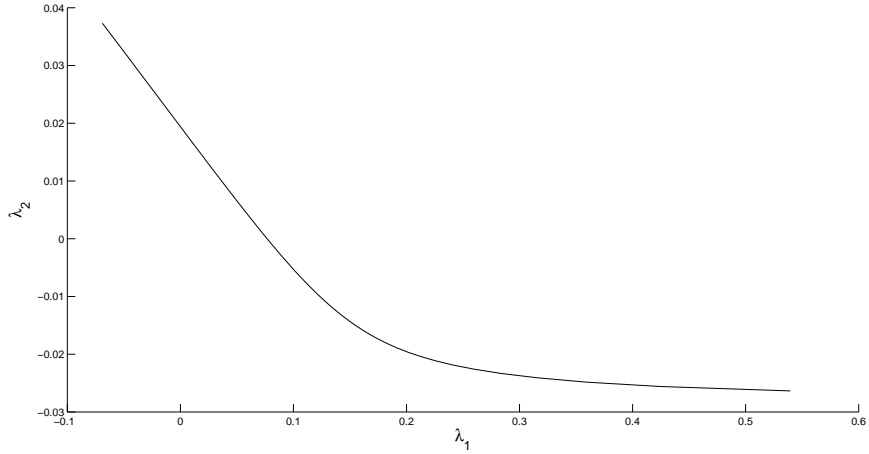

Figure 10: The curve representing the the implicit relationship between λ_1 and λ_2 having specified $\alpha_2 = 80$.

As before this leads to parameter values of $\lambda_1 = 0.1401$ and $\lambda_2 = -0.0128$. We now wish to find the range of possible values for α_3 given that α_1 and α_2 take these values.

This includes a new model parameter v_d as well as all of the dispersion model parameters from the one step example. It also includes the input from the dispersion model, x, as well as a new input, y.

We proceed as in the one step case. Initially we set $\lambda_3 = 0$ to find a starting point for α_3 of 6×10^{-4}. Working from this we find the range of possible specifications for α_3 given the two constraints, which is

$$\alpha_3 \mid (\alpha_1 = 45, \alpha_2 = 80) \in [1 \times 10^{-4}, 3.9 \times 10^{-3}].$$

Suppose that, given this range, the expert specifies $\alpha_3 = 6.5 \times 10^{-4}$. The parameter values which satisfy all three specifications are $\lambda_1 = 0.1207, \lambda_2 = -0.0121, \lambda_3 = 0.2459$. We proceed in the same manner for α_4. Setting $\lambda_4 = 0$ gives an initial value for α_4 of 2×10^{-4}. This allows us to calculate the admissible range for α_4 given these specifications for the other constraints. It is

$$\alpha_4 \mid (\alpha_1 = 45, \alpha_2 = 80, \alpha_3 = 6.5 \times 10^{-4}) \in [1 \times 10^{-4}, 6 \times 10^{-4}].$$

Suppose we select $\alpha_4 = 2.14 \times 10^{-4}$. Then

$$\lambda_1 = 0.1207, \lambda_2 = -0.0121, \lambda_3 = 0.0403, \lambda_4 = 0.6758.$$

Thus we have fully defined the minimum information distribution which satisfies the four constraints for this two-stage example in such a way that ensures the expert always gives feasible values.

7.3. Discussion

Within the example we have worked with constraints sequentially, one at a time. In doing so it is clear that the order in which the expectations are specified is of importance. In some cases the order will be clear. For example, in a chain of models in which the output of one model is the input of the next, the constraints follow a natural order. However, this will not always be the case.

In a subjective expert elicitation procedure psychological anchoring occurs when an expert is given a value for an unknown (the anchor) and is then asked to adjust this in light of new information. The adjustment the expert makes has been found, Slovic [13], to be too small in general. Hence the expert is anchored by the original value, whether or not that value is accurate.

By sequentially specifying the expectation values in this minimum information approach we are introducing another form of anchoring, which we shall refer to as quantitative anchoring. That is, having specified the initial expectation, the expert is then anchored within the reduced range for the second expectation. They are then anchored to the new range for the third expectation given the first two, and so on. Clearly the order in which this takes place could have an effect on the resulting probability distributions for the model parameters. More work is needed to investigate a suitable procedure for deciding on the order to specify the constraints in light of this anchoring.

The benefit of the procedure discussed, though, is that experts will provide coherent specifications, which eliminates the problem of over-specification of the distribution on the parameters. Clearly one could use the methods developed here interactively to explore uncertainties on output values and hence negate the possible issue of anchoring.

8. Conclusion

In this paper we have investigated an approach to uncertainty modelling on coupled models in which an expert specifies the values of expectations of observable model outputs and these are used to specify probability distributions on unobservable model parameters. Generally this problem is either over- or under-specified and the methodology developed deals with both situations.

For under-specified problems we have solved for the distribution on the unobservable model parameters which has minimum information. This has been achieved numerically by first discretizing the problem. To avoid over-specification we consider a process of specifying the constraints sequentially. This is useful as it allows us to guide the expert in choosing values that do not conflict with previously chosen values.

The extension of this methodology to coupled models has allowed us to examine the range of possible values constraints can take for a second model, which can have some model inputs and outputs in common with the first model and others unique to itself. Again, the specifications of expectations for model outputs in the first model has an effect on the possible values of constraints in the second model.

The main advantage we see for this approach to modelling is that its flexibility enables us to build in features that we would qualitatively like to see, such as the onset of ageing in the second example or the physical models of the first example. The uncertainty modelling can then be carried out around these qualitative features. Rather than working with models which are mathematically convenient, we can work with models that possess the right qualitative properties.

The main difficulties of working with minimum information distributions are that they require numerical computation. While the curse of dimensionality is applicable here as in any other approach, there are opportunities to speed up numerical computation, for example by using variable grid sizes, and interesting opportunities to develop measures of coherence associated to the Lagrange multipliers.

Acknowledgements

Appendix

Algorithm 1 Exploration in a radial direction of the feasible region of a 2 dimensional minimum information distribution.

Set radius to r
Define background distributions for x_1, x_2
for $i \leftarrow 1, \ldots, p$ **do**
  $e_i = h_i(x_1, x_2)$
end for
Define λ_1 as a discretization on $[-1, 1]$
$\lambda_2^{(p)} = \sqrt{1 - \lambda_1^2}$
$\lambda_2^{(n)} = -\sqrt{1 - \lambda_1^2}$
$\lambda_1 = r\lambda_1$
$\lambda_2^{(p)} = r\lambda_2^{(p)}$
$\lambda_2^{(n)} = r\lambda_2^{(n)}$
$\lambda_i = (\lambda_{1,i}, \lambda_{2,i}^{(p)}, \lambda_{2,i}^{(n)})$
for $i \leftarrow 1, \ldots, \text{length}(\lambda_1)$ **do**
  $\alpha_{1,i} = \mathrm{E}_{\lambda_i}[h_1(x_1, x_2)]$
  $\alpha_{2,i} = \mathrm{E}_{\lambda_i}[h_2(x_1, x_2)]$
end for

Algorithm 2 Defines the minimum information distribution satisfying $(\mathrm{E}[h_1(x)], \ldots, \mathrm{E}[h_p(x)]) = (e_1, \ldots, e_p)$.

Set background distributions for $x_1, \ldots, x_p$
$\alpha_1 = e_1$
$\lambda_1^{(1)} = 0$
for $i \leftarrow 2, \ldots, n$ **do**
  $\lambda_1^{(i)} = \lambda_1^{(i-1)} - [\psi(\lambda_1^{(i-1)}) - \alpha_1]/\psi'(\lambda_1^{(i-1)})$
end for
for $j \leftarrow 2, \ldots, p$ **do**
  $\alpha_j = (\alpha_{j-1}, e_j)$
  $\lambda_j^{(1)} = (\lambda_{j-1}^{(n)}, 0)$
  for $i \leftarrow 2, \ldots, n$ **do**
    $\lambda_j^{(i)} = \lambda_j^{(i-1)} - [\psi(\lambda_j^{(i-1)}) - \alpha_j]/\psi'(\lambda_j^{(i-1)})$
  end for
end for

Algorithm 3 Simulation from a minimum information distribution $g(x)$ of a sample of size N based on discretizations $x_{j,k}$ of each variable x_j.

---

  for $i \leftarrow 1, \ldots, N$ **do**
    $(u_1, \ldots, u_p)$ sampled from independent uniform distributions on $[0, 1]$
    for $k \leftarrow 1, \ldots, m$ **do**
      $d_1(x_1) = g(x_1)$
      $D_1(x_1) = G(x_1)$
      if $D_1(x_1) > u_1$ **then**
        $\tilde{x}_1 = x_{1,k}$
      end if
    end for
    for $j \leftarrow 2, \ldots, p$ **do**
      for $k \leftarrow 1, \ldots, m$ **do**
        $d_j(x_j) = g(x_j \mid \tilde{x}_1, \ldots, \tilde{x}_{j-1})$
        $D_j(x_j) = G(xj \mid \tilde{x}_1, \ldots, \tilde{x}_{j-1})$
        if $D_j(x_{j,k}) > u_j$ **then**
          $\tilde{x}_j = x_{j,k}$
          Stop
        end if
      end for
    end for
  end for

References

[1] T. Bedford and R. Cooke. *Probabilistic Risk Analysis: Foundations and Methods.* Cambridge University Press, Cambridge, 2001.

[2] B. Kraan and T. Bedford. Probabilistic inversion of expert judgments in the quantification of model uncertainty. *Management Science*, 51(6):995–1006, 2005.

[3] Dorota Kurowicka and Roger Cooke. *Uncertainty Analysis - with High Dimensional Dependence Modelling.* Wiley, 2006.

[4] T.J. Bedford. Interactive expert assignment of minimally-informative copulae. In B. Lindqvist, editor, *Proceedings of Mathematical Methods in Reliability 2002.* 2002.

[5] E. T. Jaynes. *Probability Theory: the Logic of Science.* Cambridge, 2003.

[6] J. Borwein and A. Lewis. *Convex analysis and nonlinear potimization: theory and examples.* Canadian Mathematical Society, 2006.

[7] F.L. Harper, L. Goosens, R. Cooke, J. Helton, S. Hora, J. Jones, B. Kraan, C. Lui, M. McKay, L. Miller, J. Pasler-Sauer, and M. Young. Joint USNRC/CEC consequence uncertainty study: Summary of objectives, approach, application, and results for the dispersion and deposition uncertainty assessments. Technical report, NUREG/CR-6244 and EUR 15855 and SAND94-1453, Sandia National Laboratories and Delft University of Technology, Washington DC.

[8] O. E. Lanford. Entropy and equilibrium states in classical statistical mechanics. In A. Lenard, editor, *Statistical Mechanics & Mathematical Problems*, volume 20 of *Springer Lecture Notes in Physics*, pages 1–113. 1973.

[9] T. Bedford and K.J. Wilson. On the construction of minimum information bivariate copula families. *Submitted to Annals of the Institute of Statistical Mathematics*, 2012.

[10] T. Bedford, A. Daneshkhah, and K.J. Wilson. Approximate uncertainty modeling with vine copulas. *Submitted to Operations Research*, 2012.

[11] P.J. van Gestel. Kmoss, a maintenance optimization support system. In *Proc. Scand. SRE Symp.*, Studsvik, Sweden, 1990.

[12] J.B. Kadane. Experiences in elicitation. *The Statistician*, 47:3–19, 1998.

[13] P. Slovic. From shakespeare to simon: speculations - and some evidence - about man's ability to process information. *Oregon Research Institute Research Bulletin*, 12, 1972.