ENVIRONMETRICS
Environmetrics 2003; **14**: 355–371 (DOI: 10.1002/env.590)

Bayesian uncertainty assessment in multicompartment deterministic simulation models for environmental risk assessment

Samantha C. Bates[1,*,†], Alison Cullen[2] and Adrian E. Raftery[3]

[1]*Department of Statistics, Virginia Tech, Hutcheson Hall (0439), Blacksburg, VA 24061, U.S.A.*
[2]*Daniel J. Evans School of Public Affairs, University of Washington, Box 353055, Seattle, WA 98195-4322, U.S.A.*
[3]*Department of Statistics, University of Washington, Box 354322, Seattle, WA 98195-3055, U.S.A.*

SUMMARY

We use a special case of Bayesian melding to make inference from deterministic models while accounting for uncertainty in the inputs to the model. The method uses all available information, based on both data and expert knowledge, and extends current methods of 'uncertainty analysis' by updating models using available data. We extend the methodology for use with sequential multicompartment models. We present an application of these methods to deterministic models for concentration of polychlorinated biphenyl (PCB) in soil and vegetables. The results are posterior distributions of concentration in soil and vegetables which account for all available evidence and uncertainty. Model uncertainty is not considered. Copyright © 2003 John Wiley & Sons, Ltd.

KEY WORDS: bayesian melding; deterministic models; sampling importance resampling; uncertainty analysis

1. BACKGROUND

Risk is the probability of an adverse outcome, and risk assessment is a process by which an estimate of this probability is obtained. A health risk assessment generally follows four steps: hazard identification, dose–response assessment, exposure assessment and risk characterization. In the hazard identification step, the potential hazard and its adverse health effects (if any) are identified. The exposure assessment step identifies populations which could be exposed and the pathways by which this may occur. The dose–response step quantifies the relationship between levels of exposure and the level of potential adverse effect (NRC, 1983). The final risk characterization step involves combining the results of the previous three steps to determine some outcome of interest, for instance, the risk of an adverse health effect for an individual exposed to a contaminant. This article focuses on problems in the exposure assessment phase of a risk assessment. We consider the case in which a deterministic model relates an output, such as contaminant concentration, to a group of inputs.

Until the 1980s most literature on quantitative methods for exposure assessment concentrated on purely deterministic approaches whereby point values were assigned to the inputs in an exposure

*Correspondence to: Samantha C. Bates, Department of Statistics, Virginia Tech, Hutcheson Hall (0439), Blacksburg, VA 24061, U.S.A.
†E-mail: sbates@vt.edu

Contract/grant sponsor: U.S. Environmental Protection Agency; contract/grant number: CR825173-01-0.

Published online 10 April 2003

Received 27 June 2001
Revised 31 July 2002

model. However, inputs are often not known precisely. This lack of information is often referred to in the risk assessment community as *uncertainty.*

In the 1980s authors began to call for a probabilistic approach to exposure assessment to account for uncertainty (Morgan *et al.*, 1985; Bogen and Spear, 1987; IAEA, 1989; Finkel, 1990). Our aim here is to present a Bayesian method which accounts for the uncertainty about the values of model inputs and propagates this through to the output. This will be done by using all available relevant information, both application-specific empirical data and expert knowledge. The result will be a posterior distribution for an outcome of interest which accounts for uncertainty and which informs decision-makers about the range of potential exposure in the population and the associated level of uncertainty.

In Section 2 we present the Bayesian approach for non-sequential and sequential multicompartment deterministic simulation models. We follow in Section 3 with an application to a one-compartment model for predicting contaminant concentration in soil, including a comparison of the Bayesian approach to the current state-of-the-art method of two-dimensional Monte Carlo (Hoffman and Hammonds, 1994; McKone, 1994; Cullen and Frey, 1999). This method separates the contributions of uncertainty due to lack of knowledge from variability due to natural heterogeneity in the population of interest or across time and space, but does not use the empirical data on model outputs. Section 4 contains an application to a sequential two-compartment model for predicting contaminant concentration in vegetables. We conclude with discussion.

2. METHODOLOGY

2.1. Bayesian melding for deterministic simulation models

Suppose we have a deterministic simulation model M which maps inputs $\boldsymbol{\theta}_1$ to outputs $\boldsymbol{\phi}_1$. The quantities $\boldsymbol{\theta}_1$ and/or $\boldsymbol{\phi}_1$ may be single or vector valued. We refer to $M(\boldsymbol{\theta}_1)$ as the induced output, i.e. the value of the output that is induced by the model from the input $\boldsymbol{\theta}_1$. Further, suppose that the available knowledge about the values of the inputs is represented by a prior distribution, which we denote by $p(\boldsymbol{\theta}_1)$.

Often in environmental risk assessments, the collection of data specific to the assessment is expensive or impossible, and so such data tend to be sparse. We would like the method used to make inference from the model to incorporate all available information on the model inputs and outputs, including the often limited data. This observed data can be used to form likelihoods for $\boldsymbol{\theta}_1$ and $\boldsymbol{\phi}_1$, which are denoted respectively by $L_{\text{inp}}(\boldsymbol{\theta}_1)$ and $L_{\text{out}}(\boldsymbol{\phi}_1)$. In any specific application, both, either or neither of these likelihoods may be present.

Let D denote the observed data. The marginal posterior distribution of the inputs, $\pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_1)$, is the conditional distribution of $\boldsymbol{\theta}_1$ given both D and the model M. Suppose we can partition D into components relating to $\boldsymbol{\theta}_1$ and $\boldsymbol{\phi}_1$, respectively: $\{D_{\boldsymbol{\theta}_1}, D_{\boldsymbol{\phi}_1}\}$. Then, by Bayes's theorem,

$$\pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_1) \;\propto\; p(\boldsymbol{\theta}_1) P(\{D_{\boldsymbol{\theta}_1}, D_{\boldsymbol{\phi}_1}\} \mid \boldsymbol{\theta}_1, M) \tag{1}$$

and $\pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_1)$ is proportional to the product of the prior on inputs and the joint likelihood evaluated at the inputs and induced output values. If $D_{\boldsymbol{\theta}_1}$ and $D_{\boldsymbol{\phi}_1}$ are conditionally independent given $\boldsymbol{\theta}_1$ and $\boldsymbol{\phi}_1$, then the likelihood of inputs and outputs factorizes into components corresponding to the inputs and outputs, and Equation (1) has the following form:

$$\pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_1) \;\propto\; p(\boldsymbol{\theta}_1)\, L_{\text{inp}}(\boldsymbol{\theta}_1) L_{\text{out}}(M(\boldsymbol{\theta}_1)) \tag{2}$$

Then $\pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_1)$ is proportional to the product of three factors: the prior on the inputs, the likelihood of the inputs, and the likelihood of the outputs evaluated at the induced values.

Inference about $\boldsymbol{\phi}_1$, or any function of it, can be made from its marginal posterior distribution, the distribution of $\boldsymbol{\phi}_1 = M(\boldsymbol{\theta}_1)$ when $\boldsymbol{\theta}_1 \sim \pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_1)$. Thus we need a method to find $\pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_1)$. Often, the analytical form of $\pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_1)$ is intractable and obtaining an exact sample from it can be difficult or impossible. One method to obtain an approximate sample from $\pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_1)$ is to use the sampling importance resampling (SIR) algorithm of Rubin (1988). This algorithm involves initially sampling from a known distribution and resampling according to SIR weights.

When, as in our application, the prior is easily sampled from and is used as the importance sampling distribution, each SIR weight is the ratio of $\pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_1)$ and $p(\boldsymbol{\theta}_1)$ for a value of $\boldsymbol{\theta}_1$ sampled from the prior. Let $\boldsymbol{\theta}_{1_i}$ denote the ith sample from the prior on the inputs, $p(\boldsymbol{\theta}_1)$, with corresponding induced value $\boldsymbol{\phi}_{1_i} = M(\boldsymbol{\theta}_{1_i})$. Using the relation in Equation (1), the SIR weight associated with $\boldsymbol{\theta}_{1_i}$ and $\boldsymbol{\phi}_{1_i}$ is

$$w_i = \frac{\pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_{1_i})}{p(\boldsymbol{\theta}_{1_i})} \propto P(\{D_{\theta_1}, D_{\phi_1}\} \mid \boldsymbol{\theta}_{1_i}, M) \tag{3}$$

Calculation of the SIR weights is straightforward regardless of whether conditional independence of D_{θ_1} and D_{ϕ_1} is assumed. More details on the use of the SIR algorithm in this context can be found in Poole and Raftery (2000). Givens (1993) and Hesterberg (1995) discuss diagnostics for the SIR algorithm.

2.2. *Approach for sequential multicompartment models*

We now consider the situation where our model, M, consists of two *compartments*, or sub-models, M_1 and M_2. In a general two-compartment model, outputs from M_1 can be inputs to M_2, and vice versa i.e. there can be feedback in the system. Inputs can also be common to both compartments. Here we will consider only *sequential multicompartment models* in which inputs to M_1 and outputs from M_1 can be inputs to M_2, but inputs and outputs of M_2 cannot be inputs to M_1.

Figure 1 represents such a scheme. The inputs to an initial deterministic model M_1 are partitioned into two groups, $\boldsymbol{\theta}_1'$ and $\boldsymbol{\beta}$. M_1 predicts some quantities of interest $\boldsymbol{\phi}_1$. $\boldsymbol{\beta}$ represents those inputs which are used only in M_1. M_2 is a secondary model which maps $\boldsymbol{\theta}_1'$, $\boldsymbol{\phi}_1$ and a distinct set of inputs $\boldsymbol{\theta}_2$ to a new output $\boldsymbol{\phi}_2$. This figure shows an example where there are data D_{θ_1} on inputs to M_1 and data (denoted by D_{ϕ_1} and D_{ϕ_2}) on the respective model outputs. There may be additional data on inputs to the second model.

We are now interested in inference on $\boldsymbol{\phi}_2$. We will apply the same general approach as before for the one-compartment model. However we now need to proceed in two stages: first sample from the posterior distribution, $\pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_1)$, of inputs to M_1 as before. Note that this induces a posterior distribution of the outputs which is easily available as a byproduct of the SIR algorithm. Second, produce a sample from the posterior distribution of the inputs and outputs to M_2. We will do this by using the posterior distribution of inputs to M_1, $\pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_1)$, as the importance sampling function for $\boldsymbol{\theta}_1'$ and $\boldsymbol{\phi}_1$.

The sample from the marginal posterior distribution of inputs to M_2 can be obtained first by sampling from the joint posterior of all inputs $\boldsymbol{\theta}_2$ and $\boldsymbol{\theta}_1 = \{\boldsymbol{\theta}_1', \boldsymbol{\beta}\}$ given all the data and sub-models. It can be shown that this posterior has the following form:

$$\begin{aligned} P(\boldsymbol{\theta}_2, \boldsymbol{\theta}_1 \mid D, M_1, M_2) &= P(\boldsymbol{\theta}_1 \mid D, M_1, M_2)\, P(\boldsymbol{\theta}_2 \mid \boldsymbol{\theta}_1, D, M_1, M_2) \\ &\propto \pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_1)\, P(\boldsymbol{\theta}_2 \mid \boldsymbol{\theta}_1, M_1, M_2)\, P(D_{\phi_2} \mid \boldsymbol{\theta}_2, \boldsymbol{\theta}_1, D_{\theta_1}, D_{\phi_1}, M_1, M_2) \end{aligned} \tag{4}$$

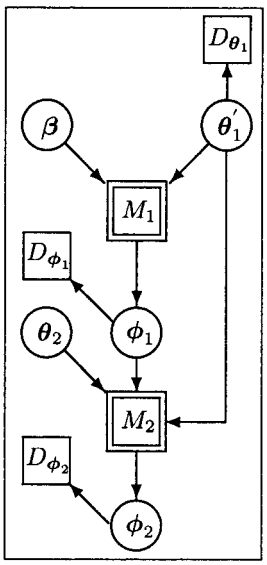

Figure 1. Map of sequential models M_1 and M_2. M_1 maps inputs $\boldsymbol{\theta}_1 = \{\boldsymbol{\theta}_1', \boldsymbol{\beta}\}$ to output $\boldsymbol{\phi}_1$. $\boldsymbol{\theta}_1'$ and $\boldsymbol{\phi}_1$ are inputs to model M_2 along with a distinct set of inputs $\boldsymbol{\theta}_2$

if inputs $\boldsymbol{\theta}_2$ are independent of D_{θ_1} and D_{ϕ_1} given $\boldsymbol{\theta}_1$ and M_1. The first factor in Equation (4) is the posterior distribution of the inputs to M_1 obtained in Equation (1). The second term is the prior for inputs to M_2 given the values of inputs and outputs to the first model. The last term is the conditional likelihood of the second set of observed data given all inputs and outputs. No independence assumptions were placed on the three data sets in deriving Equation (4) but, if conditional independence is assumed, Equation (4) may be simplified.

SIR can be used to sample from this posterior. Assuming that the conditional prior for $\boldsymbol{\theta}_2$ is easy to sample from, SIR entails sampling from $\pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_1)$, and then sampling from $P(\boldsymbol{\theta}_2 \mid \boldsymbol{\theta}_1, M_1, M_2)$ conditional on the $\boldsymbol{\theta}_1$. Resampling with weights $P(D_{\phi_2} \mid \boldsymbol{\theta}_2, \boldsymbol{\theta}_1, D_{\theta_1}, D_{\phi_1}, M_1, M_2)$ yields an approximate sample from the desired distribution. We note that the analytic form of $\pi^{[\boldsymbol{\theta}_1]}(\boldsymbol{\theta}_1)$ is not needed. It is sufficient to resample the posterior sample already obtained.

In Section 4 we will present an application to a two-compartment model for predicting concentration of a contaminant in vine produce (e.g. tomatoes) which has the form shown in Figure 1. One of the compartments in the model represents the contribution to produce concentration from the concentration in soil. The concentration in soil can itself be predicted by another deterministic model.

A distribution representing uncertainty in the output of a multicompartment model can be obtained using the method of Section 2.1, or if there are other models for predicting each compartment then the sequential approach of Section 2.2 may be used.

3. APPLICATION TO A ONE-COMPARTMENT MODEL: SOIL CONCENTRATION OF PCBS

3.1. Background

A severely contaminated region of New Bedford Harbor, MA (NBH) was designated as a Superfund site in 1982 due to a high concentration of polychlorinated biphenyls (PCBs), a probable human

carcinogen (U.S. EPA, 1997b) in marine sediments. The PCBs were introduced to the harbor in waste dumped between the 1940s and the late 1970s by local industry. PCBs in contaminated sediment or water can travel to other media, such as air, soil or food. PCBs may also be present in household electrical appliances. Dredging of the severely contaminated region began in 1994 and proceeded as part of a cleanup. Dredging disturbs the contaminated sediment, which may increase the transport of PCBs to air via volatilization or to water with tidal flushing.

In order to estimate the incremental increase in the annual risk of cancer due to exposure to PCBs, one requires an estimate of the average level of PCB a person could be exposed to in a year. For people with no occupational exposure to PCBs, indirect exposure pathways such as contaminated soil, dust or food may contribute significantly to overall exposure. An important intermediate step in determining the level of exposure through indirect pathways is estimating the average concentration of PCB in soil.

Taylor (1992) adapted a deterministic single compartment steady-state model for predicting contaminant concentration in soil from the concentration in air from its original form in Fries and Paustenbach (1990). The model has a single output, C_s, the contaminant concentration in soil (mg/g), and five inputs, and is given by

$$C_s = \frac{C_a V_d}{b \rho D} \tag{5}$$

Descriptions of the inputs are given in Table 1. The numerator in Equation (5) represents the mass flux (deposition) of the contaminant from air onto soil. The denominator represents the dilution and decay of the contaminant in the soil (Cullen, 2002).

Researchers have measured PCB concentration in 19 air samples and 18 soil samples taken from a site upwind of the harbor in Dartmouth, MA. The method of data collection and protocols followed are described in Vorhees *et al.* (1997) and Vorhees *et al.* (1999). PCB concentrations in these upwind samples do not reflect dredging induced fluxes related to the NBH cleanup which would invalidate the use of a steady-state model.

Hereafter PCB concentration refers to the sum of the 59 prevalent, persistent and toxic PCB congeners for which the samples were analyzed. Individual congener information can be obtained from Table 2 of Cullen *et al.* (1996).

3.2. Model specification

PCB concentrations in air and soil vary in space and time, and it is their marginal distribution that is of primary interest in this application. We therefore reformulate the model given in Equation (5) in terms of the marginal mean and variance of logged average air and soil PCB concentrations. Under the

Table 1. Description of inputs to the soil concentration model, Equation (5)

Input	Description	Unit
C_a	PCB concentration in air	mg/m^3
V_d	Deposition velocity	m/d
b	Decay constant in soil	1/d
ρ	Density of the soil	g/m^3
D	Mixing depth of soil	m

Table 2. Priors on inputs to the soil concentration model, Equation (5)

Input	Prior	Input	Prior
b	Uniform(0.0001, 0.0002)	μ_g	Normal(-13.67, 1.54)[§]
ρ	$\text{Log}_e\text{Normal}(1.4 \times 10^6, 1.15)$[‖]	σ_g^2	Inv-$\chi^2(1, 0.19) \mid V_t^A \geq 0$[¶]
D	Uniform(0.15, 0.25)	μ_d	Normal(-15.49, 1.54)
p	Dirichlet($20 \times (0.06, 0.01, 0.9)$)	σ_d^2	Inv-$\chi^2(1, 1.11) \mid V_t^A \geq 0$
V_d for:		V_s^A	Inv-$\chi^2(1, 0.0024) \mid V_t^A \geq 0$
Coarse particles	$\text{Log}_{10}\text{Uniform}(10, 1000)$[‡]	V_s^S	Inv-$\chi^2(1, 0.32) \mid (V_s^S > V_s^A)$
Fine particles	$\text{Log}_{10}\text{Uniform}(1, 100)$	γ^A	0.0025
Vapor phase	$\text{Log}_{10}\text{Uniform}(10, 1000)$	γ^S	0.0025

[‖]If $X \sim \text{Log}_e\text{Normal}(a, b)$ then $\ln(X) \sim \text{Normal}(\ln(a), b^2)$.
[‡]If $X \sim \text{Log}_{10}\text{Uniform}(a, b)$ then $\log_{10}(X) \sim \text{Uniform}(\log_{10}(a), \log_{10}(b))$.
[§]If $X \sim \text{Normal}(a, b^2)$ then X has mean a and standard deviation b.
[¶]If $X \sim \text{Inv-}\chi^2(a, b^2)$ then X is a scaled inverse χ^2 with a df and scale b.

lognormal assumptions we will make and justify later, this will be sufficient to give the marginal distributions we seek. We note that there are probably spatial and temporal dependencies in the air and soil concentrations, but they are of little interest for our purposes, and our data are insufficient to estimate them.

Taking an expectation of the logarithm of Equation (5) with respect to space and time we find $\mu_{\ln(C_s)} = \mu_{\ln(C_a)} + \ln(V_d) - \ln(b) - \ln(\rho) - \ln(D)$, where $\mu_{\ln(C_s)}$ and $\mu_{\ln(C_a)}$ are the mean of logged average soil and air PCB concentrations, respectively. Soil PCB concentration, has spatial dependence but no temporal dependence under Equation (5) due to the steady-state assumption. The PCB concentration in air varies over both space and time. Thus $\ln(C_a)$ has an extra component of variability for which we must account.

Assuming additivity on the log scale we can decompose the total uncertainty in soil and air PCB concentration into a spatial variability component, a component for uncertainty due to measurement error and an additional temporal variability component for air concentration. Then, $\sigma^2_{\ln(C_s)} = V_s^S + \gamma^S$ and $\sigma^2_{\ln(C_a)} = V_s^A + V_t^A + \gamma^A$, where $\sigma^2_{\ln(C_s)}$ and $\sigma^2_{\ln(C_a)}$ represent the total uncertainty in $\ln(C_s)$ and $\ln(C_a)$ respectively, V_s^S and V_s^A are the spatial components for soil and air, V_t^A is the temporal component for air, and γ^S and γ^A are the measurement error components.

In terms of the formulation in Section 2.1, we have

$$\begin{aligned}
&\text{Inputs :} \quad \boldsymbol{\theta}_1 = \left(\ln(b), \ln(\rho), \ln(D), \ln(V_d), \mu_{\ln(C_a)}, \sigma_{\ln(C_a)}, V_s^S, V_s^A, \gamma^A, \gamma^S\right) \\
&\text{Outputs :} \quad \boldsymbol{\phi}_1 = \left(\mu_{\ln(C_s)}, \sigma_{\ln(C_s)}\right) \\
&\text{Model :} \quad M: \begin{pmatrix} \mu_{\ln(C_s)} = \mu_{\ln(C_a)} + \ln(V_d) - \ln(b) - \ln(\rho) - \ln(D) \\ \sigma^2_{\ln(C_s)} = V_s^S + \gamma^S \end{pmatrix}
\end{aligned} \tag{6}$$

The measured PCB concentration in air and soil samples are represented by D_{θ_1} and D_{ϕ_1}, respectively. Note that V_t^A is an intermediate output using

$$V_t^A = \sigma^2_{\ln(C_a)} - V_s^A - \gamma^A \tag{7}$$

We assume that the model given by Equation (6) is applicable and do not take into account model uncertainty in our analysis.

3.3. Priors and likelihoods

Table 2 summarizes the priors for each input. We now outline the scientific basis for these choices.

Priors for b, ρ and D were given in Cullen (1995) in an analysis of human exposure to emissions from waste incinerators. These priors were based on appropriate available literature and expert knowledge to adapt them to the site in question. This information is presented in Cullen (1995). They are applicable also to the NBH site.

In Cullen (1995), deposition velocity was initially assigned a mixture of 13 log-uniform distributions reflecting dependence of the velocity at which a contaminant falls to the ground on the size of particle that it adheres to or whether it occurs as a vapor. This work was based on the belief that the behavior of larger particles is well described by Stokes' Law, while the settling of smaller particles is most effectively described by Brownian motion (Sehmel, 1980). Several of the particle classes were later pooled for computational efficiency (Cullen, 2002). We follow the pooled approach and assign as a prior to deposition velocity a mixture of three log-uniform densities corresponding to fine (0.1–1.0 μm) and coarse (1–10 μm) particles and a vapor phase.

Let $\mathbf{p} = (p_c, p_f, p_v)$ be the proportion of total PCB adhered to coarse and fine particles or occurring as a vapor, respectively. A Dirichlet prior was assigned to $\mathbf{p}$ that was consistent with the following ranges for the marginal distributions, $p_c \in [0.0085, 0.13]$, $p_f \in [0.0015, 0.03]$, $p_v \in [0.85, 0.99]$. Experimental measurements on which to base these ranges are only estimates due to difficulties in sampling a large enough volume of air and hence PCB in a sampling period. From Falconer and Bidleman (1994) we derived the proportion of individual congeners adhered to either coarse or fine particles in air of 10°C (50°F) and 25°C (77°F). Ranges for $p_c + p_f$ and p_v for total PCB were derived from these results by using the relative loadings of the individual congeners found in the NBH samples (Vorhees *et al.*, 1997) and by considering both air temperatures. Steinberg *et al.* (1995) also considered relative loadings in deriving priors. Whitby (1978) reports that approximately 15 per cent of the volume of background average air is in fine particles, and this was used to separate the ranges for p_c and p_f.

Vorhees *et al.* (1997) reported that measured air and soil PCB concentrations were ±10 per cent (as assessed in blanks). Assuming a normal distribution on the logged errors, the measurement error variances γ^A and γ^S were both set equal to 0.0025.

We have little prior information about V_s^A and V_s^S. One common approach to Bayesian inference in this situation is to use flat priors, i.e. priors that assign a uniform distribution to the parameter or to a function of it, over all possible values. Such a distribution integrates to infinity rather than to 1 so is not a true probability distribution; it is called an *improper prior*. Such priors often lead to problems and paradoxes.

Instead, we use a *unit information prior*, i.e. a prior that contains the same amount of information as one typical observation. Such priors are proper. However, they are diffuse relative to the likelihood and hence tend to have little effect on inference (Kass and Wasserman, 1995; Raftery, 1999), thus representing well the situation where there is little prior information.

The spatial variability in soil should be greater than that in air. Hence we impose the condition that $V_s^S > V_s^A$ as part of the prior specification. Due to this dependence of V_s^S on V_s^A, the intermediate output V_t^A given in Equation (7) now plays a role in the calculation of the output $\sigma_{\ln(C_s)}$. This quantity,

V_t^A, represents the temporal component to variability in air PCB concentration and must be non-negative. This introduces constraints on the priors for both $\sigma_{\ln(C_a)}$ and V_s^A.

The observed PCB concentrations in air are 24 h average measurements. The inputs $\mu_{\ln(C_a)}$ and $\sigma_{\ln(C_a)}$ are the mean and standard deviation of annual average air PCB concentration. We must account for this disparity in time scale.

Let μ_g and σ_g be the mean and standard deviation of 24 h average logged air PCB concentration in the growing season for vegetables (June–September). Let μ_d and σ_d be the corresponding values in the dormant season. We weight the growing and dormant season 24 h average air concentrations according to the proportion of the year in which vegetables grow in this region, to develop expressions for the mean and variance of year-round 24 h average logged air concentration. This mean and variance are on the same scale as the observed measurements, are denoted by μ_a, σ_a^2, and are given by the following expressions:

$$\mu_a = \frac{1}{365}(122\mu_g + 243\mu_d) \qquad \sigma_a^2 = \frac{1}{365^2}\left((122\sigma_g)^2 + (243\sigma_d)^2\right) \tag{8}$$

The quantities μ_g, μ_d and σ_g, σ_d are given diffuse unit information priors estimated using the observed (24 h average) concentrations. The form of each prior represents the information that a single observation would provide and is given in Table 2.

Pollutant concentrations are believed to follow a lognormal distribution. Ott (1990) justified this distribution in the general case using a dilution argument with the central limit theorem to show that the product of independent random variables has a lognormal distribution.

Assuming that the observed growing season and dormant season air data are independent and identically distributed, the combined likelihood for the observed air data given the inputs is

$$L_{inp}(\boldsymbol{\theta}_1) = \prod_{i=1}^{11} \mathrm{N}\left(Y_i^d; \mu_d, \sigma_d^2\right) \prod_{i=1}^{8} \mathrm{N}\left(Y_i^g; \mu_g, \sigma_g^2\right)$$

where Y_i^g and Y_i^d are the ith logged observed 24 h average air measurements in the growing and dormant seasons, respectively.

In the NBH region, weather is believed to occur in five-day cycles (Leith, 1973). It follows that there are approximately 73 of these cycles in a year. We follow Cullen and Frey (1999) in assuming that the expected annual average air concentration equals the expected 24 h average air concentration but variability in annual average concentration is smaller than the variability in 24 h average concentration, due to the factor of 73. Under these assumptions and using properties of the lognormal distribution, the 24 h average values can be converted to the annual average parameters $\mu_{\ln(C_a)}$ and $\sigma_{\ln(C_a)}$ and these used in model (6). The formulae for this conversion are as follows:

$$\mu_{\ln(C_a)} = \mu_a + 0.5\,\sigma_a^2 - 0.5\ln\left(1 + \frac{1}{73}\left(\exp\{\sigma_a^2\} - 1\right)\right) \tag{9}$$

$$\sigma_{\ln(C_a)}^2 = \ln\left(1 + \frac{1}{73}\left(\exp\{\sigma_a^2\} - 1\right)\right) \tag{10}$$

The observed PCB concentrations in soil reflect long-term averaging. The outputs, $\mu_{\ln(C_s)}$ and $\sigma_{\ln(C_s)}$ are on this long-term scale also. Assuming that the observed soil data are independent and

identically distributed, the likelihood for the observed soil data given the outputs and model is

$$L_{\text{out}}(\boldsymbol{\phi}_1) = \prod_{i=1}^{18} \text{N}\left(Y_{\text{i}}^{\text{S}};\, \mu_{\ln(C_{\text{s}})}, \sigma^2_{\ln(C_{\text{s}})}\right)$$

where Y_{i}^{S} is the ith logged observed soil measurement.

The assumption that the observed data are independent may not be correct, but the data are too few to provide evidence of spatial or temporal dependence, or of the form that it would take. Our modeling of the variation in air concentration in terms of five-day cycles is an effort to take account of temporal autocorrelation.

3.4. Results

Unless previously stated, prior independence of the inputs was assumed. An initial sample of size 100 000 was generated from the priors in the first step of the SIR algorithm, and in the resampling step of the SIR algorithm, a sample of size 10 000 was used. This produced approximately 550 unique values in the marginal posterior distributions for each element of $\boldsymbol{\theta}_1$ and $\boldsymbol{\phi}_1$. This was enough to provide a reasonable estimate of the posterior distributions of interest.

Figures 2(a)–(d) show estimates of the posterior distributions of mixing depth, deposition velocity and the parameters of the distribution for $\ln(C_{\text{a}})$, based on the sampled values. These were calculated by non-parametric kernel density estimation from the set of sampled values. These are overlaid on estimates of the prior densities. Both Monte Carlo variability and edge effects are evident in the

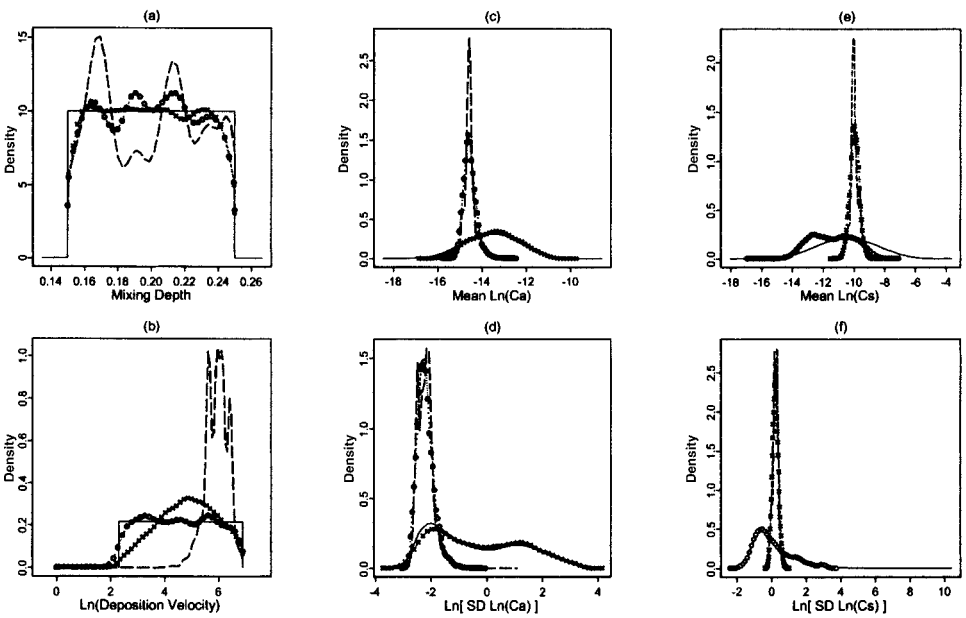

Figure 2. (a), (b): Analytic form of the prior (solid line) and non-parametric density estimates of the posterior distributions (symbols) for soil model inputs, D and V_{d}. (c), (d): Non-parametric density estimates of the prior (solid line) and posterior distributions (symbols) for the air parameters, $\mu_{\ln(C_{\text{a}})}$ and $\sigma_{\ln(C_{\text{a}})}$. (e), (f): Non-parametric density estimates of the induced (solid line) and posterior distributions (symbols) for the soil model outputs. Posterior distributions using all data (dashed), only air data ($\circ$) and only soil data ($\times$) are shown

estimates. Inspection of these plots can reveal whether the observed data in the form of likelihoods provided information on the inputs or outputs.

The posterior distribution of mixing depth shows little updating over the prior, suggesting that the observed data did not contain much additional information. There is a large amount of updating in Figure 2(b). This suggests that under this model the observed data are consistent with faster deposition velocities. This corresponds to PCB adhering to coarse particles or occurring in the vapor phase rather than adhering to fine particles. The inputs $\mu_{\ln(C_a)}$ and $\sigma_{\ln(C_a)}$ were given diffuse priors, so a large amount of updating was expected and is observed. Lower values of variability in air are favored by the data and the model.

Figures 2(e), (f) show estimates of the induced and posterior distributions for the outputs $\mu_{\ln(C_s)}$ and $\sigma_{\ln(C_s)}$. Again we expect, and observe, a large amount of updating due to the effect of diffuse priors.

The posteriors just discussed represented the additional information on $\boldsymbol{\theta}_1$ and $\boldsymbol{\phi}_1$ provided by the combined effect of the two likelihoods $L_{\text{inp}}(\boldsymbol{\theta}_1)$ and $L_{\text{out}}(\boldsymbol{\phi}_1)$. Our uncertainty about the true values of the inputs and outputs was reduced in most cases by this additional information. We may be interested to know which set of data is providing more of the information, i.e. which set is reducing our uncertainty for a particular input or output most?

The effect of the air data alone can be seen by examining the posterior distributions based on only the air likelihood. The SIR algorithm proceeds as before, but in this case the resampling is performed with weights dependent on only $L_{\text{inp}}(\boldsymbol{\theta}_1)$. The SIR weight associated with $\boldsymbol{\theta}_{1_i}$ and $\boldsymbol{\phi}_{1_i}$ is then $w_i \propto L_{\text{inp}}(\boldsymbol{\theta}_{1_i})$ rather than those given in (3). A similar approach is used to examine the relative effect of the soil data alone.

Figures 2(c), (d) show that it is the air data which contain the most information on the air parameters. This is because the posterior based on only the soil data shows almost no change from the prior distribution. Neither data set contains much information on mixing depth and it is the soil data that are consistent with the faster deposition velocities. In Figures 2(e), (f) it is the soil data that reduce the uncertainty about the soil parameters $\mu_{\ln(C_s)}$ and $\sigma_{\ln(C_s)}$.

Sensitivity to the prior distributions placed on the model inputs was investigated. In most cases, the posterior results were relatively insensitive to reasonable changes in the prior distributions. However, we found sensitivity exceeding that due to Monte Carlo variability in the prior distributions, in the posterior mean of $\mu_{\ln(C_s)}$ to perturbations in the prior median of ρ. While sensitivity to priors is generally not desired, we expect it here because our priors represent the available expert knowledge. The perturbed prior median which caused the sensitivity would be excluded by this available scientific knowledge.

3.5. Comparison to other methods

Monte Carlo, Bayesian and other probabilistic methods are relatively new in environmental risk assessment. Point estimates have traditionally been used to estimate the inputs and outputs of a deterministic model. A Monte Carlo (MC) approach pushes values of the inputs simulated from the prior distribution through the model, resulting in a distribution for the model outputs representing uncertainty (U.S. EPA, 1997a; Smith, 1994; Thompson *et al.*, 1992). The current state-of-the-art two-dimensional Monte Carlo approach (Hoffman and Hammonds, 1994; McKone, 1994; Cullen and Frey, 1999) goes further by propagating inputs in such a way as to separate the effects of uncertainty from lack of knowledge from those of variability due to natural heterogeneity in the population of interest or across time and space. In its current formulation, our Bayesian approach does not separate uncertainty and variability. The main difference between these MC approaches and the Bayesian one described here is that the MC approaches do not take account of the likelihood derived from the data on outputs, and in particular do not update our knowledge of the inputs to take account of the likelihoods.

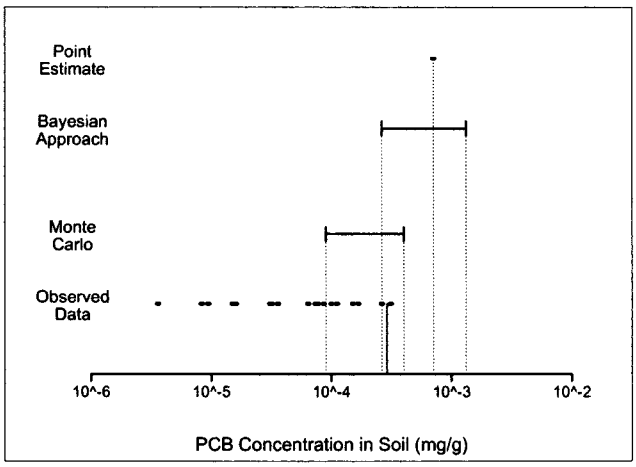

Figure 3. A point estimate of the 97.5th percentile of PCB concentration in soil as compared to 90 per cent intervals for this percentile based on the two-dimensional Monte Carlo and Bayesian approaches. The observed data are overlaid. The observed 97.5th percentile is indicated by the solid vertical line

A distribution for percentiles of soil concentration developed under a standard MC approach would be equivalent to the distribution of induced values obtained in the SIR algorithm if data-based distributions rather than diffuse distributions were used for all model inputs.

Cullen (2002) applied a two-dimensional Monte Carlo approach to Equation (5) and developed a distribution for the 97.5th percentile of soil concentration. This distribution is shown in Figure 1(d) of Cullen (2002). This MC distribution provides a 90 per cent credible interval for the 97.5th percentile of C_s. Similarly, the posterior sample can be used to develop a Bayesian approach 90 per cent credible interval for the 97.5th percentile of C_s.

Figure 3 represents a comparison of the estimates of the 97.5th percentile for C_s that would be obtained using point estimate, two-dimensional MC and Bayesian approaches. The point estimate for soil concentration of PCB is based on the 97.5th percentiles of V_d and a lognormal distribution fitted to the observed air data, and the 2.5th percentiles for b, ρ and D (Cullen, 1994). The observed data are overlaid and the observed 97.5th percentile indicated by the solid vertical line.

The Bayesian interval is more health conservative than the Monte Carlo one, reflecting the information contained in the soil likelihood and the effect of not separating uncertainty and variability. The point estimate of the 97.5th percentile of soil PCB concentration is 7.1×10^{-4} mg/g, while the Bayesian approach would estimate this to be 1.3×10^{-3} mg/g (using the upper endpoint of the interval). Thus in this case the Bayesian estimate of the upper endpoint is more health conservative than the point estimate. This ordering will not necessarily always hold.

4. APPLICATION TO A MULTICOMPARTMENT MODEL: PRODUCE CONCENTRATION OF PCBS

We now illustrate the methodology for sequential multicompartment models developed in Section 2.2. We use a two-compartment model for predicting contaminant concentration in produce from the concentration in air and soil.

Table 3. Description of inputs to the produce concentration model, Equation (11)

Inputs	Description	Unit
C_s	PCB concentration in soil	mg/g
C_{a_g}	PCB concentration in air	mg/m^3
V_d	Deposition velocity	m/d
T_g	Duration of growing season	d
I_f	Interception factor	—
W_c	Weathering rate constant	1/d
Y	Yield of crop	g/m^2
U_f	Uptake factor	—

The first compartment is the soil model already described in Section 3. The second compartment is a deterministic steady-state model adapted by Taylor (1992). The model has a single output C_p, the average contaminant concentration in produce after the growing season (mg/g), and is given by

$$C_p = (C_{a_g} V_d) \frac{(1 - \exp\{-W_c T_g\}) I_f}{W_c Y} + U_f C_s \tag{11}$$

Descriptions of the inputs are given in Table 3. The first term in Equation (11) represents the contribution from air, and the second term, the uptake from soil. Produce are exposed to air only during their growing season, so the model uses C_{a_g}, the average PCB concentration in air during the growing season, rather than an annual average PCB concentration in air. The quantities V_d and C_s are the same as in the soil model given by Equation (5).

We are interested in developing a distribution for the concentration of PCB in vine produce (e.g. tomatoes). We want to take advantage of the information contained in the growing season air and soil likelihoods as well as expert knowledge about the inputs contained in the model for predicting soil concentration. Using the formulation of Section 2.2 and Figure 1 we have the following inputs and outputs to the produce model:

$$\begin{aligned}
\text{Inputs :} \quad & \boldsymbol{\theta}_1' = (\mu_g, \sigma_g, V_d) \\
& \boldsymbol{\phi}_1 = \left(\mu_{\ln(C_s)}, \sigma_{\ln(C_s)}\right) \\
& \boldsymbol{\theta}_2 = (C_{a_g}, C_s, T_g, I_f, W_c, Y, U_f) \\
\text{Outputs :} \quad & \boldsymbol{\phi}_2 = C_p \\
\text{Model :} \quad & M_2 : C_p = (C_{a_g} V_d)(1 - \exp\{-W_c T_g\}) I_f/(W_c Y) + U_f C_s
\end{aligned}$$

The inputs to the soil model that are not direct inputs to the produce model are $\boldsymbol{\beta} = (\ln(b), \ln(\rho), \ln(D), \mu_d, \sigma_d, V_s^S, V_s^A, \gamma^A, \gamma^S)$. The sub-model M_1 is the model given in Equation (6).

In Section 3.4 we obtained an approximate sample from the posterior distribution of $\boldsymbol{\theta}_1 = (\boldsymbol{\theta}_1', \boldsymbol{\beta})$ and therefore for $\boldsymbol{\theta}_1'$ and $\boldsymbol{\phi}_1$. Table 4 details $P(\boldsymbol{\theta}_2 \mid \boldsymbol{\theta}_1, M_1, M_2)$, the prior for inputs to the produce model conditional on the values of soil model inputs and outputs just sampled. The priors for T_g, I_f, W_c, Y and U_f were given in Cullen (1995) along with the basis for these priors.

The lognormal priors placed on C_s and C_{a_g} reflect the prior belief of lognormally distributed pollutant concentrations. C_{a_g} is a seasonal average, so we must transform the 24 h logged average

Table 4. Priors on inputs $\boldsymbol{\theta}_2$ to the produce concentration model, Equation (11)

Input	Prior
T_g	Uniform(90, 110)
I_f	Uniform(0.05, 0.25)
W_c	$\text{Log}_e\text{Uniform}(0.01, 0.1)$[§]
Y	Uniform(5000, 15000)
U_f	Uniform(0.05, 0.15)
$C_{a_g}\|(\mu_g, \sigma_g, T_g)$	$\text{Log}_e\text{Normal}(\mu_{\ln(C_a)_g}, \sigma_{\ln(C_a)_g})$[‡]
$C_s\|(\mu_{\ln(C_s)}, \sigma_{\ln(C_s)})$	$\text{Log}_e\text{Normal}(\mu_{\ln(C_s)}, \sigma_{\ln(C_s)})$

[§]If $X \sim \text{Log}_e\text{Uniform}(a, b)$ then $\ln(X) \sim \text{Uniform}(\ln(a), \ln(b))$.
[‡]If $X \sim \text{Log}_e\text{Normal}(a, b)$ then $\ln(X) \sim \text{Normal}(\ln(a), b^2)$.

parameters μ_g and σ_g obtained from the soil model posterior to seasonal average parameters $\mu_{\ln(C_a)_g}$ and $\sigma_{\ln(C_a)_g}$. In a growing season of length T_g days there are $T_g/5$ five-day weather cycles, so such a conversion is done using equations similar to Equations (9) and (10) with variability due to the factor of $T_g/5$ rather than 73.

In addition to the soil and air samples collected in the NBH region, researchers measured PCB concentration in six samples of vine produce. These measurements and the model output represent total PCB concentration in a produce sample at the time of harvest. Assuming that the observed produce data are independent and identically distributed, the likelihood for the observed plant data given the output $\boldsymbol{\phi}_2 = C_p$ is

$$L_{\text{out}}(\boldsymbol{\phi}_2) = \text{Log}_e\text{Normal}(\boldsymbol{\phi}_2; \bar{x}_p, s_p)$$

where $\bar{x}_p$ and s_p are the mean and standard deviation of the six logged observed vine produce concentrations. The vine produce measurements are represented by D_{ϕ_2} in Figure 1.

An initial sample of size 20 000 was resampled from the posterior for $\boldsymbol{\theta}_1'$ and $\boldsymbol{\phi}_1$ developed in Section 3.4. An initial sample of equal size from the priors in Table 4 was then taken.

Unless previously stated, we assume independence in sampling between $\boldsymbol{\theta}_2$ and $\boldsymbol{\theta}_1$, and within $\boldsymbol{\theta}_2$. However, the a posteriori correlation between members of $\boldsymbol{\theta}_1'$ and $\boldsymbol{\phi}_1$, induced by the soil model and the air and soil data, is retained. A resample of size 2000 was then taken using the density of the induced plant concentrations under the likelihood as the SIR weights. This produced a sample from the posterior for inputs to the vine produce model with approximately 400 unique values.

The distribution induced by the priors overpredicts vine produce PCB concentration. This can be seen in Figure 4(d) in the lower values of C_p favored by the vine produce likelihood. No diffuse priors were used in producing this posterior, so the large reduction in overall uncertainty was not expected and is due to the information contained in the observed produce measurements. The overprediction is further seen in the posterior distributions for $\mu_{\ln(C_s)}$ and Y in Figures 4(a) and (c). Higher values of Y and lower values of $\mu_{\ln(C_s)}$ are consistent with lower induced produce concentrations. The priors in Figures 4(a), (b) consist of a resample from the posterior in Figures 2(e) and (b), respectively. The produce data provide little additional information on deposition velocity.

Sensitivity of the posterior results to the prior distributions placed on the model inputs was investigated. No sensitivity exceeded that due to Monte Carlo variability in the prior distributions, and details are omitted.

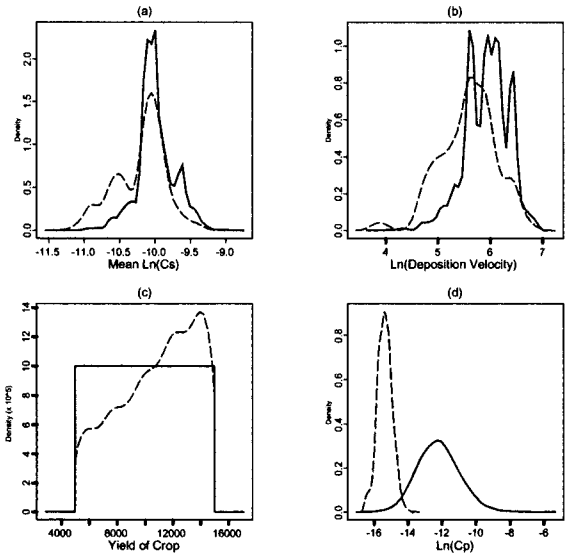

Figure 4. (a), (b): Non-parametric density estimates of the prior (solid line) and posterior distributions (dashed) for vine model inputs, $\mu_{\ln(C_s)}$ and V_d. (c): Analytic form of the prior (solid line) and non-parametric density estimate of the posterior distribution (dashed) for vine model input Y. (d): Non-parametric density estimates of the induced (solid line) and posterior distributions (dashed) for $\ln(C_p)$

5. DISCUSSION

We have presented a Bayesian approach to dealing with uncertainty in inputs and outputs of deterministic models when data on the inputs and/or the outputs is available. This takes full account of evidence and uncertainty about all the model inputs and outputs when making inference about quantities of policy or scientific interest which are functions of inputs, outputs or both. We have extended the approach to sequential multicompartment models of a kind that are common in environmental risk assessment.

The contribution of individual likelihoods to reducing uncertainty was assessed graphically in Section 3.4. This could provide valuable information on where to concentrate future data collection resources.

The Bayesian approach was compared against more traditional methods. In contrast with these methods, the Bayesian approach takes account of the observed data on outputs, and uses all observed data to update knowledge on model inputs. A choice of whether to use the Bayesian approach or Monte Carlo approach for a particular application should consider whether data on outputs are available, whether updating knowledge on the model inputs using the observed data is desired, and the amount of computer time required (the Bayesian approach will be more computationally expensive in general). The Bayesian and Monte Carlo approaches to the example of Section 3 took 140 and 20 CPU seconds, respectively. Both approaches were coded in C and run on a DEC Alpha machine.

In our work, we have taken account of prior information about inputs, but often in applications there is also prior information about outputs which could be expressed as a prior on outputs. If there is a prior on model output in addition to a likelihood, then Poole and Raftery (2000) showed how the

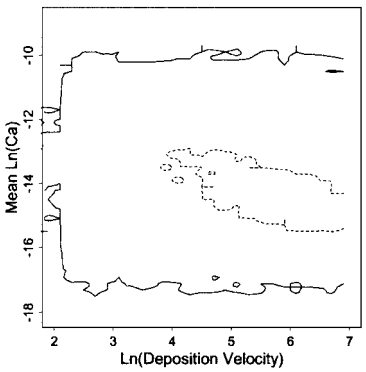

Figure 5. Solid lines are grid-based contours representing the space covered by the sample from the joint prior distribution of $(\ln(V_d),\ \mu_{\ln(C_a)})$. The dashed contours represent the space covered by the SIR sample from the corresponding posterior distribution. Both contours contain approximately 86 per cent of the non-zero mass grid points (40×40 grid). The a posteriori correlation is -0.59

present Bayesian melding approach can be extended to incorporate this additional source of information.

In related work, Draper *et al.* (1999) presented a Bayesian method for uncertainty analysis, including model uncertainty, with an application to risk assessment at a waste incineration site. In their application they had no data on which to base likelihoods or update the prior information. Figures 2 and 4 illustrated the reduction in uncertainty arising from the use of likelihoods in our application. Unlike the approach of Draper *et al.* (1999), our approach does not consider model uncertainty.

In the first application we sampled from the prior on inputs, $p(\boldsymbol{\theta}_1)$, under an assumption of independence. If inputs to a model are correlated a priori, sampling can be adapted to account for this. However, prior correlations should not be confused with the expectation that the data will induce correlations between inputs. Our approach often yields high correlations between inputs a posteriori, even when their prior distributions are independent. An example of this is given in Figure 5. The solid contours represent the region covered by approximately 86 per cent of the sample from the joint priors for $\ln(V_d)$ and $\mu_{\ln(C_a)}$. No a priori correlation was assumed in this sampling. Superimposed are dashed contours representing the region covered by approximately 86 per cent of the posterior sample. We see a posterior correlation of -0.59 induced by the model and the observed data. These a posteriori correlations can provide useful information in a multicompartment model setting. Priors based on independence will typically be more spread out than dependent priors, and will tend to cover the area covered by dependent priors fairly well. Thus it could be argued that prior independence provides a form of prior robustness, regardless of the extent to which it is scientifically justified.

The sequential approach described in Section 2.2 makes full use of all relevant expert knowledge, observed data and model structure, in developing a posterior distribution for some quantity of interest. Of course caution should be taken in using this method. If the posterior distribution developed from a particular model conflicts with physical intuition or displays behavior which cannot be explained, we must question its use as a prior distribution for inputs to a subsequent model. In addition, the underlying assumptions of the models should be compatible.

If posterior expectations rather than posterior distributions are of interest, there is no need to perform the resampling. The posterior expectations can be calculated using the samples from the prior distributions for inputs and the importance weights.

We note that SIR is only one of the possible ways to generate a sample from a posterior distribution. Other methods include Markov chain Monte Carlo (MCMC), and this could be used in principle to simulate from the posterior distributions here. However, posterior distributions in this kind of problem are often highly dependent, often being concentrated close to a submanifold or ridge in the high-dimensional space. Thus MCMC output will tend to be highly correlated. Designing a successful MCMC algorithm for the present application is an extremely delicate problem, and current generic MCMC methods are unlikely to do well, or even to be feasible.

ACKNOWLEDGEMENT

This research was supported by the Environmental Protection Agency through the National Research Center for Statistics and the Environment through Agreement CR825173-01-0.

REFERENCES

Bogen K, Spear R. 1987. Integrating uncertainty and interindividual variability in environmental risk assessment. *Risk Analysis* **7**(4): 427–436.

Cullen A. 1994. Measures of compounding conservatism in probabilistic risk assessment. *Risk Analysis* **14**(4): 389–393.

Cullen A. 1995. The sensitivity of probabilistic risk assessment results to alternative model structures: a case study of municipal waste incineration. *Journal of Air and Waste Management Association* **45**: 538–546.

Cullen A. 2002. Comparison of model predictions and measurements for PCB concentration in the vicinity of Greater New Bedford superfund site. *Environmental Science and Technology* **36**: 2033–2038.

Cullen A, Frey H. 1999. *Probabilistic Techniques in Exposure Assessment: A Handbook for Addressing Variability and Uncertainty in Models and Inputs*. Plenum Press: New York.

Cullen A, Vorhees DJ, Altshul LM. 1996. Influence of harbor contamination on the level and composition of polychlorinated biphenyls in produce in Greater New Bedford, Massachusetts. *Environmental Science and Technology* **30**(5): 1581–1588.

Draper D, Pereira A, Prado P, Saltelli A, Cheal R, Eguilior S, Mendes B, Tarantola S. 1999. Scenario and parametric uncertainty in {GESAMAC}: a methodological study in nuclear waste disposal risk assessment. *Computer Physics Communications* **117**: 142–155.

Falconer R, Bidleman T. 1994. Vapor pressures and predicted particle/gas distributions of polychlorinated biphenyl congeners as functions of temperature and ortho-chlorine substitution. *Atmospheric Environment* **28**(3): 547–554.

Finkel A. 1990. *Confronting Uncertainty in Risk Management: A Guide for Decision Makers.* Center for Risk Management, Resources for the Future: Washington DC.

Fries G, Paustenbach D. 1990. Evaluation of potential transmission of 2,3,7,8-tetrachlorodibenzo-p-dioxin contaminated incinerator emissions to humans via foods. *Journal of Toxicology and Environmental Health* **29**: 1–43.

Givens GH. 1993. A Bayesian framework and importance sampling methods for synthesizing multiple sources of evidence and uncertainty linked by a complex mechanistic model. *Ph.D. Thesis*, University of Washington, Seattle, WA.

Hesterberg T. 1995. Weighted average importance sampling and defensive mixture distributions. *Technometrics* **37**(2): 185–194.

Hoffman F, Hammonds J. 1994. Propagation of uncertainty in risk assessments: the need to distinguish between uncertainty due to lack of knowledge and uncertainty due to variability. *Risk Analysis* **14**(4): 707–712.

IAEA. 1989. *Evaluating the Reliability of Predictions Made Using Environmental Transfer Models: A Safety Practice*. International Atomic Energy Agency, Safety Series No. 100, Vienna, Austria.

Kass R, Wasserman L. 1995. A reference Bayesian test for nested hypotheses and its relationship to the Schwarz criterion. *Journal of the American Statistical Association* **90**: 928–934.

Leith C. 1973. The standard error of time-average estimates of climatic means. *Journal of Applied Meteorology* **12**: 1066–1069.

McKone T. 1994. Uncertainty and variability in human exposures to soil contaminants through home-grown food: a Monte Carlo assessment. *Risk Analysis* **14**(4): 449–463.

Morgan MG, Henrion M, Morris SC, Amaral DAL. 1985. Uncertainty in risk assessment. *Environmental Science and Technology* **19**(8): 662–667.

NRC. 1983. *Risk Assessment in the Federal Government: Managing the Process*. National Research Council, Committee on the Institutional Means for Assessment of Risks to Public Health. National Academy Press: Washington DC.

Ott W. 1990. A physical explanation of the log-normality of pollutant concentrations. *Journal of Air and Waste Management Association* **14**: 1378–1383.

Poole D, Raftery AE. 2000. Inference for deterministic simulation models: the Bayesian melding approach. *Journal of the American Statistical Association* **95**(452): 1244–1255.

Raftery AE. 1999. Bayes factors and BIC—Comment on 'A critique of the Bayesian information criterion for model selection'. *Sociological Methods and Research* **27**(3): 411–427.
Rubin D. 1988. Using the SIR algorithm to simulate posterior distributions. In *Bayesian Statistics 3*, Bernardo JM, DeGroot MH, Lindley DV, Smith AFM (eds). Oxford University Press: Oxford; 395–402.
Sehmel G. 1980. Particle and gaseous dry deposition: a review. *Atmospheric Environment* **14**: 983–1011.
Smith R. 1994. Use of Monte Carlo simulation for human exposure assessment at a Superfund site. *Risk Analysis* **14**(4): 433–439.
Steinberg L, Reckhow K, Wolpert RL. 1995. A Bayesian model for the fate and transport of polychlorinated biphenyl in the upper Hudson river. *Technical Report No. 37*, National Institute of Statistical Sciences.
Taylor A. 1992. Addressing uncertainty in the estimation of environmental exposure and cancer potency. *Ph.D. Thesis*, Harvard School of Public Health, Boston, MA.
Thompson K, Burmaster D, Crouch E. 1992. Monte Carlo techniques for quantitative uncertainty analysis in public health risk assessments. *Risk Analysis* **12**(1): 53–63.
U.S. EPA. 1997a. Guiding principles for Monte Carlo analysis. *EPA/630/R-07/001*. U.S. Environmental Protection Agency, Risk Assessment Forum: Washington DC.
U.S. EPA. 1997b. IRIS (Integrated Risk Information System): Polychlorinated biphenyls (PCBs). *CASRN 1336-36-3*. U.S. Environmental Protection Agency: Internet.
Vorhees DJ, Cullen A, Altshul LM. 1997. Exposure to polychlorinated biphenyls in residential indoor air and outdoor air near a Superfund site. *Environmental Science and Technology* **31**: 3612–3618.
Vorhees DJ, Cullen A, Altshul LM. 1999. Exposure to polychlorinated biphenyls in house dust and yard soil near a Superfund site. *Environmental Science and Technology* **33**: 2151–2156.
Whitby K. 1978. The physical characteristics of sulfur aerosols. *Atmospheric Environment* **12**: 135–159.

