

# RISK ASSESSMENT WITH SUBJECTIVELY DERIVED DOSES

Raymond J. Carroll \*

June 22, 1998

## Abstract

In situations involving radiation exposure and acute inhalation exposure of chemicals, among others, actual exposure to a target organ is not available. Instead, interdisciplinary teams build complex, often differential equation-based models relating a measure of exposure to actual exposure at a target organ. These complex models are themselves filled with uncertainties, e.g., in transition rates from one site to the next, in environmental measures, etc. The uncertainties then impact analysis of dose-response. Uncertainties can be quantified either subjectively or empirically, but in either case, the result is that the models provide not an actual target-organ dose, but a distribution for this dose. Although many might base analysis on a representative value of the target-organ dose, we argue that the fact that these mechanistic models give a distribution for target-organ dose is easily accommodated in a statistical dose-response analysis.

*Key words and phrases:* Bayesian Analysis; Likelihood Analysis; Maximum Likelihood; Measurement Error; Missing Data; Radiation; Regression; Risk Assessment; Subjective Assessment of Priors; Toxicology.

**Short title.** Subjectively Derived Doses.

---

\*Raymond J. Carroll is University Distinguished Professor, and Professor of Statistics, Nutrition and Toxicology, Department of Statistics, Texas A&M University, College Station, TX 77843-3143. Carroll's research was supported by a grant from the National Cancer Institute (CA-57030). The author thanks Owen Hoffman, Elaine Ron, Jay Lubin and Andre Bouville of the National Cancer Institute for many helpful conversations.

# 1 INTRODUCTION

One aim of radiation studies is to estimate the risk that a specific dose to a target organ (e.g., the thyroid) has on a health outcome (e.g., thyroid cancer). The difficulty with such analyses is that for a variety of reasons, the exact dose to the target organ cannot be determined. The purpose of this article is to suggest an alternative to the common practice of constructing a “representative value” of a target–organ dose.

This article is structured as follows. In this section, we set the basic framework for the problem. Section 3 describes the relationship of our approach with Berkson modeling. Section 2 briefly describes the basic idea of a Bayesian analysis in the context. Section 4 gives details of statistical analysis. Concluding remarks are made in Section 5.

The classical measurement error literature is only tangentially relevant to this problem. The typical idea in this literature is that the target–organ dose is measured with error. Generally, the error is assumed to be either additive or multiplicative, so that either

$$\text{Observed target–organ dose} = \text{True target–organ dose} + \text{error (additive)} ; \quad (1)$$

$$\text{Observed target–organ dose} = \text{True target–organ dose} \times \text{error (multiplicative)} . \quad (2)$$

Within the constraints of such error models, a wide variety of fitting methods are available, see Carroll, et al. (1995) for a recent review. The most popular idea is to replace the true target–organ dose by its best estimate given the observed data, and then apply a standard analysis. Forming a “representative value” in order to apply classical regression models is a technique known in the measurement error model literature as *regression calibration*, and is a method with a long history. The method has been suggested in the statistics literature relevant to radiation problems by Prentice (1982), Pierce, et al. (1992) and Schafer (1992), who call it a “replacement” method because one replaces the true target–organ dose by its best estimate. Other researchers call this method “single imputation”.

Replacement methods are especially well–suited to linear regression, with one caveat described below. In logistic regression, replacement methods work reasonably well as long as relative risks are small to moderate, with a tendency to underestimate risk, sometimes severely, see equation (3.24) of Carroll, et al. (1995). In other models, replacement methods can have drawbacks. In mixed models, replacement methods overestimate the size of variance components (Wang, et al., 1998), sometimes severely. In loglinear models, they tend to overestimate the intercept (Carroll, et al., section 3.9.3), even in simple cases, and with non-constant error structure can even lead to biased

relative risk estimates. Overestimation of the intercept is serious, because it tends to overestimate the beneficial impact of ameliorative strategies. Thus, for example, consider a model in which one believes that the probability of an event is related to target–organ dose by the equation

$$\text{pr}(\text{event}) = \exp \{ \beta_0 + \beta_1 (\text{target–organ dose}) \}.$$

Even if the risk parameter  $\beta_1$  is accurately estimated by a replacement method, if  $\beta_0$  is overestimated, then the total number of events reduced by a change in target–organ dose is overestimated. Other methods of estimation are required to overcome these problems.

It is not the point of this article to argue about the best method of fitting models under highly structured measurement error models such as (1) and (2). Our point is to highlight an emerging class of problems where the use of replacement methods seems unduly restrictive.

In radiation dosimetry and in quantifying environmental exposure (Hoffman, et al., 1995, 1996), the dose that a specified individual actually receives is unknown. In some instances, researchers undertake a sophisticated uncertainty analysis combining observables, subjectively derived distributions and Monte–Carlo simulation to arrive at a *distribution* for the true dose which reflects the uncertainty in the problem. While the dose received by each individual is a single true but unknown value, for each individual (and hence for each individual true but unknown dose) there is a subjective probability distribution that quantitatively expresses the analyst’s state of knowledge about that value. In their Figure 2, Hoffman, et al. (1995) give a particularly nice illustration of this process. The key point here is that the outcome of the process is not the actual target–organ dose for an individual, but the distribution of this dose. In effect, the covariate is a distribution and not a number. To quote Hoffman, et al.,

*Today, Monte–Carlo simulation is the preferred method for error propagation, particularly for complex models whose input parameters are interdependent. Monte–Carlo simulation ... is usually necessary when equations used for dose reconstruction increase in complexity beyond simple additive and multiplicative chains ... Based on the best information available, each uncertain model component is expressed as a subjective distribution of values ... a single value is selected at random for each and every uncertain model component to produce a single estimate of the dose ... This process is repeated numerous times, producing a distribution of many realizations.*

Sometimes, for dose–response modeling, the results of this uncertainty analysis are summed up into a so–called “best estimate” or “best guess” (Hoffman, et al., 1995, pages 110–111) for the target–organ dose (for example, the median of the distribution), which is then used in a classical

regression analysis. We prefer the term “representative value”, which we use throughout. One can argue that the median of the distribution leads to biases in a replacement analysis, and that the mean is theoretically preferable, although in practice the mean may be very much more variable than the median. We argue below that the distribution itself can be used to form a better estimate of dose–response than is obtainable from any single “best guess”.

A similar issue arises in risk assessment for the acute inhalation of chemicals. Simpson, et al. (1996) and Guth, et al. (1997) describe the problem of relating acute inhalation exposure to PERC (tetrachloroethylene) to adverse outcomes. The data are a compilation of various studies in which the animals are exposed in a chamber to a known amount of PERC for a known duration. The analyses done to date model adverse outcome as a response predicted as a linear function of log-dose and log-duration. Such modeling is useful, but there is considerable interest in relating exposure to specific adverse outcomes, e.g., problems with lung functioning. In such a case, it makes sense to relate adverse outcome for a target–organ to exposure at that target–organ. To do this, we require a means of converting a given amount of PERC in the chamber for a given duration to animals of given sex, age and weight to actual target–organ dose. One potentially important way of doing this is through the use of physiologically–based pharmacokinetic (PBPK) modeling. Any PBPK model is a complex mechanistic compartment model which involves many layers of uncertainty, not least of which is in the transition probabilities from compartment to compartment. An excellent statistical description of PBPK modeling for PERC is given by Gelman, *et al.* (1997). Recognizing the impossibility of setting a fixed target–organ dose in the light of all the sources of uncertainty, Gelman, *et al.* use Bayesian techniques and subjective information to arrive at a *distribution* for the target–organ dose. While this distribution can be summed up into a “best guess” for target–organ dose, again we will argue that there is an alternative analysis which directly incorporates the uncertainty in target–organ dose into the dose–response model.

Both examples illustrate cases where the target–organ dose cannot be summarized from data as a number, but instead as a distribution. The situation can be described formally as follows.

- There is a response  $Y$ , a true but unknown target–organ dose  $X$ , and additional variables (called covariates here)  $Z$  such as gender, age, etc. In the PERC example, concentration and duration of exposure would be included in the covariates.
- Interest lies in establishing a dose–response relationship between the response and target–organ dose, which is summarized via a statistical model with probability mass or density function  $f(y|x, z, \beta)$ . For example,  $\beta$  might be the coefficients in a logistic regression model.

- Target–organ dose cannot be observed, but information about its distribution conditional on the covariates can be obtained via simulation, subjective analysis, etc. In what follows, we denote the probability mass or density function of this distribution by  $f(x|z)$ . We call this the distribution of the target–organ dose, or DTOD for short.

If the DTOD were known in analytical form (but not the actual doses, only their distribution), then a simple analysis can be performed. This is discussed in Section 4, where the essential statistical likelihood functions are summarized in equations (7) and (8).

## 2 BAYESIAN ANALYSIS OF RISK WHEN THE DISTRIBUTION OF TARGET–ORGAN DOSE HAS NO KNOWN FUNCTIONAL FORM

In practice, of course, it is more likely that the DTOD will not be known in analytical form. In the Bayesian framework, we discuss a likely scenario which leads to a simple albeit computationally complex analysis. More formally, we assume that one can sample from the distribution of an individual target–organ dose given a set of covariates. Building models for which this distribution is either known analytically or can be sampled from is a nontrivial exercise in practice, but it can be done and it is facilitated by the fact that the necessary distribution involves only target–organ dose and the covariates but not the response. Gelman, *et al.* (1997) describe a substantial formal Bayesian analysis of PERC to which these ideas are immediately applicable. Somewhat less directly Bayesian is a hypothetical case described by Hoffman, *et al.* (1995). While theirs is not a formal Bayesian analysis, it is very close to being one, relying as it does on subjective decisions and outside data.

In this instance, one can draw from a computer Monte–Carlo samples from the (subjective) distribution of each individual’s target–organ dose given only the covariates. This is for example what is envisaged by Hoffman, *et al.*, *i.e.*, for each individual one can generate as many observations as one wishes from the (subjective) distribution of exposures. It is important here to note that the DTOD itself need not be known analytically in a functional form, but only that one can sample from it. This is a nontrivial consideration, because in many examples the underlying mechanistic model will be so complex that to specify the DTOD itself in a functional form is impossible.

Bayesian analysis under this first scenario is simple in principle, based on combining Gibbs sampling with the Metropolis–hastings algorithm. The basic idea is to use the computer to repeatedly sample both  $\beta$  the true target–organ dose’s of all individuals. One can think of this as

first “filling-in” the true target-organ dose’s via Monte-Carlo simulation, then using a standard analysis to estimate  $\beta$ , then again fill-in the true target-organ dose’s, then get a new estimate of  $\beta$ , etc. At the end, the mean of all the sampled  $\beta$ ’s is an estimate of the posterior mean of  $\beta$  given all the observed data. The key point here though is how one performs the “filling-in” step. One might think that this is just using the DTOD, but this is not quite the case. The details of such an implementation are given in section 4.3.

### 3 RELATIONSHIP WITH THE BERKSON MODEL

Except for the fact the the DTOD is derived from a weight-of-evidence approach, the problem as we have described it is essentially a Berkson-type model. It is not a Berkson model in the usual sense of a mechanism which is set to a nominal value with the true value varying around this nominal value, but in a strictly mathematical sense it can be written in such a way as to correspond to the usual Berkson mathematics. The true target-organ dose is  $X$ , and let the mean of the DTOD be  $W$ , the “representative value”. Then, by definition,

$$X = W + \text{Berkson error.} \tag{3}$$

The distribution of the Berkson error is just the DTOD centered to have mean zero.

Thus, while formally the problem we have described is of Berkson-type, there is one important difference with the usual Berkson literature. The difference is that the Berkson literature usually assumes that the DTOD is known analytically (even more usually that it is a normal distribution), while we recognize that the DTOD is unknown, and one can at best only sample from it using Monte-Carlo.

## 4 DETAILS OF STATISTICAL ANALYSIS

### 4.1 The Basic Statistical Approach

In this subsection, our goal is to describe the mathematical formulation of the problem. *We emphasize that through equation (8), we are pretending that the DTOD is known analytically. After that, and specifically in subsection 4.2, we discuss an approach that takes into account the fact the this distribution is not known analytically.*

Let us recall the basic statistical setup, which consists of the following components.

- A dose-response model, relating an individual’s health response  $Y$  to the target-organ dose  $X$  as well as other predictors  $Z$ . Generally these models have a parameter,  $\beta$ , and can be

summarized in the form of a probability density or mass function  $f(y|x, z, \beta)$ .

- A distribution of target–organ doses, which we call the DTOD. In this section, we assume that (given the covariates) the distributions of the target–organ doses are independent across individuals. Exceptions to this are described in section 4.4.

Since the true target–organ doses are unknown, the question is how to use the DTOD to estimate the dose–response model.

An example might help to understand what is going on. Suppose that  $X$  merely represents whether a person gets a positive dose ( $X = 1$ ) or not ( $X = 0$ ). If a person gets a positive dose, no matter what their covariates are, the probability of an adverse outcome is  $p_{A|EX}$ , while those with no dose have an adverse outcome probability of  $p_{A|NE}$ . If  $X$  were known, the dose–response model (namely the parameters  $p_{A|EX}$  and  $p_{A|NE}$ ) would be estimated as the fraction with adverse health outcomes among those with and without positive doses, respectively.

Here is where the DTOD comes into play. In our framework, every person has a unique set of covariates  $Z$ . For each person, this subjective probability of a positive dose is  $p_{EX|Z}$ , i.e.,  $\text{pr}(X = 1|Z) = p_{EX|Z}$ .

We now seek to link the dose–response model to the DTOD. The only data we actually have is the response  $Y$  and the covariates  $Z$ , and thus we can only compute the probability of an adverse health outcome for each individual given that individual’s covariates. Specifically, by elementary probability,

$$\begin{aligned}
 \text{pr}(Y = 1|Z) &= \text{pr}(\text{adverse response given covariates } Z) \\
 &= \text{pr}(Y = 1, X = 0|Z) + \text{pr}(Y = 1, X = 1|Z) \\
 &= \text{pr}(Y = 1|X = 0, Z)\text{pr}(X = 0|Z) + \text{pr}(Y = 1|X = 1, Z)\text{pr}(X = 1|Z) \\
 &= \text{pr}(Y = 1|X = 0, Z)(1 - p_{EX|Z}) + \text{pr}(Y = 1|X = 1, Z)p_{EX|Z} \\
 &= p_{A|NE}(1 - p_{EX|Z}) + p_{A|EX}p_{EX|Z}.
 \end{aligned} \tag{4}$$

If the (subjective) probability of exposure  $p_{EX|Z}$  were known, it is generally a simple matter to use the observed data to obtain the dose–response function.

For example, suppose that  $Z$  itself takes on only two values, say male and female. Then in a sample, the observed data of  $Y$ ’s and  $Z$ ’s for a  $2 \times 2$  contingency table, as illustrated in Table 1. In this table, a sample of size  $n$  is split up into four cell counts,  $n_{00}, n_{10}, n_{01}, n_{11}$  corresponding to the possibilities (non–adverse outcome, male), (adverse outcome, male), (non–adverse outcome,

	Not Adverse ( $Y = 0$ )	Adverse $Y = 1$
Male, $Z = 0$	$n_{00}$	$n_{10}$
Female, $Z = 1$	$n_{01}$	$n_{11}$

Table 1: A hypothetical  $2 \times 2$  table for adverse outcomes and gender in a sample of size  $n$ , when exposure status is unknown.

female), (adverse outcome, female), respectively. The estimate of the probability of an adverse outcome for males is  $n_{10}/(n_{10} + n_{00})$  and the estimate for females is  $n_{11}/(n_{11} + n_{01})$ . We can then estimate the dose–response probabilities by solving the equations

$$n_{10}/(n_{10} + n_{00}) = p_{A|NE}(1 - p_{EX|male}) + p_{A|EX}p_{EX|male}; \quad (5)$$

$$n_{11}/(n_{11} + n_{01}) = p_{A|NE}(1 - p_{EX|female}) + p_{A|EX}p_{EX|female}. \quad (6)$$

Any statistical analysis is basically a generalization of this idea. One uses the observed data (health outcomes and covariates) to construct an estimate of the dose–response function, even though individual doses are not available.

Here is how the mathematics work. The dose–response probability density or mass function is  $f(y|x, z, \beta)$ , while the DTOD is  $f(x|z)$ . In our simple example, we have the identifications

$$\begin{aligned} \beta &= (p_{A|NE}, p_{A|EX}); \\ p_{A|NE} &= f(\text{adverse, i.e., } y = 1 | \text{not exposed, i.e., } x = 0); \\ p_{A|EX} &= f(\text{adverse, i.e., } y = 1 | \text{exposed, i.e., } x = 1); \\ p_{EX|male} &= f(\text{exposed, i.e., } x = 1 | \text{male, i.e., } z = 0); \\ p_{EX|female} &= f(\text{exposed, i.e., } x = 1 | \text{female, i.e., } z = 1). \end{aligned}$$

If  $X$  is a discrete random variable, then the probability distribution or mass function for adverse outcomes in the presence of covariates is

$$f(y|z, \beta) = \sum_{\text{possible values of } x} f(y|x, z, \beta)f(x|z). \quad (7)$$

With a little checking of algebra, one can see that (4) is a special case of (7).

If  $X$  is a continuous random variable, then the probability density function for adverse outcomes in the presence of covariates is

$$f(y|z, \beta) = \int f(y|x, z, \beta)f(x|z)dx. \quad (8)$$

This discussion has proceeded as if the DTOD were known analytically. However, in the context of Hoffman, et al., this is not the case, and instead the best that one can do is to simulate observations from the DTOD. The question is how to use simulated observations to obtain dose–response estimates. There are a wide variety of possibilities, depending on the problem and whether one wants to use a Bayesian or a frequentist approach. However, it is important to see that at least in principle, one can do enough simulation to eliminate any Monte–Carlo effect from the analysis. Consider our simple example of a binary response (adverse or not), binary dose level (exposed or not), and binary covariate (male or female). In order to do the analysis (5)–(6), we have assumed that we knew the probability of dose level for males and females. In the Monte–Carlo context, these probabilities are not known, but we can simulate them. Thus, we could for example, simulate dose levels many thousands of times for males, and the fraction of the times that the simulated values indicate a non–zero dose would be an essentially error–free estimate of the probability of non–zero dose for males. A similar calculation can be done for females. Thus, brute–force simulation at the order of many thousands of times can, at least in principle, allow us to compute quantities such as (7) and (8) essentially exactly.

Once one realizes that brute–force simulation can allow one to compute the essential quantities (7) and (8), it then becomes a matter of how to do this efficiently. From the frequentist point of view, the maximum likelihood estimate can be computed via a Monte–Carlo EM algorithm. This is closely related to the Bayesian analysis described below.

## 4.2 Gibbs Sampling

The Bayesian approach to statistics treats all parameters as random variables, with the randomness of a parameter representing uncertainty about its value. In this section, we give a quick introduction to the Bayesian paradigm. The reader is referred to Berger (1985) for a thorough introduction.

Bayesian analysis of parametric models requires specifying a likelihood that is then interpreted as the conditional density of the data given the parameters. It also requires a prior distribution for the parameters, representing knowledge about the parameters prior to data collection. The product of the prior and likelihood is the joint density of the data and the parameters. Often, one uses non-informative priors, meaning that the prior tells us extremely little about the parameters, relative to what is learned from the sample. However, if there is substantial prior knowledge about some parameters, then using informative priors for them leads to a more effective analysis.

Given the joint density of the data and parameters, one can integrate out the parameters to get the marginal density of the data. One can then divide the joint density by this marginal density

to get the posterior density, i.e., the conditional density of the parameters given the data. The posterior summarizes all of the information about the values of the parameters and is the basis for all Bayesian inference. For example, the mean, median, or mode of the posterior density are all suitable point estimators. A region with probability  $(1 - \alpha)$  under the posterior is called a “credible set”, and is a Bayesian analog to a confidence region.

Computing the posterior distribution is often a non-trivial problem, because it usually requires high-dimensional numerical integration. This computational problem is the subject of much recent research, with many major advances. The method currently receiving the most attention in the literature is the Gibbs sampler, good introductions to which are given by Smith & Gelfand (1992) and Casella & George (1992).

The Gibbs sampler generates a Markov chain whose stationary distribution is the posterior distribution. The key feature of the Gibbs sampler is that this chain can be simulated using only the joint density of the parameters and the data, e.g., the product of the likelihood and the prior, and not the unknown posterior density. If the chain is run long enough, then the observations in a sample from the chain are approximately identically distributed with common distribution equal to the posterior. Thus posterior moments, the posterior density, and other posterior quantities can be estimated from a sample from the chain.

The Gibbs sampler is most easily understood when there is no uncertainty about the doses, and when there are no covariates. Consider a sample of size  $n$  independent individuals with known doses  $X_1, \dots, X_n$  and responses  $Y_1, \dots, Y_n$ . The responses and the doses are related through a model, which has probability density or mass function  $f(y|x, \beta)$ . In this case the likelihood is

$$\prod_{i=1}^n f(Y_i|X_i, \beta).$$

Let  $(\tilde{\mathbf{Y}}, \tilde{\mathbf{X}})$  refer to the ensemble of complete data. If  $\pi(\beta)$  denotes a prior distribution for  $\beta$ , then the density of  $(\tilde{\mathbf{Y}}, \beta)$  given  $\tilde{\mathbf{X}}$  is

$$\pi(\beta) \prod_{i=1}^n f(Y_i|X_i, \beta).$$

The posterior distribution of  $\beta$  is then

$$\pi(\beta|\tilde{\mathbf{Y}}, \tilde{\mathbf{X}}) = \frac{\pi(\beta) \prod_{i=1}^n f(Y_i|X_i, \beta)}{\int \pi(v) \prod_{i=1}^n f(Y_i|X_i, v) dv}. \quad (9)$$

The practical problem is that the denominator of (9) may be very difficult to compute. Numerical integration typically fails to provide an adequate approximation even when there are as few as three or four components to  $\beta$ .

The Gibbs sampler is one solution to the dilemma, although other methods are possible. The Gibbs sampler is an iterative, Monte-Carlo method consisting of two main steps:

- Form a sequence of computer-generated observations  $\beta_1, \beta_2, \dots$  from the posterior distribution of  $\pi(\beta|\tilde{\mathbf{Y}}, \tilde{\mathbf{X}})$ ;
- Quantities such as the posterior mean are estimated by the sample mean of  $\beta_1, \beta_2, \dots$ , while kernel density estimates are used to approximate the entire posterior density or the marginal posterior density of a single parameter or subset of parameters.

Generating pseudo random observations is the first step of the Gibbs sampler, and is often done using the Metropolis–Hastings algorithm, see Smith & Gelfand (1992), and the details listed below.

The mechanics of stopping the Gibbs sampler, and whether one should use one long sequence as described here or a number of shorter sequences, are currently a matter of controversy and is not discussed here.

In practice, the doses  $X$  are not known exactly. The usual device is to treat the unobserved  $X$ 's as unobserved random effects (parameters), and at each stage of the Gibbs sampler to generate new versions of these doses. We describe this more explicitly below.

### 4.3 Gibbs Sampling When There Are No Parameters In The Distribution Of Target–Organ Dose, and Observations are Independent

If the target–organ dose ( $X$ ) were known exactly, then the model relating the response ( $Y$ ) to dose and covariates ( $Z$ ) has a probability density function  $f(y|x, z, \beta)$ . The distribution of the target–organ dose given the covariates is denoted by  $f(x|z)$ . In this section, we assume that there are no parameters in this latter distribution, and that one can use a computer to draw Monte–Carlo samples from it. For the sake of simplicity, we will assume that the study consists of  $n$  independent individuals. Handling more complex situations is difficult in practice but mathematically requires only more complex notation, which we outline below.

Bayesian analysis requires that we have a prior distribution for  $\beta$ , and we will let  $\pi(\beta)$  be the prior density or mass function for  $\beta$ . Then the likelihood for a sample of size  $n$  is

$$\left\{ \pi(\beta) \prod_{i=1}^n f(Y_i|X_i, Z_i, \beta) \right\} \prod_{i=1}^n f(X_i|Z_i). \quad (10)$$

The Gibbs sampling algorithm in this context works in a series of steps, as follows.

1. Construct a series of starting target–organ doses,  $X_{1,o}, \dots, X_{n,o}$ . For example, in the Hoffman, et al. context, these might be the means from simulations of target–organ doses.

2. Construct a starting value  $\beta_o$ , which would typically be a standard regression pretending that the starting target–organ doses are exact.
3. Given the target–organ doses, and the current value  $\beta_o$ , generate a new value of  $\beta$ . Record the new value and then replace the old  $\beta_o$  by the new value. This step is typically specific to the problem, but it has a useful property. Specifically, given the current target–organ doses, updating  $\beta$  is a standard exercise in Bayesian Gibbs sampling using only the terms inside the curly brackets in (10) and is not repeated here. Basically one can use algorithms from the literature for logistic regression, parametric survival analysis, etc.
4. Now one must update the target–organ doses, starting from the new value  $\beta_o$  generated in the step 3 and from the current values of the target–organ doses, namely  $X_{1,o}, \dots, X_{n,o}$ . We describe here how to update the target–organ doses themselves, using the Metropolis–Hastings algorithm. The steps are as follows.

- First, remember that as in Hoffman, et al., one is able to draw samples of observations of target–organ doses given only the covariates, i.e., from  $f(x|z)$ . Do this, calling them "candidate" target–organ doses,  $X_{1,c}, \dots, X_{n,c}$ .
- The new target–organ doses will either be the candidates or the current values  $X_{1,o}, \dots, X_{n,o}$ . To decide which one, for each of  $i = 1, \dots, n$ , define

$$\alpha_i = \min\{1, f(Y_i|X_{i,c}, Z_i, \beta_o)/f(Y_i|X_{i,o}, Z_i, \beta_o)\}.$$

- Then with probability  $\alpha_i$ , once accepts the candidate and changes from  $X_{i,o}$  to  $X_{i,c}$ ; otherwise, the current value of the target–organ dose remains unchanged.

Relabel the new values as  $X_{1,o}, \dots, X_{n,o}$ .

5. Repeat the steps 3 and 4 a large number of times to obtain the posterior distribution of  $\beta$ .

#### 4.4 Gibbs Sampling When There Are No Parameters In The Distribution Of Target–Organ Dose, and Observations are Not Independent

In section 4.3, we assumed for simplicity that one could treat all individuals in the study as independent. In practice though, this might not be the case.

For example, the population may be stratified by a variable which either may itself have uncertainty in it, or the effect of which might be uncertain. For instance, consider a study made up of individuals who work in mines. The levels of radiation in a mine are a source of uncertainty in

dose which would be correlated for all individuals who work in the mine. The uncertainty analysis might then contain a component for the mine, as well as components for individuals within a mine.

In our notation, it is fairly easy to handle correlations in uncertainties caused by stratification, at least mathematically. The unit of analysis is the stratification variable (the mine), and what we have called  $Y$ ,  $X$  and  $Z$  become the ensemble of responses, target–organ doses and covariates for all the individuals within the mine. Not a thing changes mathematically in this framework. In other words, once the hard work of uncertainty analysis is done, the Bayesian analysis is itself mathematically straightforward.

## 5 DISCUSSION

The main points of this article have been the following.

1. In estimating dose–response, there is an emerging class of problems where the information about the target–organ dose is in the form of a distribution, often subjectively derived.
2. These problems are formally (mathematically) in the framework of the Berkson error model, with the important exception that the distribution of target–organ dose is unknown.
3. One can draw samples (via computer) from the distribution of target–organ dose.
4. It is relatively straightforward conceptually to go from being able to draw samples from the distribution of target–organ dose to a dose–response analysis, using either Bayesian or frequentist techniques. One need not be forced to restrict analysis of dose–response to a “representative value” of dose.

## REFERENCES

- Berger, J.O. (1985). *Statistical Decision Theory and Bayesian Analysis, second edition*, Springer-Verlag, New York.
- Carroll, R. J., Ruppert, D. & Stefanski, L. A. (1995). *Measurement Error in Nonlinear Models*. Chapman & Hall, London.
- Casella, G. & George, E. I. (1992). Explaining the Gibbs sampler. *American Statistician*, 46, 167–174.
- Gelman, A., Bois, F. & Jiang, J. (1996). Physiological pharmacokinetic analysis using population modeling and informative prior distributions. *Journal of the American Statistical Association*, 91, 1400–1412.

- Hoffman, F. O. and the NCRP Scientific Committee 64-17 (1996). *A Guide for Uncertainty Analysis in Dose and Risk Assessments Related to Environmental Contamination*. NCRP Commentary #14, National Council on Radiation Protection and Measurement, Bethesda MD.
- Hoffman, F. O., Simon, S. L. and Thiessen, K. M. (1995). The role of uncertainty analysis in dose reconstruction and risk assessment. In *Environmental Dose Reconstruction and Risk Implications, Proceedings of the Thirty-First Annual Meeting of the National Council on Radiation Protection and Measurement*, J. E. Till, editor, National Council on Radiation Protection and Measurement, Bethesda MD.
- Pierce, D. A., Stram, D. O., Vaeth, M., Schafer, D. W. (1992). Some insights into the errors in variables problem provided by consideration of radiation dose-response analyses for the A-bomb survivors. *Journal of the American Statistical Association*, 87, 351-359.
- Prentice, R. L. (1982). Covariate measurement errors and parameter estimation in a failure time regression model. *Biometrika*, 69, 331-342.
- Schafer, D. W. (1992). Replacement methods for measurement error models. Technical Report #151, Department of Statistics, Oregon State University.
- Smith, A. F. M. & Gelfand, A. E. (1992) Bayesian statistics without tears: a sampling-resampling perspective. *American Statistician*, 46, 84-88.
- Wang, N., Lin, X., Gutierrez, R. G. & Carroll, R. J. (1998). Generalized Linear Mixed Measurement Error Models. *Journal of the American Statistical Association*, to appear.