LA-UR-01-1573
submitted to Radiation Protection Dosimetry

Bayesian Internal Dosimetry Calculations Using Markov Chain Monte Carlo

by

Guthrie Miller, Harry F. Martz, Tom T. Little, and Ray Guilmette

October 4, 2001

Summary

A new numerical method for solving the inverse problem of internal dosimetry is described. The new method uses Markov Chain Monte Carlo and the Metropolis algorithm. Multiple intake amounts, biokinetic types, and times of intake are determined from bioassay data by integrating over the Bayesian posterior distribution. The method appears definitive, but its application requires a large amount of computing time.

[0]Key words: internal dosimetry, intake estimation, Bayesian analysis, bioassay, internal dosimetry algorithms, Markov-Chain Monte Carlo

1 Introduction

Internal dosimetry is concerned with the problem of determining the radiation dose to workers caused by forms of radiation that cannot be measured directly (as with a dosimetry badge). If, for example, the α-emitting nuclide Pu-239 is inhaled, it will impart radiation dose to the lungs, and after dissolving will be absorbed to blood and deposited in the bone and liver, imparting dose to these organs. Monitoring for exposure to Pu-239 is done by making bioassay measurements (for example, urine, fecal, lung count, etc.).

The measurements are interpreted using biokinetic models that describe how Pu-239 is transported through the body. The biokinetic models describe how a unit amount of material taken into the body in a certain way (for example, inhalation) will later in time appear in various bioassay compartments (for example, the lungs, urinary excretion) and how radiation dose will be accumulated in the course of time in the different body organs and tissues. Standard biokinetic models have been proposed by the International Commission on Radiation Protection (ICRP) (e.g., ICRP publications 30, 54, 66, and 78)[1, 2, 3, 4, 5].

Given a set of agreed-upon biokinetic models, the inverse problem of internal dosimetry is to use the bioassay measurements to infer if and when intakes may have occurred and the magnitude of the resultant radiation dose to the worker. In using intake-based biokinetic models we are required to determine the time and amount of intakes and to assess the 50-yr effective whole body dose to the worker (the CEDE) associated with each intake. The process obviously entails considerable uncertainty, so quantitatively assessing uncertainty is also of great importance. We have been pursuing a Bayesian statistical approach to this problem.[6, 7, 8, 9, 10, 11] The present work describes a method that extends our previous work and appears to be definitive.

2 Formulation of the Problem

In the problem of internal dosimetry there are M bioassay data y_j taken at times t_j for $j = 1, M$. From these data we seek to determine N possible intakes with amounts ξ_i, biokinetic types l_i, times of intake t_i, for $i = 1, N$. The intake times are ordered, so that $t_1 < t_2 < \ldots t_N$. The domain of time t_i is the time interval Δt_i. That is, t_i is in the interval Δt_i. The intervals Δt_i cover the time domain of all possible intakes in a non-overlapping and ordered way. The time intervals are often chosen to be the times between successive bioassay measurements, in which case $N = M - 1$. The time intervals are chosen to be sufficiently small so that multiple intakes are unlikely in any interval.

Using the notation

$$Y \equiv \{y_1, y_2, \ldots y_M\} \tag{1}$$

$$\Xi \equiv \{\xi_1, l_1, t_1, \ldots \xi_N, l_N, t_N\}, \tag{2}$$

the problem is to determine the parameters Ξ from the data Y. Using Bayes theorem, the Probability distribution of Ξ given Y can be immediately written down as

$$P(\Xi|Y) \propto P(Y|\Xi)P(\Xi), \tag{3}$$

that is, the probability of particular values of the parameters given the data is proportional to the probability of the measured values of the data given the parameters (the likelihood function) times the prior probability of the parameter values. The calculational problem is then to integrate (or sum) over the full detailed posterior probability distribution function in order to determine the marginal probability distribution of quantities of interest. The multi-dimensional integration problem is well suited to the Markov Chain Monte Carlo Method[12] using the Metropolis algorithm (see Appendix A).[13]

The likelihood function $P(Y|\Xi)$ is of the form

$$P(Y|\Xi) \propto \exp\left(\sum_{j=1}^{M} \mathcal{L}_j(\Xi)\right) \tag{4}$$

because of the assumed independence of the M measurements, where $\mathcal{L}_j(\Xi)$ is the log-likelihood function for the j^{th} measurement.

The prior probability distribution $P(\Xi)$ is taken to be of the form

$$P(\Xi)\, d\Xi = \prod_{i=1}^{N} P(\xi_i)\, d\xi_i\, P(l_i) P(t_i)\, dt_i. \tag{5}$$

The prior probability distribution of biokinetic types l is a discrete probability distribution over $\{l_1, l_2, \ldots l_{n_l}\}$, usually uniform except that the ICRP-recommended default model is given a higher probability.

The prior probability distribution for ξ_i and t_i depends on whether or not a known incident has occurred in the intake time interval Δt_i. Two cases are considered, an incident reported in the time interval (incident) and no report of an incident (non-incident).

incident The prior probability distribution over intake time $P(t_i)$ is assumed to be

$$P(t_i) = \delta(t_i - t_i^{(inc)}), \tag{6}$$

where $\delta(.)$ is the delta function, and $t_i^{(inc)}$ is the known time of the incident. The prior probability distribution over intake amount ξ_i is assumed to be a broad log normal (standard deviation of the log of ξ_i equal to 3) with median determined by incident indicators (for example, nose swipe results, air monitor readings) as discussed in Ref. [11].

non-incident If no incident has occurred in the intake interval Δt_i, the prior probability distribution of intake time t_i is assumed to be uniform in the

interval Δt_i, and the distribution of intake amount ξ_i is assumed to be given by the following (a special case of the gamma distribution),

$$P(\xi_i)\, d\xi_i = \frac{\alpha_i \Delta t_i}{\xi_i} \left(\frac{\xi_i}{\xi_i^{(\max)}} \right)^{\alpha_i \Delta t_i}, \tag{7}$$

which we have called the "alpha distribution".[11] The parameter $\xi_i^{(\max)}$ specifies the maximum intake allowed and otherwise is unimportant. The parameter α_i, which can be interpreted as the "intake" probability per unit time in the i^{th} time interval, is meant to be determined empirically using population averages. Using Los Alamos plutonium historical data from 1980 to the present α was found to be very small, 0.001 yr^{-1} or less.[11]. The smallness of α shows that the internal dosimetry problem for non-incident-related intakes is a "needle-in-the-haystack" problem of detecting very rare events. In such cases Bayesian methods avoid an inordinate number of false positives.

The likelihood function $P(Y|\Xi)$ gives the probability of measuring data values Y given parameters Ξ, considered as a function of Ξ. In this paper it is assumed that the Gaussian approximation for the likelihood function is applicable. In this case

$$\mathcal{L}_j(\Xi) = -\frac{1}{2} \left(\frac{(y_j - \psi_j^{(l)})^2}{\sigma_j^2} + \log(\sigma_j^2) \right), \tag{8}$$

Here y_j is measurement value, $\psi_j^{(l)}$ is the calculated value based on the parameters, in particular the biokinetic type l, and σ_j is the uncertainty standard deviation associated with the j^{th} measurement. The calculated value is given by

$$\psi_j^{(l)} = \sum_{i=1}^{N} \xi_i f^{(l)}(t_j - t_i), \tag{9}$$

where ξ_i is the magnitude of the i^{th} intake, $f^{(l)}(t)$ is the biokinetic retention fraction for biokinetic type l at time t after the intake, and t_j and t_i are the times of the j^{th} measurement and the i^{th} intake. Note that $f^{(l)}(t_j - t_i) = 0$ for $t_j < t_i$. The uncertainty σ_j is composed of measurement uncertainty $\sigma_j^{(m)}$ and a multiplicative factor uncertainty $\sigma_j^{(f)}$ (for example, for a lung count, the measurement uncertainty would be the counting statistics uncertainty, while the multiplicative factor uncertainty would be the estimated uncertainty of the calibration factor, which is mostly associated with chest wall absorption uncertainty).

$$\sigma_j^2 = (\sigma_j^{(m)})^2 + (\sigma_j^{(f)} \psi_j^{(l)})^2. \tag{10}$$

3 Numerical Algorithm

Using the Metropolis algorithm[13] a Markov chain of the parameter values Ξ is generated that has as its stationary distribution the joint posterior distribution (a multivariate distribution, because Ξ is a vector) of Eq. 3 (see Appendix A).[12] A Markov chain is a sequence of random variables Ξ_k such that Ξ_{k+1} depends on Ξ_k and does not depend further on the history of the chain. Given such a chain we can effectively integrate over the posterior distribution by using the relation

$$\int f(\Xi)P(\Xi|Y)\,d\Xi \rightarrow \frac{1}{N_k}\sum_{k=1}^{N_k} f(\Xi_k) \tag{11}$$

for $N_k \rightarrow \infty$, where $f(.)$ is an arbitrary function of Ξ.

We lump together the three parameters (ξ_i, l_i, t_i) of intake i as the i^{th} component of the intake vector Ξ and chain update these components one by one. The intake components are selected for updating probabilistically, with probability given by

$$P_i \propto \text{Max}(\text{CEDE}_i, \text{CEDE}_{\text{min}}), \tag{12}$$

that is, the attention given to the i^{th} component is proportional to the current CEDE associated with the i^{th} intake. However, for CEDE's below the lower limit CEDE_{min}, all components are given equal attention. The lower limit CEDE_{min} is usually chosen to be 0.1 mSv (0.01 rem).

The components Ξ_i are updated using a probabilistic random walk scheme, where with some given probability (a parameter of the code) the new value is selected within a small neighborhood of the current value (random walk), or, with the complement of that probability, the new value is generated from the entire domain.

The chain has a starting value Ξ_0 that influences to some extent average values obtained from Eq. 11 for finite numbers of trials N_k. The pseudo-random numbers used to generate the chain also have a seed value that determines the sequence. Our approach to convergence of results is to compare results from two chains with different random number seeds, one starting from the minimum allowed value of Ξ and the other starting from the maximum allowed value.

A code validation test case is described in Appendix B.

4 Example Using Actual Data

In Fig. 1 are shown actual Los Alamos urine data for Pu-238 urine excretion from a single individual over a number of years. The error bars represent plus or minus one standard deviation of the measurement uncertainty. In addition, it is known that this person was involved in an incident on 31-October-1980. The incident classification indicates a relatively low probability of an inhalation intake resulting from the incident. This case was discussed in Ref. [14] as

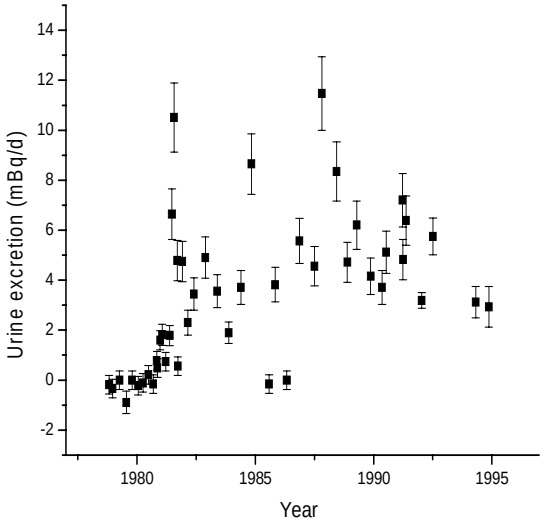

Figure 1: Actual Pu-238 urine excretion data.

example 3, and the data is available in the file BIOASSAY.333 downloadable from our web site[15] (in the software package BayesII). In interpreting these data, the "wing-9 accident" biokinetic model[9] was used in addition to the six standard ICRP-30 biokinetic models (class W and class Y, 0.2, 1.0, and 5.0 μm AMAD particle size).

In Fig. 2 is shown the urine excretion data together with the calculated expected value. The median of the log-normal prior for the 1980 incident was chosen to be 37 Bq (1 nCi). There is actually only one "positive" intake, using the definition of "positive"

$$P(\text{CEDE} > 1\,\text{mSv}(0.1\text{rem})|\text{data}) > 0.95. \tag{13}$$

The most probable biokinetic type for this intake is the special wing-9 type (IEE) with 55 % probability.

In Fig. 3 are shown the year-by-year CEDE's lumping together all intakes occurring in a given year. The square dots represent the magnitude of the expected CEDE for each year while the shaded bars show the 90% credible interval (5% to 95%) for those cases where the upper limit exceeds 1 mSv (0.1 rem). For 1980, we do not have confidence that the CEDE exceeds 1 mSv since the lower credible limit does not exceed 1 mSv. The total expected CEDE for all years is 520 mSv (52 rem) with 5% and 95% credible limits of 410 to 620 mSv.

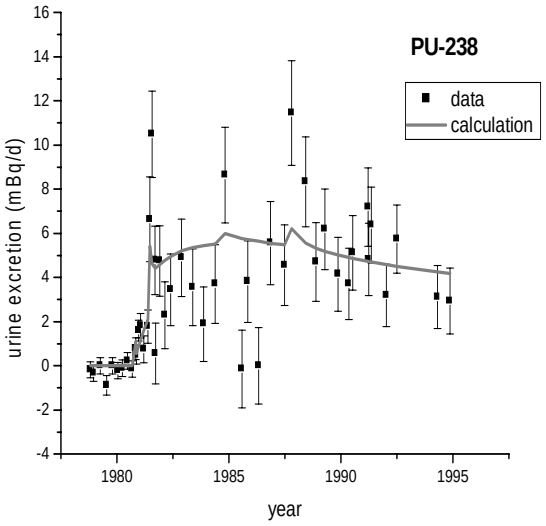

Figure 2: Urine data and calculated expected value of urine excretion.

The error bars in Fig. 2 are larger than those in Fig. 1 because they include multiplicative uncertainty in addition to measurement uncertainty. The quantity $\sigma^{(f)}$ in Eq. 10 was assumed to have the value 0.3, which is the value we normally assign for simulated 24-hour urine samples.

In Fig. 3 note that in some cases the expected value of CEDE exceeds 1 mSv while the 95% limit does not. This is possible for distributions mostly concentrated at small values but having a tail extending to large values.

If the incident information is not used to analyze the data, a very similar year-by-year intake scenario is calculated, as shown in Fig. 4. However, when the incident information is not used, no single intake is "positive" (however the sum of all intakes in 1980 is "positive").

The prior probability parameter α representing a worker's intake probability per unit time for non-incident situations was assumed to be 0.001 per year in the foregoing (acute intake situation). If the example data are analyzed assuming a value of α 100 times larger (chronic intake situation), the results shown in Figs. 5 and 6 are obtained. Many more intakes are now possible, although no individual intake is actually "positive". That is, it seems likely that many intakes have occurred, but it is not possible to identify with certainty the times of these intakes.

The total CEDE from all intakes is well determined by the data in all of these cases. For example, Fig. 7 shows the result assuming the normal small

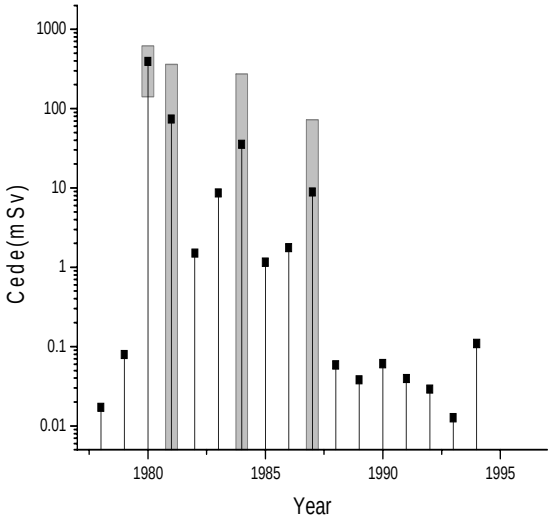

Figure 3: Calculated year-by-year expectation value of CEDE's. The shaded bars represent the 90% credible interval (5% to 95%) for those cases where the upper limit exceeds 1 mSv.

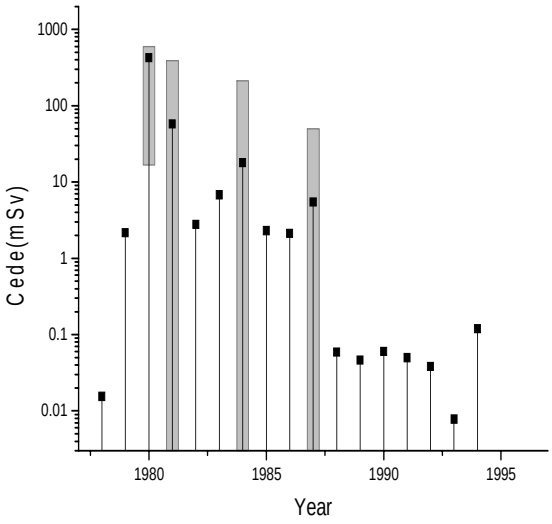

Figure 4: Calculated year-by-year expectation value of CEDE's when prior information about incident is not used. The shaded bars represent the 90% credible interval (5% to 95%) for those cases where the upper limit exceeds 1 mSv.

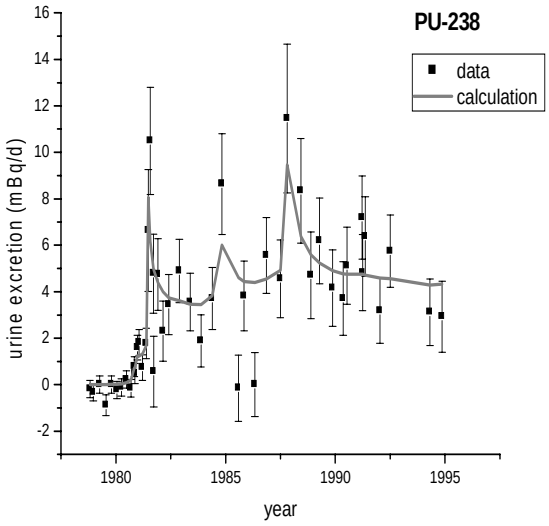

Figure 5: Urine data and calculated expected value of urine excretion when prior probability of intake per unit time 100 times larger than normal is assumed ($\alpha \rightarrow 100 \times \alpha$). Larger α corresponds to a chronic rather than acute intake situation.

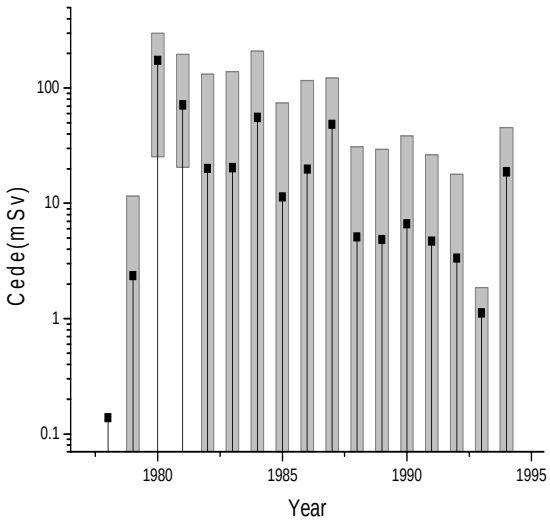

Figure 6: Calculated year-by-year expectation value of CEDE's when prior probability of intake 100 times larger is assumed. The shaded bars represent the 90% credible interval (5% to 95%) for those cases where the upper limit exceeds 1 mSv. Many more intakes are now possible, however the intake dates are not identified with certainty.

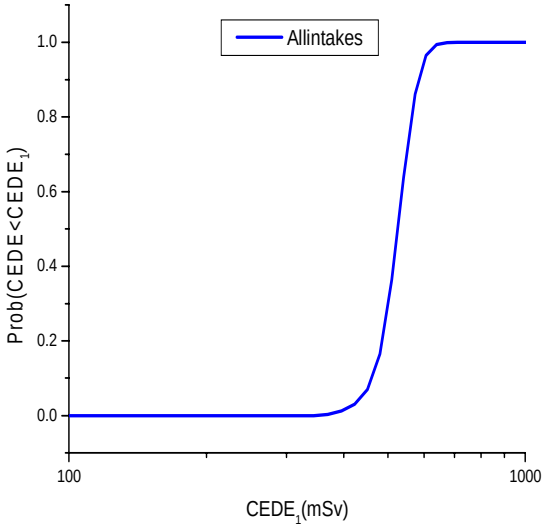

Figure 7: Calculated cumulative distribution of total CEDE.

value of α without using the incident information.

5 Discussion

The Markov Chain algorithm appears to provide a definitive solution of the inverse problem of internal dosimetry, that of calculating the intake scenario given the bioassay data and an agreed-upon set of biokinetic models. By a definitive solution we mean an exact solution of the problem without simplifying assumptions.

The Bayesian method allows us to directly address the question of interest ("what is the dose?") and to quantify the uncertainties. The quantitative assessment of uncertainty, which is based on calculation of the probability distribution of intake parameters given the data-an inherently Bayesian entity-is not possible using non-Bayesian methods. Not surprisingly, it is simply not possible to identify the times of intakes with certainty in many cases, although other quantities, such as annual dose or total CEDE are usually relatively well determined by the data.

The drawback of this method is that it requires a large amount of computer time. Our rule of thumb for convergence is at least 1 million to 10 million chain iterations per possible intake, which translates to a 1 to 10 hour run for a case such the example discussed here (about 50 possible intakes) using a 1 Ghz Pentium processor. Population studies involving thousands of cases, such as those carried out to determine the prior parameter α[11] are then not practical on a desktop workstation. Our future plans are to use massively parallel supercomputers to carry out such studies.

A Appendix–Markov Chain Monte Carlo using the Metropolis-Hasting Algorithm

Suppose we are interested in making statistical inference about a parameter (possibly vector valued) Ξ. We characterize our information (or lack of information) about the distribution of $\Xi = \{\xi_1, \xi_2, \ldots, \xi_n\}$ as $P(\Xi)$ (prior distribution). Data are collected and represented by the likelihood or $P(Y|\Xi)$. In any Bayesian analysis, inference on the parameters depends on the calculated posterior distribution

$$P(\Xi|Y) = \frac{P(\Xi)P(Y|\Xi)}{\int_{\Xi} P(\Xi)P(x|\Xi)\, d\Xi}. \tag{14}$$

In many situations, use of the posterior distribution given by (14) requires numerical calculation. Monte Carlo integration evaluates the expectation value of an arbitrary function $f(.)$ of Ξ, $E[f(\Xi)]$, by drawing samples $\{\Xi_k, k = 0, \ldots, N_k\}$ from the posterior distribution and then approximating

$$E[f(\Xi)] \approx \frac{1}{N_k} \sum_{k=0}^{N_k} f(\Xi_k). \tag{15}$$

So the population mean of $f(\Xi)$ is estimated by a sample mean. Markov Chain Monte Carlo method is a powerful tool in such cases.

The following description of the Metropolis-Hastings algorithm[13, 16] closely follows that given in Ref. [12]. Using the Metropolis-Hastings algorithm, for each state k, the next state Ξ_{k+1} is chosen by first sampling a candidate point Ξ' from a proposal distribution $q(.|\Xi_k)$. Note that the proposal distribution may depend on the current point Ξ_k. The candidate point Ξ' is then accepted with probability $\alpha(\Xi_k, \Xi')$ where

$$\alpha(\Xi, \Xi') = \min\left(1, \frac{P(\Xi')q(\Xi|\Xi')}{P(\Xi)q(\Xi'|\Xi)}\right). \tag{16}$$

If the candidate point is accepted, the next state becomes $\Xi_{k+1} = \Xi'$. If the candidate is rejected, the chain does not move, i.e. $\Xi_{k+1} = \Xi_k$.

Thus the Metropolis-Hastings algorithm is extremely simple:

1 Initialize Ξ_0 and set $k = 0$.

2 Generate an observation Ξ' from a candidate distribution $q(\Xi'|\Xi_k)$.

3 Generate a uniform (0,1) random variable u.

4 If $u \leq \alpha(\Xi_k, \Xi')$ set $\Xi_{k+1} = \Xi'$, otherwise set $\Xi_{k+1} = \Xi_k$.

5 Increment k, go to step 2

Remarkably, the proposal distribution $q(.|.)$ can have practically any form and the stationary distribution of the chain will be $P(.|Y)$

The Metropolis[13] algorithm considers only symmetric proposals, having the form $q(\Xi|\Xi') = q(\Xi'|\Xi)$ for all Ξ and Ξ'. A special case of the Metropolis algorithm is random-walk Metropolis, for which $q(\Xi'|\Xi) = q(|\Xi - \Xi'|)$. Typically $q(\Xi'|\Xi)$ is a constant for Ξ' within some given prescribed neighborhood of Ξ.

Typical implementation of the algorithm generates an initial "large" number of iterations (called the burn-in) until the influence of the initial value of the chain has subsided. The burn-in samples are discarded, and the samples generated thereafter are used as samples from the posterior distribution of Ξ.

B Appendix–Code Validation

Several test cases where the correct result is known were used to validate the computer code (ID1.1).

single measurement In the case of a single measurement, we know from previous work[10] that for a prior probability distribution describing rare non-incident-related intakes, an intake is "positive" only when the measurement is about 4 or more standard deviations from zero. Figure 8

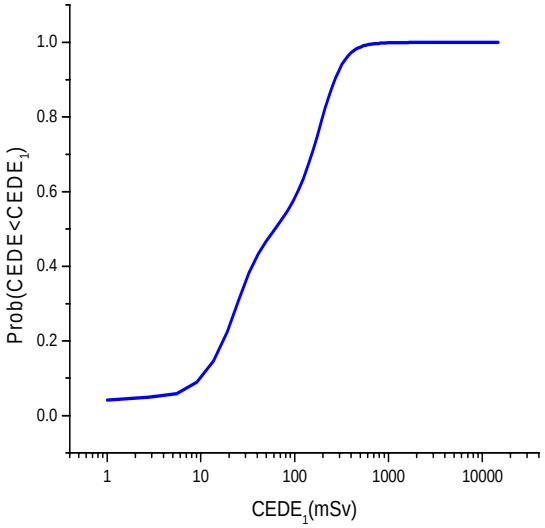

Figure 8: Cumulative posterior distribution of CEDE for a single measurement 4.6 standard deviations above zero.

shows the cumulative posterior probability distribution of CEDE for a measurement of Pu-239 urinary excretion of 1.7 ± 0.37 mBq/d (0.046 ± 0.01 pCi/d) assuming $\alpha = 0.001$/yr. The distribution is slightly positive using the definition of Eq. 13. This agrees with the result obtained using the unfolding algorithm UF3.5 (described in Ref. [9] and downloadable from our web site[15] as the BayesII software package). The UF3.5 calculation should be exact in this case. The ID1.1 and UF3.5 calculations assumed a set of six ICRP-30 inhalation models (class Y and class W, 0.2, 1, and 5 μm AMAD) and an intake time interval of one year preceding the measurement. Using only a single biokinetic model rather than a set of six produces a simpler looking cumulative posterior distribution, but does not change the number of standard deviations required for "positive". Similarly, using a fixed intake date of 6 months preceding the measurement (as is done in the UF3.5 code) rather than allowing intake date to be variable (the ID1.1 code allows both possibilities) does not change the number of standard deviations required for "positive".

calculated data In this case we use calculated urine bioassay data for nine samples in an eight-month period following an intake of 370 Bq (10 nCi) of class Y, 1μm AMAD Pu-239 (this data is in the file BIOASSAY.TST in the BayesII software package[15]). There are a number of possibilities for

running the code, for example: 1) the data can be treated as resulting from a known incident, in which case the median of the log normal prior needs to be specified (the parameter a shown in Table 1), 2) the data can be treated as non-incident related ($\alpha = 0.001$ yr^{-1} assumed for nonincident intake time intervals), and 3) the measurement uncertainty of the data ($\sigma^{(m)}$ in Eq. 10) can be decreased, which makes the urine excretion pattern more significant. The UF3.5 code does not allow a variable intake date, but otherwise treats these cases exactly. Table 1 shows results for the calculated CEDE for various run parameters.

Table 1: Calculation results using simulated test data corresponding to single 28.7 mSv CEDE intake.

run parameters	CEDE(mSv) ID1.1	UF3.5
incident		
a = 370 Bq[a]	21 (1.6, 35)[b]	21 (1.9, 34)
a = 0.37 Bq	10 (1.2, 32)	9.1 (1.4, 31)
$\sigma^{(m)} \rightarrow \sigma^{(m)} \times \frac{1}{10}$ [c]	27 (5.2, 33)	26 (3.7, 33)
non incident		
	18 (1.4, 38.5)	
$\sigma^{(m)} \rightarrow \sigma^{(m)} \times \frac{1}{10}$	28 (21, 37)	

[a] median of the log-normal prior–see text
[b] expectation value and 5% and 95% credible limits
[c] $\sigma^{(m)}$ is standard deviation of measurement uncertainty–see text

The CEDE results are the calculated expectation value and the 5% and 95% credible limits of the posterior distribution, given the data. The calculated result misses the mark in some cases (as it should) because the excretion pattern from the data is not well enough determined to rule out class W material. In the small measurement uncertainty cases the data are clearly consistent only with class Y, and the posterior probability of the correct biokinetic model is over 90 %.

In the two non-incident cases, the calculated expectation value of the intake date matches the correct intake date exactly. The UF3.5 code does not determine the date of intake from the data (but uses the midpoint of the preceding bioassay data interval), so it is not expected to reproduce the correct CEDE in these cases. The data and calculated expectation value of urine excretion for the first non-incident case appearing in Table 1 is shown in Fig. 9.

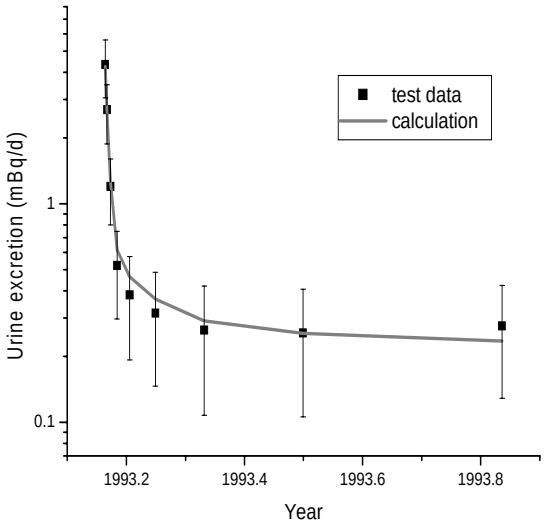

Figure 9: Test data and ID1.1 calculation result.

References

[1] International Commission on Radiological Protection. Limits for Intakes of Radionuclides by Workers, volume 2 of Annals of the ICRP. Pergamon Press, New York, 1979. ICRP Publication 30, Part 1.

[2] International Commission on Radiological Protection. Limits for Intakes of Radionuclides by Workers: an Addendum, volume 19 of Annals of the ICRP. Pergamon Press, New York, 1988. ICRP Publication 30, Part 4.

[3] Internation Commission on Radiological Protection. Recommendations of the International Commission on Radiological Protection, Individual Monitoring for Intakes of Radionuclides by Workers: Design and Interpretation, volume 19 of Annals of the ICRP. Pergamon Press, Oxford, 1988. ICRP Publication 54.

[4] International Commission on Radiological Protection. Human Respiratory Tract Model for Radiological Protection, volume 24 of Annals of the ICRP. Pergamon Press, New York, 1994. ICRP Publication 66.

[5] International Commission on Radiological Protection. Individual Monitoring for Internal Exposure of Workers, Replacement of ICRP Publication54, volume 27 of Annals of the ICRP. Pergamon Press, New York, 1998. ICRP Publication 78.

[6] G. Miller, W. C. Inkret, and H. F. Martz. Bayesian Detection Analysis for Radiation Exposure. Radiation Protection Dosimetry, 48:251–256, 1993.

[7] G. Miller, W. C. Inkret, and H. F. Martz. Bayesian Detection Analysis for Radiation Exposure, II. Radiation Protection Dosimetry, 58:115–125, 1995.

[8] G. Miller and W. C. Inkret. Bayesian Maximum Posterior Probability Method For Interpreting Plutonium Urinalysis Data. Radiation Protection Dosimetry, 63:189–196, 1996.

[9] G. Miller, W. C. Inkret, and H. F. Martz. Internal Dosimetry Intake Estimation Using Bayesian Methods. Radiation Protection Dosimetry, 82:5–17, 1999.

[10] G. Miller, W. C. Inkret, M. E. Schillaci, H. F. Martz, and T. T. Little. Analyzing Bioassay Data Using Bayesian Methods–A Primer. Health Physics, 78(6):598–613, 2000.

[11] G. Miller, W. C. Inkret, T. T. Little, H. F. Martz, and M. E. Schillaci. Bayesian Prior Probability Distributions for Internal Dosimetry. Radiation Protection Dosimetry, 94:347–352, 2001.

[12] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter. Markov Chain Monte Carlo in Practice. Chapman and Hall, New York, 1996.

[13] N Metropolis, Rosenbluth A. W., Rosenbluth M. N., and A. H. Teller. Equation of state calculations by fast computing machines. Journal of Chemical Physics, 21(6):1087–1092, 1953.

[14] G. Miller. Applications of Probability and Statistics in Health Physics, chapter 9, pages 323–330. Medical Physics Publishing, 2000. (downloadable from www.lanl.gov/bayesian).

[15] Los Alamos Center for the Application of Bayesian Statistics to Bioassay, Radiochemistry, and Internal Dosimetry. URL:www.lanl.gov/bayesian.

[16] W. K. Hastings. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. Biometrika, 57:97–109, 1970.