

Uncertainty Quantification for Dose-Response Models Using Probabilistic Inversion with Isotonic Regression: Bench Test Results

Roger Cooke
Resources for the Future
Dept. Math. Delft Univ. of Technology
Oct.22, 2007

1. Introduction

This technical paper describes an approach for quantifying uncertainty in the dose-response relationship¹ to support health risk analyses. This paper demonstrates uncertainty quantification for bioassay data using the mathematical technique of probabilistic inversion (PI) with isotonic regression (IR).

The basic PI technique applied here was developed in a series of European uncertainty analyses. (Illustrative studies from the EU-USNRC uncertainty analysis of accident consequence models are available on line at http://cordis.europa.eu/fp5-euratom/src/lib_docs.htm; the main report is ftp://ftp.cordis.europa.eu/pub/fp5-euratom/docs/eur18826_en.pdf.) Previous applications concerned atmospheric dispersion, environmental transport, and transport of radionuclides in the body, and were based on structured expert judgment. This report transfers these techniques to bioassay data. The focus of this evaluation is on mathematical techniques, the experimental data are not analyzed from a toxicological viewpoint. The analyses should be seen as illustrative rather than definitive. Background on isotonic regression (IR), iterative proportional fitting (IPF) and probabilistic inversion (PI) aimed at the uninitiated is provided.

In this report, observational uncertainty is discussed first, followed by a brief introduction to PI, iterative proportional fitting (IPF) and IR. Four cases of simple data are then analyzed to demonstrate how uncertainty in the DR modeling could be done in such simple cases. In reality, many more factors will contribute to the uncertainty, but we must first get such simple cases right, before attacking more difficult issues.

For two cases, benchmark dose (BMD) and the example chemical Nectorine, standard dose-response (DR) models proved suitable. For two other cases, example chemicals Frambozadrine and Persimonate, the data suggest a threshold model and a barrier model, respectively, for recovering observational uncertainty. Appendix 1 provides some results with alternative models and alternative optimization strategies. Complete specification of starting distributions and calculation scripts are provided in Appendix 2.

¹ Different DR relations may be seen as instantiations of a general functional form, for example a Taylor expansion of $\log(\text{dose})$, therefore the distinction between ‘parameter’ and ‘model’ uncertainty is more semantic than real, and is not maintained in this report.

2. Tent Poles

Tent poles for this approach are:

1. *Observational Uncertainty*: There must be some antecedent uncertainty over observable phenomena which we try to capture via distributions on parameters of DR models. This is the target we are trying to hit. When the target is defined, we can discuss (i) whether this is the right target, and (ii) whether we got close enough. Without such a target, the debate over how best to quantify uncertainty in DR models is under-constrained.

2. *Integrated Uncertainty Analysis*: The goal is to support integrated uncertainty analysis, in which potentially large numbers of individuals receive different doses. For this to be feasible, the uncertainty in DR must be captured via a joint distribution over parameters which does not depend on dose.

3. *Monotonicity*: Barring toxicological insights to the contrary, we assume that the probability of response cannot decrease as dose increases. It is not uncommon that data at increasing doses show a decreasing response rate, simply as a result of statistical fluctuations. A key feature in this approach is to remove this source of noise with techniques from isotonic regression.

3. Observational Uncertainty

Suppose we give 49 mice a dose of 21 [units] of some toxic substance. We observe a response in 15 of the 49 mice. Now suppose we randomly select 49 new mice from the same population and give them dose 21. What is our uncertainty in the number of mice responding in this new experiment?

Without making some assumptions, there is no “right” answer to this question, but the customary answer would be: “*our uncertainty in the number of responses in the second experiment is binomially distributed with parameters (49, 15/49)*”. Thus, there is a probability of 0.898 of seeing between 9 and 19 responses in the second experiment. A Bayesian with no prior knowledge (all probabilities are equally likely) would update a uniform $[0, 1]$ prior for the probability of response to obtain a $Beta(16, 35)^2$ posterior for the probability p of response, with expectation $16/51 = 0.314$, and variance 0.0038 . His “predictive distribution” on the number of responses is obtained as a $p \sim Beta(16, 35)$ mixture of binomial distributions $Bin(49, p)$. He would have approximately a 76% chance of seeing the result of the second experiment between 9 and 19. The two distributions are shown in Figure 1.

Since

- (i) These two distributions tend to be close, relative to the other differences encountered further on,
- (ii) The distribution $Bin(49, 15/49)$ is much easier to work with, and

² To recall, the $B(16, 35)$ density is $p^{16-1}(1-p)^{35-1}/B(16, 35)$, where $B(16, 35) = 15! \times 34! / 50!$.

- (iii) The assumption that before the first experiment, all probability values are equally likely, regardless of dose, is rather implausible,

we will take the binomial distribution to represent the observable uncertainty in this simple case.




Figure 1: Binomial and Bayes Non-Informative Predictive Distributions.

3.1 Binomial Uncertainty

In the example from the Benchmark Dose (BMD) Technical Guidance document (EPA 2000), shown in Table 1, the different doses are far apart, and the respective binomial distributions have little overlap.

Table 1: BMD Technical Guidance Example

Dose	Number of Subjects	Number of Responses
0	50	1
21	49	15
60	45	20

In such cases, we interpret the observable uncertainty as “binomial uncertainty.” For the above example, the observable uncertainty is three binomial distributions $Bin(50, 1/50)$, $Bin(49, 15/49)$ $Bin(45, 20/45)$. If we take this to represent our uncertainty for the results of repetitions of the above three experiments, then our uncertainty modeling should re-capture, or recover this uncertainty.

3.2 Isotonic Uncertainty³

It may happen that the doses in two experiments are close together, and that the observed percentage response for the *lower* dose is actually *higher* than the observed percentage response at the higher dose. Table 2 gives an example from the combined Frambozadrine data. The percentage response at dose 21 is lower than that at doses 1.2, 1.8, and 15.

Table 2: Combined Male-Female Data for Frambozadrine

Dose	0	1.2	1.8	15	21	82	109
Number exposed	95	45	49	44	47	47	48
Number of responses	5	6	5	4	3	24	33
% of responses	5.26	13.3	10.20	9.09	6.38	51.06	68.75

According to tent pole 3, we *know* that this decreasing relation is merely due to noise and does not reflect the relation between dose and response. We should therefore adapt our observational uncertainty to reflect the fact that probability of response is a non-decreasing function of dose. The method for doing this is called *isotonic regression*, and the maximum likely probabilities are found by the *Pooling Adjacent Violators* (PAV) algorithm (Barlow et al. 1972). We illustrate its use for the above data..

Suppose we draw five samples from the binomial distributions based the data in Table 2, for the lower five doses. For each sample value we estimate the binomial probability (with three responses in 95 trials, the estimated probability is $3/95 = 0.0316$). The results are shown in Table 3:

Table 3: Binomial Samples for Frambozadrine

Dose:	0	1.2	1.8	15	21	Estimated Binomial Probabilities				
Number exposed	95	45	49	44	47					
Number of responses	5	6	5	4	3					
Samples	5	5	6	5	6	0.0526	0.1111	0.1224	0.1136	0.1277
	3	4	5	4	3	0.0316	0.0889	0.102	0.0909	0.0638
	5	7	5	3	3	0.0526	0.1556	0.102	0.0682	0.0638
	5	4	4	4	2	0.0526	0.0889	0.0816	0.0909	0.0426
	6	4	7	6	3	0.0632	0.0889	0.1429	0.1364	0.0638

³ I am grateful to Eric Cator of TU Delft for suggesting this approach.

For each row of probabilities the PAV algorithm changes the probabilities to the most likely non-decreasing probabilities in the following way. In Table 3, move through each row of the probability matrix from left to right. At cell i , if its right neighbor is smaller, average these numbers and replace cells i and $i+1$ with this average. Now move from i to the left; if you find a cell j which is bigger than the current cell i , average all cells from j to $i+1$. Table 4 shows the successive PAV adjustments (in bold) to the binomial probabilities in the third row of Table 3.

Table 4: PAV Algorithm for First Row Probabilities of Table 3

	i=	1	2	3	4	5
	Dose	0	1.2	1.8	15	21
	Starting Probability	0.0526	0.1556	0.102	0.0682	0.0638
PAV	i=2	0.0526	0.1288	0.1288	0.0682	0.0638
PAV	i=3	0.0526	0.1086	0.1086	0.1086	0.0638
PAV	i=4	0.0526	0.0974	0.0974	0.0974	0.0974

3. Statistics as Usual (BMD)

Usual statistical methods aim at estimating the parameters of a best fitting model, where goodness of fit is judged with a criterion like the Akaike Information Criterion (AIC). A joint distribution on the parameters reflects the sampling fluctuations of the maximum likelihood estimates (MLEs), *assuming the model is true*.

Suppose we sample parameter values from their joint distribution (assume this distribution is joint normal with means and covariance obtained in the usual way). For each sample vector of parameters, we sample binomial distributions⁴ with probabilities computed from the sampled parameters with doses and with a number of subjects as given in the bioassay data (e.g., Table 1). Repeating this procedure we obtain a mixture of binomial distributions for each bioassay experiment. The statistics as usual (SAU) approach would regard these mixture binomial distributions as representing our uncertainty on the numbers of responses in each experiment. We take the best fitting model, assume it is true, consider the distribution of the maximum likelihood parameter estimates we would obtain if the experiments were repeated many times, and plug these into the appropriate binomial distributions. In the sequel, the distributions obtained in this way are called BMD distributions, as this approach is compatible with the Benchmark Dose Technical Guidance Document (EPA 2000); and the best fitting models and their MLE parameter distributions are obtained with the BMD software.

Do these mixture binomial distributions bear any resemblance to the binomial uncertainty distributions? We illustrate with the BMD data from the bench test. The preferred model in this case is the log logistic model with $\beta \geq 1$:

$$Prob(dose) = \gamma + (1-\gamma)/(1+e^{(-\alpha-\beta \ln(dose))})$$

⁴ Where possible, we used the normal approximation to the binomial.

Figure 2 shows the BMD cumulative distributions on the number of responses for the three dose levels (0, 21, 60) with number of experimental subjects as (50, 49, 45). Bd^* denotes the binomial uncertainty and id^* the isotonic uncertainty at dose level “*”. Since the doses are widely spread, bd and id are nearly identical.




Figure 2: Binomial (bd), Isotonic (id) and BMD Uncertainty Distributions; on the horizontal axis is number of animals responding, on the vertical axis is cumulative probability.

The BMD uncertainty for 22 or fewer responses in the experiment at dose 60 is about 0.45, and the binomial and isotonic uncertainties are about 0.79. Apparently, the SAU approach does not recover the binomial or isotonic observational uncertainty.

The BMD distributions are mixtures of binomials, so are the Bayesian predictive distributions. It may therefore be possible for a Bayesian to make a teleological choice of prior distributions which, after updating, would agree with the BMD distributions. Of course this choice would depend on the dose (contra tent pole 2) and on the preferred model. The practice of retrochoosing one’s prior is not unknown in Bayesian statistics, but is frowned upon. Indeed, it comes down to defining the target distribution as whatever it was we hit. We shall see that these BMD distributions may differ greatly from the isotonic distributions when the doses are close together. The isotonic distributions are not generally binomial.

The SAU approach is not wrong; it is simply solving a different problem, namely estimating parameters in a preferred model. That is not the same as quantifying observational uncertainty. The latter is the goal of uncertainty analysis, according to tent pole 1.

Overall, the best fitting model was the logistic, and we will concentrate BMD results with this model. Alternatives for Frambozadrine and Nectorine are presented in Appendix 1.

4. Probabilistic Inversion

We know what it means to invert a function at a point or point set. Probabilistic inversion denotes the operation of inverting a function at a distribution or set of distributions. Given a distribution on the range of a function, we seek a distribution over the domain such that pushing this distribution through the function returns the target distribution on the range. In dose-response uncertainty quantification, we seek a distribution over the parameters of a dose-response model which, when pushed through the DR model, recovers the observational uncertainty. The conclusion from Figure 2 is that the BMD distributions in Figure 2 are not the inverse of the observational distributions, for the log logistic model with $\beta \geq 1$.

Applicable tools for PI derive from the Iterative Proportional Fitting (IPF) algorithm (Kruithof 1937, Deming and Stephan 1940). A brief description is given below, for details see, e.g., Du et al. (2006), Kurowicka and Cooke (2006). In the present context, we start with a DR model and with a large sample distribution over the model's parameters. This distribution should be wide enough that its push-forward distribution covers the support of the observable uncertainty distributions⁵. If there are N samples of parameter vectors, each vector has probability $1/N$ in the sample distribution. We then re-weight the samples such that if we re-sample the N vector samples with these weights, we recover the observational distributions. In practice, the observational distributions are specified by specifying a number of quantiles or percentiles. In all the cases reported here, three quantiles of the observational distributions are specified, as close as possible to (5%, 50%, 95%). Technically speaking, we are thus inverting the DR model at a set of distributions, namely those satisfying the specified quantiles. Of course, specifying more quantiles would yield better fits in most cases, at the expense of larger sample distributions and longer run times.

The method for finding the weights for weighted resampling is IPF. IPF starts with a distribution over cells in a contingency table and finds the maximum likelihood estimate of the distribution satisfying a number of marginal constraints. Equivalently, it finds the distribution satisfying the constraints which is minimally informative relative to the starting distribution. It does this by iteratively adjusting the joint distribution.

The procedure is best explained with a simple example. Table 5a shows a starting distribution in a 2×2 contingency table. The “constraints” are marginal probabilities which the joint distribution should satisfy. The “results” are the actual marginal distributions. Thus, $0.011+0.022$

⁵ In the studies reported here, we begin with a wide uniform or loguniform distribution on the parameters, and conditionalize this on a small box containing the values of the observational uncertainty distributions. Each observational uncertainty distribution is based on 500 samples; in other words, on tables like Table 3, with 500 sample rows.

+ 0.32 = 0.065; and we want the sum of the probabilities in these cells to equal 0.5. In this example, the constraints (0.5, 0.3, 0.2) would correspond to the 50th and 80th percentiles; there is a 50% chance of falling beneath the median, a 30% chance of falling between the median and the 80th percentile, and a 20% chance of exceeding the 80th percentile. In the first step (Table 5b), each row is multiplied by a constant so that the row constraints are satisfied. In the second step (Table 5c), the columns of Table 5b are multiplied by a constant such that the column constraints are satisfied; the row constraints are now violated. The next step (not shown) would re-satisfy the row constraints. After 200 such steps we find the distribution in Table 5d. The zero cells are shaded; each iteration has the same zero cells as the starting distribution, but the limiting distribution may have more zeros.

Table 5a: Starting Distribution for IPF

Result	Constraint			
0.065	0.500	0.011	0.022	0.032
0.097	0.300	0.054	0.043	0.000
0.838	0.200	0.000	0.000	0.838
	Constraint	0.100	0.200	0.700
	Result	0.065	0.065	0.870

Table 5b: First Iteration, Row Projection

Result	Constraint			
0.500	0.500	0.083	0.167	0.250
0.300	0.300	0.167	0.133	0.000
0.200	0.200	0.000	0.000	0.200
	Constraint	0.100	0.200	0.700
	Result	0.250	0.300	0.450

Table 5c: Second Iteration, Column Projection

Result	Constraint			
0.533	0.500	0.033	0.111	0.389
0.156	0.300	0.067	0.089	0.000
0.311	0.200	0.000	0.000	0.311
	Constraint	0.100	0.200	0.700
	Result	0.100	0.200	0.700

Table 5d: Result after 200 Iterations

Result	Constraint			
0.501	0.500	0.000	0.002	0.499
0.298	0.300	0.100	0.198	0.000
0.201	0.200	0.000	0.000	0.201
	Constraint	0.100	0.200	0.700
	Result	0.100	0.200	0.700

To apply this algorithm to the problem at hand, we convert the N samples of parameter vectors and the specified quantiles of the observational distributions into a large contingency table. With 3 specified quantiles, there are 4 inter-quantile intervals, for each observational distribution. If there are 3 doses, there are $4^3 = 64$ cells in the contingency table. Let NR_i be the number of animals exposed to dose d_i , and let $P(\alpha, \beta, \gamma, \dots, d_i)$ be the probability of response if the parameter values are $(\alpha, \beta, \gamma, \dots)$ and the dose is d_i . The initial probability in each cell is proportional to the number of sample parameter vectors $(\alpha, \beta, \gamma, \dots)$ for which $NR_i \times P(\alpha, \beta, \gamma, \dots, d_i)$; $i = 1, 2, 3$ lands in that cell. IPF is now applied and the final cell probabilities are proportional to the weights for the sample vectors in that cell.

Csiszar (1975) (finally⁶) proved that *if* there is a distribution satisfying the constraints whose zero cells include the zero cells of the starting measure, then IPF converges, and it converges to the minimal information (or maximum likelihood) distribution relative to the starting measure which satisfies the constraints. If the starting distribution is not feasible in this sense, then IPF does not converge. Much work has been devoted to variations of the IPF algorithm with better convergence behavior. An early attempt, termed PARFUM (PARAmeter Fitting for Uncertain Models, Cooke 1994), involved simultaneously projecting onto rows and columns, averaging the result. It has recently been shown (Matus 2007, Vomlel 1999) that this and similar adaptations always converge, and that PARFUM converges to a solution if the problem is feasible (Du et al 2006). In fact, geometric averaging seems to have some advantages over arithmetic averaging; but PARFUM is used in this report, in case of infeasibility. In case of infeasibility PARFUM is largely – though not completely – determined by the support of the starting measure.

The best fitting distribution in the sense of AIC is not always the best distribution to use for probabilistic inversion. Figure 3 illustrates this for the BMD bioassay data. At each dose level $*$, it shows the binomial observational uncertainty (bd $*$) the isotonic observational uncertainty (id $*$), the BMD distributions from Figure 2, and also the result of PI on the log logistic model with $\beta = 1$ (corresponding to the best fitting model) and PI on the log logistic model with $\beta > 0$.

For the figures throughout this report:

- Binomial distributions (bd $*$) are red
- Isotonic distributions (id $*$) are blue
- PI distributions (d $*$ pi) are black
- BMD distributions (BMD $*$) are green.

(Note that while parameters are usually given in Greek letters, the graphics software used does not support this. Therefore common letters are used, e.g., the background response parameter γ is rendered as g , and the intercept parameter α is rendered as a .)

⁶ There had been many incorrect proofs, and proofs of special cases, see Fienberg (1970), Kullback (1986, 1970), Haberman (1974) and Ireland and Kullback (1968). More general results involving the “Bregman distance” were published earlier in Russian (Bregman 1967), for a review see (Censor and Lent, 1981). A simple sketch of Csiszar’s proof is in Csiszar (undated).




Figure 3: PI Distributions for the BMD Bench Test, with Log Logistic Model and $\beta=1$ (orange) and $\beta>0$ (black). Number of responses are on the horizontal axis, cumulative probability on the vertical.

Benchmark Dose

Computing the Benchmark dose for the extra risk model is easy in this case. We solve the equation

$$\text{BMR} = (\text{Prob}(\text{BMD}) - \text{Prob}(0)) / (1 - \text{Prob}(0))$$

for values $\text{BMR} = 0.1, 0.05, 0.01$. $\text{Prob}(\text{BMD})$ is probability of response at dose BMD, and it depends on the parameters of the DR model, in this case the log logistic model. $\text{Prob}(0)$ is the background probability, and is taken to be the observed percentage response at zero dose. In this case, the log logistic model can be inverted and we may write

$$\text{BMD} = (1 + e^{-\alpha - \beta \ln(\text{dose})})^{-1}$$

Sampling values from (α, β) after PI, we arrive at the following values:

Table 6: BMD Tech. Guidance; Benchmark dose calculations

BMR	BMD					BMD50/BMDL
	Mean	Variance	5% Perc	50% Perc	95% Perc	
0.1	3.12E+00	4.52E+00	8.49E-01	2.53E+00	7.54E+00	2.98E+00
0.05	9.31E-01	4.85E-01	2.10E-01	7.19E-01	2.38E+00	3.42E+00
0.01	6.70E-02	3.78E-03	9.34E-03	4.64E-02	1.91E-01	4.97E+00

5. Frambozadrine: Threshold Model

The data for this case is reproduced below as Table 7.

Table 7: Frambozadrine Data

	Dose (mg/kg-day)	Total Number of Rats	Hyperkeratosis
Male			
	0	47	2
	1.2	45	6
	15	44	4
	82	47	24
Female			
	0	48	3
	1.8	49	5
	21	47	3
	109	48	33

We see strong non-monotonicities for male and female rats. A higher percentage of male rats respond at dose 1.2 than at dose 15, and similarly for female rats at doses 1.8 and 21 [mg/kg-day]. The logistic model fit best according to the AIC.

For the male and female tests, analyzed individually, a standard multistage model with linear parameters gave best results, among the standard models (adding quadratic and cubic terms did not improve the results). This model is:

$$Nr*(gamma+(1-gamma)*(1-exp(-b1*dose)))$$

The results for these two cases are shown below.





Figure 4: Fitting Results for Frambozadrine Data for Males

Note the differences between the binomial observational uncertainties (bd) and isotonic observational uncertainties (id) for doses 1.2 and 15. The PI distributions are fit to the id distributions, and are closer than the BMD distributions, but the fit is bad for the lower doses. A similar picture emerges for female rats.






Figure 5: Fitting Results for Frambozadrine Data for Females

If dose really does scale as $[mg/kg\text{-day}]$, and barring countervailing toxicological evidence, we should be able to pool these data. However, the problems of fitting a multistage model, or any other models in the BMD library, became more severe. The pooled data seems to indicate a response plateau. Beneath a certain threshold s the background probability g applies. Above another threshold t , the dose follows a multistage model. The threshold model is:

$$Nr*(g+(1-g)*il\{s,dose,\infty\}*(1-\exp(-b1*\max\{t,dose\}-b2*\max\{t,dose\}^2-b3*\max\{t,dose\}^3)))$$

Here $il\{X<Y<Z\}$ is a random indicator function which returns 1 if $X \leq Y \leq Z$, and 0 otherwise. The starting distributions for the parameters in this model were:

$$g \sim \text{uniform}[0.005, 0.1]$$

$$s \sim \text{uniform}[0, 21]$$


$$t \sim (21 - s)$$

$$b1 \sim \text{uniform}[0, 0.01]$$

$$b2 \sim \text{uniform}[0, 0.0001]$$

$$b3 \sim \text{loguniform}[10^{-13}, 5 \times 10^{-6}]$$

These distributions were conditionalized on values falling within a box around the observable uncertainty distributions. The results shown in Figure 6 indicate a decent fit.






Figure 6: Results for Fitting Frambozadrine, Pooled Data

Note the strong differences between the bd and id distributions, and note that the BMD distributions are close to neither, except for the higher dose levels. The parameter distributions after PI are shown in Figure 7. The lower threshold s concentrates on values near 0, while the higher threshold t concentrates on values near 21. The distributions of $b1$ and $b2$ are strongly affected by the PI; the distribution of $b3$ is less affected, indicating that this parameter is less important to the fitting.




Figure 7: Parameter Densities for Frambozadrine MF after PI

Figure 8 shows the parameter distributions before PI. The thresholds s and t are most affected by the fitting, $b3$ is changed very little.




Figure 8: Parameter Densities for Frambozadrine MF, before PI

The joint distribution of all parameters in the threshold model, after PI, is shown in Figure 9 as a cobweb plot. Taking 172 samples, each sample is a value; connecting each sample with a jagged line, the picture in Figure 8 emerges. We see strong negative correlations between $b1, b2, b3$, as well as the concentration of t toward values near 21.




Figure 9: Cobweb Plot for Parameters for Frambozadine, Male and Female Data

Once we have a joint distribution over the model parameters, we can easily compute the uncertainty distributions (smoothed) for probabilities at arbitrary doses, see Figure 10. The distributions tend to get wider as dose increases up to 100.

These results show that a good fit to isotonic observational uncertainty is possible with a threshold model. The same may hold for other more exotic models. Of course, the toxicological plausibility of such models must be judged by toxicologists, not by mathematicians.




Figure 10: Uncertainty Distributions for Probability of Response as Function of Dose; p^* is the Probability of Response at Dose $*$. Probability of response is on the horizontal axis.

Computing the Benchmark dose was complicated in this case by the fact that the DR model could not be analytically inverted, in contrast to the previous example. We therefore proceed as follows. For any value of the parameters, we can compute the dose d^* which satisfies

$$\text{BMR} = (\text{Prob}(d^*) - \text{Prob}(0)) / (1 - \text{Prob}(0))$$

where again $\text{Prob}(0)$ is taken to be the observed background rate, 0.0526. Consider the cumulative distribution function of d^* ; if a number d_L is the 5%-tile of this distribution, this means that with probability 0.95

$$\text{Prob}(d_L) < \text{BMR}(1 - \text{Prob}(0)) + \text{Prob}(0).$$

Hence, to find the 5% value of d^* , we must find the value d_L for which $\text{BMR}(1 - \text{Prob}(0)) + \text{Prob}(0)$ is the 95%-tile. This is BMD_L , and we compare it to the value, denoted BMD , for which $\text{BMR}(1 - \text{Prob}(0)) + \text{Prob}(0)$ is the median. The results are shown in Table 8. Values for $\text{BMR} = 0.01$ were not obtainable, as the 95% of the background parameter g was already larger than $0.01 \times (1 - 0.0526) + 0.0526 = 0.062$.

Table 8: Frambozadrine, Benchmark Dose Calculations

Frambozadrine			
BMR	BMDL	BMD50	BMDL/BMD50
0.1	1.68E+01	2.35E+01	1.40E+00
0.05	1.30E+00	1.53E+01	1.18E+01
0.01	-----	-----	-----

6. Nectorine

The data for this case is shown in Table 9 below.

Table 9: Nectorine Data

	Concentration (ppm)			
	0	10	30	60
Lesion	# response / # in trial			
Respiratory epithelial adenoma in rats	0/49	6/49	8/48	15/48
Olfactory epithelial neuroblastoma in rats	0/49	0/49	4/48	3/48

As we are unable at present to analyze zero response data, the data do not support an analysis of the neuroblastoma lesions, and we consider only pooling. The text accompanying these data indicates that these two lesions occur independently. Therefore, the probability of showing either lesion is given by $P_{a \text{ or } n} = P_a + P_n - P_a P_n$; where a and n denote adenoma and neuroblastoma respectively. These two endpoints are taken to be distinct (with summation of risks from multiple tumor sites when tumor formation occurs independently in different organs or cell types considered superior to the calculation of risk from individual tumor sites alone).

Under these assumptions, the bioassay for either lesion is shown in Table 8.

Dose	10	30	60
Number of rats	49	48	48
Number responding	6	12	18

The model used for PI was multistage with 2 parameters:

$$NR*(g+(1-g)*(1-\exp(-b1*dose)))$$

The results are shown in Figure 12. In this case both PI and BMD yield decent fits.




Figure 11: Results for Fitting Nectorine

The calculation of the Benchmark dose and the ratio BMD / BMDL proceeded as in the case of the BMD Technical Guidance example. The results are shown in Table 10 below:

Table 10: Nectorine: Benchmark Dose Calculation

Nectorine, Adenoma or Neuroblastoma						
	BMD					
BMR	Mean	Variance	5% Perc	50% Perc	95% Perc	BMDL/BMD50
0.1	1.72E+01	3.44E+01	1.02E+01	1.58E+01	2.94E+01	1.55E+00
0.05	8.43E+00	8.21E+00	4.98E+00	7.74E+00	1.44E+01	1.55E+00
0.01	1.64E+00	3.12E-01	9.70E-01	1.51E+00	2.80E+00	1.56E+00

7. Persimonate: Barrier Model

The data are shown in Table 11.

Table 11: Persimonate Data

	Continuous Equivalent Dose	Total Metabolism (mg/kg-day)	Survival-Adjusted Tumor Incidence
B6C3F1 male mice inhalation	0	0	17/49
	18ppm	27	31/47
	36ppm	41	41/50
Crj:BDF1 male mice inhalation	0	0	13/46
	1.8ppm	3.4	21/49
	9.0ppm	14	19/48
	45 ppm	36	40/49

Here again there is non-monotonicity. The issues that arise here are similar to those with Frambozadrine, and the analysis will be confined to the pooled data, assuming that the two types of mice can be pooled. The doses used are in units of *ppm*. The PI with the standard DR models from the BMD library was unsuccessful in this case. The data suggests that there may be 2 or 3 plateaus of response. This might arise if there were successive biological barriers which, when breached, cause increasing probabilities of response. After reaching the first barrier, additional dose has no effect until the second barrier is breached, etc. A barrier model was explored, using 2 barriers:

$$NR*(g+(1-g)*(1-\exp(-b1*i\{0,dose,t\}-b2*i\{t,dose,s\}-b3*i\{s,dose,\infty\})))$$

The variable and function distributions are reproduced below from Appendix 2.

Random Variables

```

Random variable name: g Distribution type: Uniform
Parameters      Main quantiles      Moments
a = 0.200      5%q = 0.210      Mean = 0.300
b = 0.400      50%q = 0.300     SDev = 0.058
               95%q = 0.390

Random variable name: a1 Distribution type: Uniform
Parameters      Main quantiles      Moments
a = 0.100      5%q = 0.145      Mean = 0.550
b = 1.000      50%q = 0.550     SDev = 0.260
               95%q = 0.955

Random variable name: a2 Distribution type: Uniform
Parameters      Main quantiles      Moments
a = 0.100      5%q = 0.145      Mean = 0.550
b = 1.000      50%q = 0.550     SDev = 0.260
               95%q = 0.955

Random variable name: a3 Distribution type: Uniform
Parameters      Main quantiles      Moments
a = 0.100      5%q = 0.145      Mean = 0.550
    
```

$b =$	1.000	50%q =	0.550	SDev =	0.260
		95%q =	0.955		
Random variable name: T Distribution type: Uniform					
Parameters		Main quantiles		Moments	
$a =$	0.000	5%q =	1.000	Mean =	10.000
$b =$	20.000	50%q =	10.000	SDev =	5.774
		95%q =	19.000		
Random variable name: s Distribution type: Uniform					
Parameters		Main quantiles		Moments	
$a =$	20.000	5%q =	21.250	Mean =	32.500
$b =$	45.000	50%q =	32.500	SDev =	7.217
		95%q =	43.750		

Formulas

1. $d_0: 95 * g$
2. $d_{1.8}: 49 * (g + (1-g) * (1 - \exp(-b_1 * i_1\{0, 1.8, t\} - b_2 * i_1\{t, 1.8, s\} - b_3 * i_1\{s, 1.8, >>\})))$
3. $d_9: 48 * (g + (1-g) * (1 - \exp(-b_1 * i_1\{0, 9, t\} - b_2 * i_1\{t, 9, s\} - b_3 * i_1\{s, 9, >>\})))$
4. $d_{18}: 47 * (g + (1-g) * (1 - \exp(-b_1 * i_1\{0, 18, t\} - b_2 * i_1\{t, 18, s\} - b_3 * i_1\{s, 18, >>\})))$
5. $d_{36}: 50 * (g + (1-g) * (1 - \exp(-b_1 * i_1\{0, 36, t\} - b_2 * i_1\{t, 36, s\} - b_3 * i_1\{s, 36, >>\})))$
6. $d_{45}: 49 * (g + (1-g) * (1 - \exp(-b_1 * i_1\{0, 45, t\} - b_2 * i_1\{t, 45, s\} - b_3 * i_1\{s, 45, >>\})))$
7. $b_1: a_1$
8. $b_2: b_1 + a_2$
9. $b_3: b_2 + a_3$
10. $condition: i_1\{24, d_0, 37\} * i_1\{16, d_{1.8}, 25\} * i_1\{16, d_9, 42\} * i_1\{21, d_{18}, 36\} * i_1\{34, d_{36}, 47\} * i_1\{36, d_{45}, 48\}$

Here, t and s are random barriers with $t < s$, and $i_1\{X, Y, Z\}$ is a random indicator function returning 1 if $X \leq Y \leq Z$ and 0 otherwise. The parameters b_1, b_2 and b_3 are random but increasing: $b_1 < b_2 < b_3$. Thus, doses less or equal to t have coefficient b_1 , those between t and s have b_2 , and those greater or equal to s have b_3 .

In fitting this barrier model, the PI chooses maximum likely weights for resampling the starting distribution, and thus chooses a maximum likely distribution for the barriers and their associated coefficients, based on the starting sample distribution. The results are shown in Figure 12.

Here again, we see pronounced differences between the binomial observational uncertainty and the isotonic observational uncertainty. Again, we see decent agreement between the id and pi distributions, while the BMD distributions don't agree with either bd or id, except at the highest dose (45).

The parameter distributions for the barrier model, after PI, are shown in Figure 13. Graph (ii) shows the graphs for the thresholds t and s . The starting distributions were uniform on $[0, 21]$ and $[21, 45]$ respectively. The distribution of t concentrates near 21 and that of s concentrates between 21 and 36 (there is no dose between 21 and 36). The cobweb plot, graph (i) shows less interaction between b_1, b_2, b_3 than in Figure 8 for Frambozadrine. Figure 14 shows the starting densities for b_1, b_2, b_3 ; comparing these with Figure 13 (iii) shows the action of the PI.




Figure 12: Results for Fitting Persominate, Pooled Data




Figure 13: Distributions after PI from upper left to bottom right: (i) all parameters, (ii) t and s (smoothed), (iii) $b1$, $b2$, $b3$, and (iv) g .




Figure 14: Starting Distributions for $b1$, $b2$, $b3$




Figure 15: Persimonate (smoothed) Probability as Function of Dose; p^* is the Probability of Response a Dose*. Probabilities of response are on the horizontal axis.

The difference between the barrier model for Persimonate and the threshold model for Frambozadrine can be appreciated by comparing the graphs for uncertainty in probability as a function of dose (Figures 10 and 15). Whereas Figure 10 showed a smooth shift of the probability distributions as dose increased, Figure 15 reflects the uncertainty of the positions of the barriers. Doses 0, 5, and 10 fall beneath the first barrier with high probability and their probabilities of response are similar. The same holds for doses 35, 40, 45 with high probability they fall above the upper barrier. For intermediate doses, the uncertainty in the probability of response reflects the uncertainty in the positions of s and t .

As in the case of Frambozadrine, this exercise shows that a good fit to isotonic observable uncertainty is possible. The toxicological plausibility of such models is a matter for toxicologists.

The Benchmark dose was not stable in this case, owing to the form of the DR model. In the region of interest, the 95%-tile of the probability of response exhibited discontinuities making it impossible to find a dose whose 95%-tile was equal to $BMR \times (1 - \text{Prob}(0)) + \text{Prob}(0)$.

8. Concluding Remarks

The quantification of uncertainty in DR models can only be judged if there is some external, observable uncertainty which this quantification should recover. We have used the isotonic uncertainty distributions for this purpose. With PI it is possible to find distributions over model

parameters which recover these observable distributions. For the example from the BMD technical guidance document and for Nectorine, PI was successful with standard DR models. For Frambozadrine and Persimonate it was necessary to introduce threshold and barrier models. With these models, decent fits with PI were obtained. Better fits could be obtained by stipulating more than 3 quantiles in the observable distributions, and perhaps by exploring different models. The point of this analysis is to show that good fits to observable uncertainty distributions are achievable.

The SAU approach, as reflected in the BMD distributions does not recover observable uncertainty as reflected either in binomial uncertainty distributions or isotonic uncertainty distributions. This is not surprising, as the SAU approach aims at quantifying uncertainty in the parameters of a model assumed to be true, it is not aimed at capturing observable uncertainty.

The number of parameters in the threshold and barrier models for Frambozadrine and Persimonate is larger than the corresponding BMD logistic distributions. Is this not “stacking the deck”? A number of remarks address this question.

- (i) Indeed, the AIC criterion for goodness of fit punishes rather severely for additional parameters, and this is partly responsible for the discrepancy between the BMD and the observable uncertainties (esp. in Figure 2).
- (ii) Simply adding more parameters to a monotonic model, like the logistic, probit or multistage, will not produce better fits for the Frambozadrine and Persimonate data; it seems more important to get a right model type than to add parameters to a wrong model. We are fitting 3 quantiles per dose, thus for Persimonate there are $3 \times 6 = 18$ quantiles to be fit, whereas the barrier model has 6 parameters. Bayesian models, by comparison, have many more parameters.
- (iii) The fits are judged – by occlusion – with respect to the whole distributions, not just the 3 fitted quantiles.
- (iv) The question of judging unimportant parameters and finding a parsimonious model within the PI approach is a subject which needs more attention.

Having addressed the problem of capturing observational uncertainty, we can move on to more difficult problems such as extrapolation of animal data to humans, allometric scaling, sensitive subgroups, the use of incomplete data including human data from accidental releases, and extrapolation to low dose. Approaches to these more difficult problems should be grounded in methods that are able to deal with the relatively “clean” problems typified by these bench test exercises.

The PI approach is based on antecedently defined observable uncertainty distributions. These may be based on isotonic regression, as was done here, but they may also incorporate distributions from structured expert judgment or from anecdotal incident data. This may provide a method for attacking the more difficult extrapolation issues identified above.

The calculations and simulations displayed here are based on 500 samples for the isotonic uncertainty distributions, and 10,000 to 20,000 samples for the probabilistic inversion. The processing was done with the uncertainty analysis package UNICORN available free from <http://dutiosc.twi.tudelft.nl/~risk/>. The probabilistic inversion software used here is a significant

improvement over the version currently available at this website, but is not yet ready for distribution. After initial set-up, processing time per case is in the order of a few minutes. More time is involved in exploring different models and different starting distributions.

APPENDIX 1: Alternative models and optimization strategies


Frambozadrine log probit

For Frambozadrine, pooled data, the logistic model had an AIC of 290.561. The log probit model was slightly better at 290.396. The multistage model with 2nd degree polynomial was better still with AIC 289.594, but the standard deviations of the parameters were not calculated with the BMD software. The log probit model is:

$$Prob(dose) = \gamma + (1-\gamma)\Phi(\alpha + \beta \times \ln(dose))$$

Where Φ is the cdf of the standard normal distribution.

Figure 16 compares the observable uncertainty distributions for the log probit (lprob) and logistic model (BMD) models, together with the binomial and isotonic uncertainties. The (log) probit models tend to be wider, a feature noted on many exercises (not reported here).






Figure 16: Logistic (BMD), log probit (lprob), binomial (bd) and isotonic (id) uncertainties for Frambozadrine, male and female,

Nectorine Probit

For the Nectorine example, the logistic model had an AIC of 158.354, whereas the probit model had 158.286. Log logistic and log probit models were slightly worse. The probit model is:

$$Prob(dose) = \Phi(\alpha + \beta \times dose); \quad \Phi \text{ the standard normal CDF.}$$

Figure 17 compares the previously obtained results with those of the probit model. As in Figure 16, we see that the probit model tends to have higher variance than the others.




Figure 17: Nectorine, pooled, binomial (bd) isotonic (id) BMD, probabilistic inversion (pi) and probit (probit)

Optimization Strategies

Probabilistic inversion is an optimization strategy applied to quantiles of the observable uncertainties. As with all optimization strategies in multi-extremal problems, its solution may depend on the starting point. Compared to other strategies, it is more labor intensive.

A standard optimization routine might work as follows.

1. Choose a preferred, smooth model.
2. Sample from the binomial uncertainty distributions and isotonicize this sample (as in Table 4).
3. Find values of the parameters for the preferred model which minimize $\sum_i (NR_i \times P_m(d_i) - NR_i \times P_{is}(d_i))^2$, where NR_i is the number of animals given dose d_i ; $P_m(d_i)$ is the probability of response given by the model at dose d_i , and $P_{is}(d_i)$ is the isotonic probability of response at dose d_i .
4. Store these parameters, and repeat steps 2 and 3.

Figure 18 compares the results of optimizing⁷ the log logistic model (lglgopt), $\beta \geq 1$ with the BMD distributions - which are also based on the log logistic model, but use the MLE parameter distributions given by the BMD software. The binomial and isotonic uncertainties are also

⁷ The EXCEL solver with default settings was used for all the optimization results.

shown. The optimization results in a somewhat different, but not overwhelmingly better fit to the isotonic uncertainties.




Figure 18: BMD example with binomial (bd) isotonic (is) bmd, and log logistic with optimization (lglgopt) with $\beta \geq 1$

Whereas the previous example concerned a case where a smooth model (log logistic) yielded a good fit with PI, the Persimonte case did not yield a good fit with any smooth model. The Barrier model used in Figure 12 is not differentiable. This means that the regularity conditions for the asymptotic normality, and even for the strong consistency of the maximum likelihood estimators do not hold (Cox and Hinkley, p.281, 288)

In Figure 19 we show compare optimization applied to the logistic and to the log logistic models with the other alternatives of Figure 12. The log logistic ($\beta \geq 0$) model has three parameters, whereas the logistic model has two. Figure 19 shows that this extra parameter does not significantly improve the fits, and that the attempt to minimize square difference to the isotonic uncertainty does not produce better results than simply using the logistic model with MLE parameter distributions (BMD).




Figure 19: Persimionate results of Figure 12 compared with optimization of logistic (lgstop) and log logistic (llgstop)

APPENDIX 2: Start distributions and calculation script

BMD

Random variables

Random variable name: a

Distribution type: Uniform

Parameters

a = -4.000

Main quantiles

5%q = -3.850

Moments

Mean = -2.500

b = -1.000 50%q = -2.500 SDev = 0.866
 95%q = -1.150

Random variable name: b

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.500	5%q = 0.510	Mean = 0.600
b = 0.700	50%q = 0.600	SDev = 0.058
	95%q = 0.690	

Random variable name: g

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.000	5%q = 0.030	Mean = 0.300
b = 0.600	50%q = 0.300	SDev = 0.173
	95%q = 0.570	

Formulas

1. id0: $50 * g$
2. id21: $49 * (g + (1 - g) / (1 + \exp(-a - b * \ln(21))))$
3. id60: $45 * (g + (1 - g) / (1 + \exp(-a - b * \ln(60))))$
4. condition: $i1\{0, id0, 4\} * i1\{6, id21, 25\} * i1\{12, id60, 31\}$

Frambozadrine Female

Random variables (Note, b3 == 0)

Random variable name: GAMMA

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.001	5%q = 0.016	Mean = 0.150
b = 0.300	50%q = 0.150	SDev = 0.086
	95%q = 0.285	

Random variable name: b1

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.000	5%q = 0.001	Mean = 0.005
b = 0.010	5q = 0.005	SDev = 0.003
	95%q = 0.010	

Random variable name: b3

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.000	5%q = 0.000	Mean = 0.000
b = 0.000	50%q = 0.000	SDev = 0.000
	95%q = 0.000	

Formulas

1. d0: $48 * (\text{gamma})$
2. d1.8: $49 * (\text{gamma} + (1 - \text{gamma}) * (1 - \exp(-b1 * 1.8 - b3 * 1.8^3)))$
3. d21: $47 * (\text{gamma} + (1 - \text{gamma}) * (1 - \exp(-b1 * 21 - b3 * 21^3)))$
4. d109: $48 * (\text{gamma} + (1 - \text{gamma}) * (1 - \exp(-b1 * 109 - b3 * 109^3)))$
5. condition: $i1\{0, d0, 7\} * i1\{1, d1.8, 7\} * i1\{1, d21, 7\} * i1\{24, d109, 41\}$

Frambozadrine Male

Random variables (Note, b2 == 0)

Random variable name: GAMMA

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.001	5%q = 0.016	Mean = 0.150
b = 0.300	50%q = 0.150	SDev = 0.086
	95%q = 0.285	

Random variable name: b1

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.000	5%q = 0.003	Mean = 0.025
b = 0.050	50%q = 0.025	SDev = 0.014
	95%q = 0.048	

Random variable name: b2

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.000	5%q = 0.000	Mean = 0.000
b = 0.000	50%q = 0.000	SDev = 0.000
	95%q = 0.000	

Formulas

1. do: $47 * (\text{gamma})$
2. d1.2: $45 * (\text{gamma} + (1 - \text{gamma}) * (1 - \exp(-b1 * 1.2 - b2 * 1.2^2)))$
3. d15: $44 * (\text{gamma} + (1 - \text{gamma}) * (1 - \exp(-b1 * 15 - b2 * 15^2)))$
4. d82: $47 * (\text{gamma} + (1 - \text{gamma}) * (1 - \exp(-b1 * 82 - b2 * 82^2)))$
5. condition: $i1\{0, \text{do}, 7\} * i1\{0, \text{d1.2}, 9\} * i1\{2, \text{d15}, 9\} * i1\{15, \text{d82}, 34\}$

Frambozadrine Male & Female

Random variables

Random variable name: g

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.005	5%q = 0.010	Mean = 0.053
b = 0.100	50%q = 0.053	SDev = 0.027
	95%q = 0.095	

Random variable name: b1

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.000	5%q = 0.001	Mean = 0.005
b = 0.010	50%q = 0.005	SDev = 0.003
	95%q = 0.010	

Random variable name: b2

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.000	5%q = 0.000	Mean = 0.000
b = 0.000	50%q = 0.000	SDev = 0.000
	95%q = 0.000	

Random variable name: b3

Distribution type: LogUniform

Parameters	Main quantiles	Moments
a = 0.000	5%q = 0.000	Mean = 0.000
b = 0.000	50%q = 0.000	SDev = 0.000
	95%q = 0.000	

Random variable name: s

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.000	5%q = 1.050	Mean = 10.500
b = 21.000	50%q = 10.500	SDev = 6.062
	95%q = 19.950	

Formulas

1. condition: $i1\{0, \text{d0}, 10\} * i1\{0, \text{d1.2}, 7\} * i1\{1, \text{d1.8}, 8\} * i1\{1, \text{d15}, 8\} * i1\{2, \text{d21}, 8\}$

```

*il{24,d109,42}*il{14,d82,34}
2. d0: 95*g
3. d1.2: 45*(g+(1-g)*il{s,1.2,>>}*(1-exp(-b1*max{t,1.2}-b2*max{t,1.2}^2-
b3*max{t,1.2}^3)))
4. d1.8: 49*(g+(1-g)*il{s,1.8,>>}*(1-exp(-b1*max{t,1.8}-b2*max{t,1.8}^2-
b3*max{t,1.8}^3)))
5. d15: 44*(g+(1-g)*il{s,15,>>}*(1-exp(-b1*max{t,15}-b2*max{t,15}^2-b3*max{t,15}^3)))
6. d21: 47*(g+(1-g)*il{s,21,>>}*(1-exp(-b1*max{t,21}-b2*max{t,21}^2-b3*max{t,21}^3)))
7. d82: 47*(g+(1-g)*il{s,82,>>}*(1-exp(-b1*max{t,82}-b2*max{t,82}^2-b3*max{t,82}^3)))
8. d109: 48*(g+(1-g)*il{s,82,>>}*(1-exp(-b1*max{t,109}-b2*max{t,109}^2-
b3*max{t,109}^3)))
9. t: 21-s

```

Nectorine

Random variables

Random variable name: g

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.000	5%q = 0.005	Mean = 0.050
b = 0.100	50%q = 0.050	SDev = 0.029
	95%q = 0.095	

Random variable name: b1

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.000	5%q = 0.002	Mean = 0.015
b = 0.030	50%q = 0.015	SDev = 0.009
	95%q = 0.029	

Formulas

```

1. d10: 49*(g+(1-g)*(1-exp(-b1*10)))
2. d30: 48*(g+(1-g)*(1-exp(-b1*30)))
3. d60: 49*(g+(1-g)*(1-exp(-b1*60)))
4. condition: il{1,d10,10}*il{2,d30,18}*il{8,d60,26}

```

Persimonate

Random variables

Random variable name: g

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.200	5%q = 0.210	Mean = 0.300
b = 0.400	50%q = 0.300	SDev = 0.058
	95%q = 0.390	

Random variable name: a1

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.100	5%q = 0.145	Mean = 0.550
b = 1.000	50%q = 0.550	SDev = 0.260
	95%q = 0.955	

Random variable name: a2

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.100	5%q = 0.145	Mean = 0.550
b = 1.000	50%q = 0.550	SDev = 0.260
	95%q = 0.955	

Random variable name: a3

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.100	5%q = 0.145	Mean = 0.550
b = 1.000	50%q = 0.550	SDev = 0.260
	95%q = 0.955	

Random variable name: T

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 0.000	5%q = 1.000	Mean = 10.000
b = 20.000	50%q = 10.000	SDev = 5.774
	95%q = 19.000	

Random variable name: s

Distribution type: Uniform

Parameters	Main quantiles	Moments
a = 20.000	5%q = 21.250	Mean = 32.500
b = 45.000	50%q = 32.500	SDev = 7.217
	95%q = 43.750	

Formulas

- d0: $95 \cdot g$
- d1.8: $49 \cdot (g + (1-g) \cdot (1 - \exp(-b1 \cdot i1\{0, 1.8, t\} - b2 \cdot i1\{t, 1.8, s\} - b3 \cdot i1\{s, 1.8, >>\})))$
- d9: $48 \cdot (g + (1-g) \cdot (1 - \exp(-b1 \cdot i1\{0, 9, t\} - b2 \cdot i1\{t, 9, s\} - b3 \cdot i1\{s, 9, >>\})))$
- d18: $47 \cdot (g + (1-g) \cdot (1 - \exp(-b1 \cdot i1\{0, 18, t\} - b2 \cdot i1\{t, 18, s\} - b3 \cdot i1\{s, 18, >>\})))$
- d36: $50 \cdot (g + (1-g) \cdot (1 - \exp(-b1 \cdot i1\{0, 36, t\} - b2 \cdot i1\{t, 36, s\} - b3 \cdot i1\{s, 36, >>\})))$
- d45: $49 \cdot (g + (1-g) \cdot (1 - \exp(-b1 \cdot i1\{0, 45, t\} - b2 \cdot i1\{t, 45, s\} - b3 \cdot i1\{s, 45, >>\})))$
- b1: a1
- b2: b1+a2
- b3: b2+a3
- condition: $i1\{24, d0, 37\} \cdot i1\{16, d1.8, 25\} \cdot i1\{16, d9, 42\} \cdot i1\{21, d18, 36\} \cdot i1\{34, d36, 47\} \cdot i1\{36, d45, 48\}$

References

- Barlow, R. E. Bartholomew, D. J., Bremner, J. M. and Brunk, H. D. (1972), *Statistical Inference Under Order Restrictions*, New York, Wiley.
- Benchmark Dose Technical Guidance Document (External Review Draft). 2000.
- Bregman, L.M. (1967) The relaxation method to find the common point of convex sets and its applications to the solution of problems in convex programming, *USSR Computational Mathematics and Mathematical Physics*, 7, 200-217.
- Censor, Y. and Lent, A (1981). An iterative row-action method for interval convex programming, *Journal of Optimization Theory and Applications*, vol. 34, No. 3.
- Cooke, R. (1994) "Parameter fitting for uncertain models: modeling uncertainty in small models", *Reliability Engineering and System Safety*, vol 44, 89-102.
- Cox, D.R. and Hinkley, D.V., (1974) *Theoretical Statistics*, Chapman and Hall, London.
- Csiszar, I and Tusnady (1984) Information geometry and alternating minimization procedures, *Statistics and Decisions*, Supplement Issue No.1, 205-237, R. Oldenbourg Verlag, Munchen.
- Csiszar, I., (1975). I-divergence geometry of probability distributions and minimization problems. *Ann. Probab.* 3, 146–158.
- Csiszar, I. (undated) *Information Theoretic Methods in Probability and Statistics*.
- Deming W.E. and Stephan F.F.. (1940) On a last squares adjustment of sampled frequency table when the expected marginal totals are known. *Ann. Math. Statist.*, 11:427–444.
- Du C., Kurowicka D., Cooke R.M., (2006) Techniques for generic probabilistic inversion, *Comp. Stat. & Data Analysis* (50), 1164-1187.
- Fienberg S.E. (1970). An iterative procedure for estimation in contingency tables. *Ann. of Math. Statist.*, 41:907–917.

13. Fienberg S.E. (1970). An iterative procedure for estimation in contingency tables. *Ann. of Math. Statist.*, 41:907–917.
14. Haberman S.J. (1974) *An Analysis of Frequency Data*. Univ. Chicago Press.
15. Haberman S.J. (1984) Adjustment by minimum discriminant information. *Ann. of Statist.*, 12:971-988.
16. Ireland, C.T., Kullback, S., (1968). Contingency tables with given marginals. *Biometrika* 55, 179–188.
17. Kraan, B., Bedford, (2004), Probabilistic inversion of expert judgements in the quantification of model uncertainty *Management Sci.*,
18. Kraan, B.C.P., Cooke, R.M., (2000), Processing expert judgements in accident consequence modeling. *Radiat. Prot. Dosim.* 90 (3), 311–315.
19. Kraan, B.C.P., Cooke, R.M., (2000). Uncertainty in compartmental models for hazardous materials—a case study. *J. Hazard. Mater.* 71, 253–268.
20. Kruithof, J., (1937), Telefoonverkeersrekening. *De Ingenieur* 52 (8), E15–E25.
21. Kullback, S., (1968), Probability densities with given marginals. *Ann. Math. Statist.* 39 (4), 1236–1243.
22. Kullback, S., (1971), Marginal homogeneity of multidimensional contingency tables. *Ann. Math. Statist.* 42 (2), 594–606..
23. Kurowicka, D. and Cooke, R.M. (2006) *Uncertainty Analysis with High Dimensional Dependence Modeling*, Wiley.
24. Matus, F. (2007) On iterated averages of I-projections, *Statistik und Informatik*, Universit’at Bielefeld, Bielefeld, Germany, matus@utia.cas.cz.
25. *Radiation Protection and Dosimetry special issue*, vol. 90 no 3, 2000.
26. Vomlel, J. (1999) *Methods of Probabilistic Knowledge Integration* Phd thesis Czech Technical University, Faculty of Electrical Engineering