LA-UR-01-2443
submitted to Health Physics (revised)

Using Exact Poisson Likelihood Functions in Bayesian Interpretation of Counting Measurements

by

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

December 6, 2001

SUMMARY

A technique for computing the exact marginalized (integrated) Poisson likelihood function for counting measurement processes involving a background subtraction is described. An empirical Bayesian method for determining the prior probability distribution of background count rates from population data is recommended and would seem to have important practical advantages. The exact marginalized Poisson likelihood function may be used instead of the commonly used Gaussian approximation. Differences occur in some cases of small numbers of measured counts, which are discussed. Optional use of exact likelihood functions in our Bayesian internal dosimetry codes has been implemented using an interpolation-table approach, which means that there is no computation time penalty except for the initial setup of the interpolation tables.

[0] Key words: Bayesian analysis, likelihood function, Poisson distribution, internal dosimetry algorithms

1 Introduction

Many health physics measurements involve counting, usually in conjunction with a background count, which is subtracted in some way. The likelihood function in these cases then involves Poisson distributions (for long decay times). In most analyses this fact is not used in a detailed way. For example the formulas given in the ANSI standard HPS13.30 (Health Physics Society 1996) are essentially those obtained using the Gaussian approximation to the Poisson distribution (Brodsky(Brodsky 1992) argues that this is valid). Recently, the Gaussian approximation has been challenged in low-level counting situations, and various non-Gaussian formulas for the decision level have been proposed.(Strom & MacLellan 2001) Other treatments of the problem of low-level paired counting have been given by Little(Little 1982) and Potter(Potter 1999).

The calibration or normalization factor (having dimensions of physical units per count) has important uncertainties in addition to the counting statistics uncertainties. In some cases these uncertainties are known to approximately follow a log-normal distribution (see, for example, (Moss et al. 1969)), and this would generally seem a reasonable assumption. In this paper we assume the calibration factor has an arbitrarily large uncertainty that follows a log-normal distribution.

Optional use of exact likelihood functions in our Bayesian internal dosimetry codes has been implemented using an interpolation-table approach. This means that the exact likelihood functions can be used with no computation time penalty except for the initial setup of the interpolation tables.

In cases with only a few measurements involving low-level counts, we find that sometimes significant errors are made using the Gaussian approximation rather than the exact likelihood function. Such cases will be discussed later on in this paper.

This paper extends the work of Little(Little 1982) in several important ways. 1) We conceptually define the marginalized likelihood function (taking into account variability of the calibration factor), which encapsulates the information content of the measurement. 2) We are able to study the differences between the exact calculation and the Gaussian approximation of the likelihood function for internal dosimetry calculations using the present understanding of realistic, empirically determined prior probability distributions. 3) We describe the interpolation-table computational technique, which makes the exact likelihood calculations almost equivalent to the Gaussian approximation in terms of computation time. 4) We describe an empirical Bayesian method for determining the prior probability distribution of background count rates that has important practical advantages.

2 Review of the likelihood function

A measurement process producing data y (for example, detected counts) is understood in terms of parameters Θ (for example, sample activity). The prob-

ability distribution of y given particular values of Θ is assumed to be known (for example, a Poisson distribution) and is denoted by $P(y|\Theta)$ (the vertical bar is read as "given"). In statistics we are interested in inferring the value of Θ from measurements y. If the value of y is given (the measurement value), while Θ is considered a variable, $P(y|\Theta)$ is known as the likelihood function, which is a fundamental statistical function. In classical statistics, the maximum of the likelihood function is often taken as the estimate of Θ. In Bayesian statistics, the likelihood function times the prior probability distribution of Θ is proportional to the probability distribution of Θ given the measurement y. The first step in interpreting a measurement is, in any case, the evaluation of the likelihood function.

3 Exact likelihood function for counting measurements

We imagine a counting measurement, where N counts are registered in time t. With a blank sample, N_b background counts are obtained in background counting period t_b. The true count rate from activity in the sample is λ, while λ_b is the true count rate from background. The quantity of interest (for example, true activity in the sample) is denoted by ψ, and f is the multiplicative factor (a calibration or normalization factor having the dimensions of physical units per count) that relates counts to the quantity of interest,

$$\psi = f\lambda t. \tag{1}$$

The multiplicative factor is assumed to have a known uncertainty distribution given by $P(f)$. Given ψ, λ_b, f, and the counting times, the probability of jointly obtaining N counts and N_b background counts is assumed to be described by independent Poisson distributions,

$$P(N, N_b|\psi, \lambda_b, f) = \frac{(\psi/f + \lambda_b t)^N}{N!} e^{-(\psi/f + \lambda_b t)} \frac{(\lambda_b t_b)^{N_b}}{N_b!} e^{-\lambda_b t_b}. \tag{2}$$

By Bayes theorem,

$$P(\psi|N, N_b) \propto \int P(N, N_b|\psi, \lambda_b, f) P(\lambda_b) P(f) P(\psi)\, d\lambda_b\, df, \tag{3}$$

where $P(\psi)$ and $P(\lambda_b)$ are the (assumed independent) prior probability distributions of ψ and λ_b. Little(Little 1982) gives a similar formula, except that he does not include variability of the multiplicative factor. The marginalized (integrated) "log-likelihood" function $\mathcal{L}(\psi)$ is defined as

$$\mathcal{L}(\psi) = \log\left(\int P(N, N_b|\psi, \lambda_b, f) P(\lambda_b) P(f)\, d\lambda_b\, df\right) \tag{4}$$

and is a fundamental quantity in the interpretation of the measurement.

This problem is somewhat confusing in that two prior probability distributions are potentially involved, one on the true background count rate λ_b, which we consider, and another on the true quantity of interest ψ, which we do not consider, instead separating out the determination of the marginalized likelihood function from the full inference problem.

Note that, strictly speaking, Eq. 4 is not a likelihood function because of the integrations involved. Equation 4 is a natural generalization of the likelihood function in the case of no background and no uncertainty of the calibration factor. Within the Bayesian context of this paper the marginalized likelihood function encapsulates the entire information content of the measurement of ψ for subsequent use, whether Bayesian or Classical, just as a likelihood function does. The problem of inferring ψ might be done Classically (for example using the maximum likelihood method) without requiring a prior on ψ. Naturally, we would argue that a better approach would be to use the prior probability distribution of ψ to calculate the posterior probability distribution of ψ.

The prior probability distribution of background counting rate λ_b is assumed to be given by the (conjugate) Gamma distribution,(Little 1982, Martz 2000)

$$P(\lambda_b) = \frac{\beta^\alpha}{\Gamma(\alpha)} \lambda_b^{(\alpha-1)} e^{-\beta\lambda_b}, \tag{5}$$

with $\Gamma(\alpha)$ the gamma function

$$\Gamma(\alpha) = \int_0^\infty x^{\alpha-1} e^{-x}\, dx. \tag{6}$$

The gamma distribution allows a large range of possibilities. For example, if $\alpha = 1/2$ and in the limit of $\beta \to 0$ we obtain an improper prior that is proportional to $\lambda^{-1/2}$ (known as Jeffrey's noninformative prior for the Poisson distribution).(Robert 1997, Martz 2000) The term "improper" refers to the fact that the distribution has a divergent normalization integral, which is often irrelevant in Bayesian analysis. If $\alpha = 1$ and again taking the limit $\beta \to 0$ produces another improper prior, a constant, called the "flat prior", which is commonly used. In this paper we recommend that the prior parameters α and β be determined empirically from data, which results in values of $\alpha > 1$ and $\beta > 0$. The peak (mode) of the gamma distribution occurs at $(\alpha - 1)/\beta$. For β large the Gamma distribution is narrow, and conversely.

The probability distribution of multiplicative factor f is assumed to be log normal,

$$P(f) = \frac{1}{\sqrt{2\pi}\sigma_f f} \exp(-\frac{1}{2\sigma_f^2}(\ln \frac{f}{f_0})^2), \tag{7}$$

with median value f_0 and (geometric) standard deviation σ_f. Making the change of variables

$$f = f_0 e^\phi, \tag{8}$$

then ϕ has a normal distribution centered at 0.

The log-likelihood function consists of the following double integral (ignoring an additive constant),

$$\begin{aligned}\mathcal{L}(\psi) &= \log\left(\int_0^\infty d\lambda_b \int_{-\infty}^{\infty} d\phi\, e^{-\frac{\phi^2}{2\sigma_f^2}}\right. \\ &\left.(\frac{\psi}{f_0}e^{-\phi}+\lambda_b t)^N e^{-(\frac{\psi}{f_0}e^{-\phi}+\lambda_b t)}[\lambda_b(t_b+\beta)]^{N_b+\alpha-1}e^{-\lambda_b(t_b+\beta)}\right). \quad (9)\end{aligned}$$

Our numerical method for evaluating Eq. 9 is straightforward summation using Monte Carlo generation of λ_b from a Gamma distribution and ϕ from a normal distribution.(Press et al. 1986)

4 Gaussian Approximation

A Gaussian probability density function with mean x_0 and standard deviation σ is of the form

$$P(x) = \frac{\exp\left(\frac{-(x-x_0)^2}{2\sigma^2}\right)}{\sqrt{2\pi}\sigma} \quad (10)$$

The Gaussian approximation to the exact likelihood function results from the fact that

$$x^n e^{-x} \approx \text{const} \times e^{-\frac{1}{2n}(x-n)^2} \quad (11)$$

for n large. Using this relationship and the fact that the convolution of two Gaussians is again a Gaussian one finds that

$$\mathcal{L}(\psi) \approx -\frac{1}{2}\left[\frac{(y-\psi)^2}{\sigma_y^2(\psi)} + \log(\sigma_y^2(\psi))\right], \quad (12)$$

for N and $N_b+\alpha-1$ sufficiently large, where

$$\begin{aligned} y &= f\left[N - \frac{N_b+\alpha-1}{R}\right] \\ \sigma_y^2(\psi) &= f^2\left[N + \frac{N_b+\alpha-1}{R^2}\right] + (\sigma_f\psi)^2, \end{aligned} \quad (13)$$

using the notation

$$R = \frac{t_b+\beta}{t}. \quad (14)$$

In Eqs. 13, if $\alpha = 1$ and $\beta = 0$, one recognizes the net counting rate $N/t - N_b/t_b$ in y and the familiar components of the total uncertainty variance associated with gross counts, background counts and multiplicative uncertainty in $\sigma_y^2(\psi)$. By comparing with the exact log-likelihood function given by Eq. 9, we can establish how large N and $N_b+\alpha-1$ need be for the Gaussian approximation to be valid.

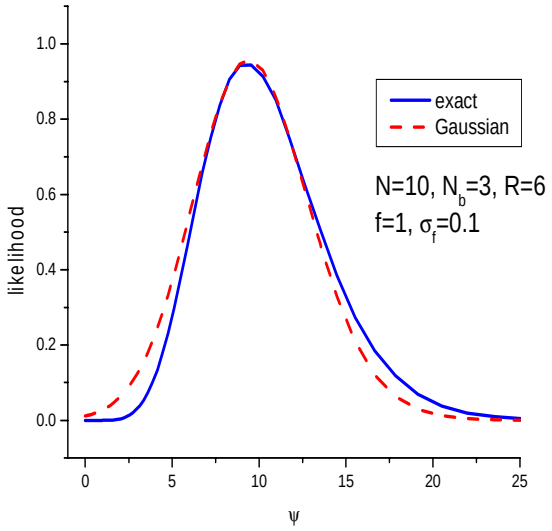

Figure 1: Comparison of exact and Gaussian approximation for the likelihood function, for a case with 10 counts and 3 background counts.

For $\alpha = 1$ and $\beta = 0$, the Gaussian approximation has an obvious difficulty when $N = N_b = 0$, giving $\sigma_y^2(0) = 0$ (implying a perfect measurement). To avoid this problem, we redefine the Gaussian by the replacement

$$N + \frac{N_b + \alpha - 1}{R^2} \to 1, \tag{15}$$

in Eq. 13 for σ_y^2 whenever the left hand side of Eq. 15 is less than 1. Rather than such an ad hoc procedure, we prefer the Empirical Bayes determination of α and β described in the next section, which results in values of $\beta > 0$ and $\alpha > 1$, avoiding this problem in a natural way.

It turns out the Gaussian approximation is surprisingly accurate, even for a fairly small number of counts. For example, Fig. 1 shows a comparison between Eq. 12 and Eq. 9 for 10 counts and 3 background counts with the ratio of background count time to count time $R = t_b/t = 6$, calibration factor $f = 1$ with small uncertainty $\sigma_f = 0.1$. A flat prior probability distribution of background count rate is assumed ($\alpha = 1$, $\beta = 0$).

As a practical example of the use of the exact likelihood function, consider a hypothetical radiochemical alpha spectrometric (RAS) measurement of plutonium in urine, assuming $f = 0.15\,\mathrm{Bq/d}$ per count (corresponding to a typical RAS calibration factor), and $R = 6$. The likelihood functions are as shown in Fig. 1. The measurement is assumed to be one year after the preceding bioassay sample. Figure 2 shows the posterior distribution of committed effective dose

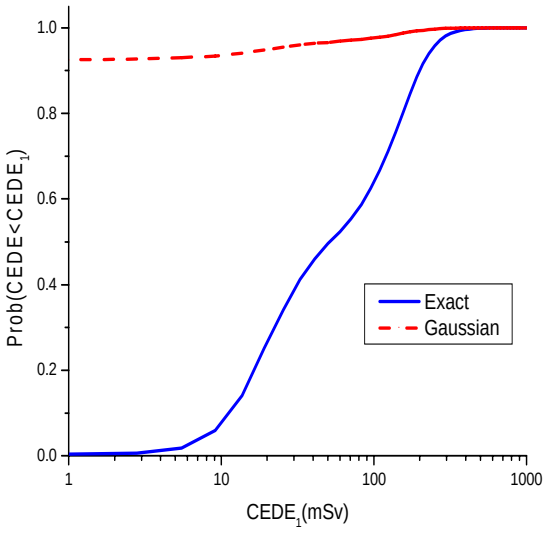

Figure 2: Comparison of exact and Gaussian approximation for the cumulative posterior probability distribution of CEDE. A single ^{239}Pu RAS urine bioassay measurement giving 10 count and 3 background counts taken one year after the preceding sample is assumed.

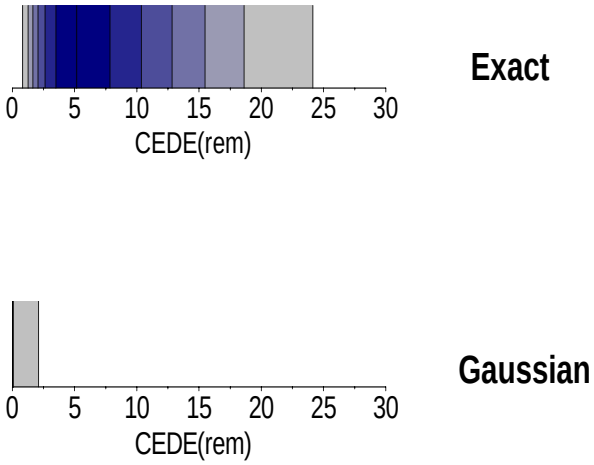

Figure 3: Shaded contour plot representation of the comparison of exact and Gaussian approximation for the cumulative posterior probability distribution of CEDE.

equivalent (CEDE) calculated using the method described in ref. (Miller, Martz, Little & Guilmette 2001), assuming an "alpha" prior on intake amount, with $\alpha = 0.001\,\mathrm{y}^{-1}$(Miller, Inkret, Little, Martz & Schillaci 2001). The ICRP-30 family of biokinetic models are used. The calculated CEDE is quite different in this case. The reason for this disagreement is that the exact likelihood function is zero at $\psi = 0$ while the Gaussian approximation is not. If the likelihood function is nonzero at zero, the alpha prior weights the posterior strongly toward zero dose.

Figure 3 shows a contour plot representation of the cumulative posterior probability distributions shown in Fig. 2. The darker areas of the plot are near the median probability, and the shading extends out to the 5% and 95% credible limits. This type of graphical representation was developed to more effectively communicate the idea of dose uncertainty to the workers.

We define the quantity $r = \mathrm{CEDE}_{\mathsf{Gaussian}}/\mathrm{CEDE}_{\mathsf{exact}}$, which is the ratio of the expectation value of CEDEs (note that expectation value is the same as mean or average value) calculated in the two ways assuming a single bioassay

measurement. Instead of CEDE we could equally well assume any calculated quantity of interest. Table 1 shows values of r for small numbers of counts and different values of $R = t_b/t$ (small values of R typically pertain to in vivo measurements with a subtracted body background). Small calibration factor uncertainty $\sigma_f = 0.1$ is assumed. The "alpha" prior on intake amount is assumed, with $\alpha\Delta t = 0.1$(Miller, Inkret, Little, Martz & Schillaci 2001). The counts N and N_b are 0, 1, 2, 3 in all possible combinations. The average r is shown together with the minimum and maximum values. Table 1 shows some cases of substantial disagreement for R small.

Table 1: Ratio of quantity of interest calculated using Gaussian likelihood to that using exact likelihood, r, as a function of the ratio R of background to sample count time.

R	r^*_{average}	min $(N, N_b)^\dagger$	max $(N, N_b)^\dagger$
6	0.79	0.48 (3, 0)	1.38 (0, 0)
1	0.77	0.46 (3, 0)	1.06 (0, 3)
1/6	3.54	1.37 (0, 0)	6.29 (0, 3)

* average for N, N_b in {0,1,2,3}

† minimum or maximum attained at (N, N_b)

5 Role of the prior probability distribution

The prior probability distribution $P(\lambda_b)$ summarizes additional information we may have about the background count rate λ_b before making the background measurement. Using the gamma function with parameters α and β, the prior distribution has expectation value (mean) α/β and variance α/β^2.(Rothschild & Logothetis 1986) The peak (mode) of the gamma distribution occurs at $(\alpha - 1)/\beta$. For β small the Gamma distribution is broad, which corresponds to assuming that we have little knowledge of what the value of λ_b might be (other than that it is positive). A commonly used limit of parameter values is $\alpha \to 1$, $\beta \to 0$, which results in a very broad (flat) prior. The interpretation of this choice will be discussed later on.

It is also possible to use population data to determine α and β. This is preferable because more relevant information is used in interpreting the measurement. For a representative set of background count data, the observed distribution of measured counts is given by the following convolution integral:

$$P(N_b) = \int P(N_b|\lambda_b)P(\lambda_b)\, d\lambda_b, \tag{16}$$

where $P(\lambda_b)$ is the prior probability distribution (the distribution of λ_b in the measured population), and $P(N_b|\lambda_b)$ is the Poisson likelihood of measuring N_b

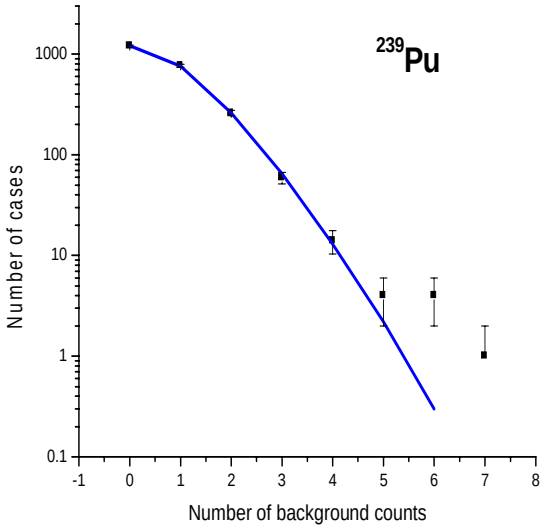

Figure 4: Distribution of ^{239}Pu background counts together with fit assuming a gamma-function prior.

background counts when the true background count rate is λ_b. Assuming a gamma prior distribution,

$$P(N_b) = \frac{\Gamma(\alpha + N_b)\beta^\alpha t_b^{N_b}}{N_b!\Gamma(\alpha)(t_b + \beta)^{\alpha + N_b}}. \tag{17}$$

That this convolution integral exists in closed form (provided $\beta > 0$) shows the usefulness of assuming a gamma prior distribution with a Poisson likelihood function.(Martz 2000)

Figures 4, 5, and 6 show population data for the frequency distribution of background counts for α-particle energy spectrometry in the ^{239}Pu, ^{238}Pu, and ^{242}Pu energy regions.

Nonlinear fits to these distributions were obtained by minimizing χ^2 given by

$$\chi^2 = \sum_{j=1}^{M} \frac{(N_j - NP(N_b(j)))^2}{\sigma_j^2}, \tag{18}$$

where N_j is the number of cases that have background counts $N_b(j)$, N is the total number of cases, and $P(N_b)$ is given by Eq. 17. The uncertainty σ_j^2 was assumed to be given by the Poisson variance, except that it is not allowed to be zero when the background counts are zero.

$$\sigma_j^2 = \max(N_j, 1). \tag{19}$$

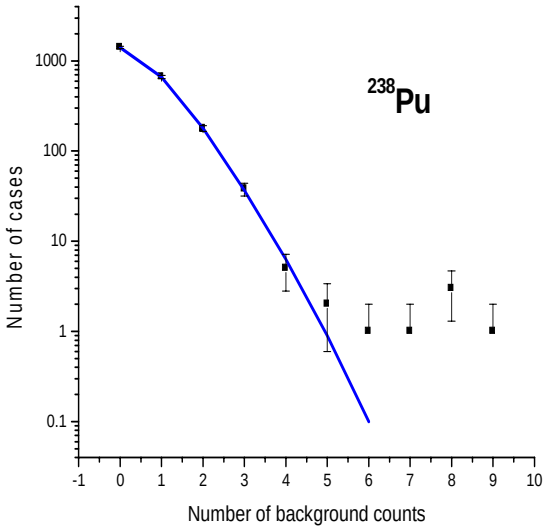

Figure 5: Distribution of ^{238}Pu background counts together with fit.

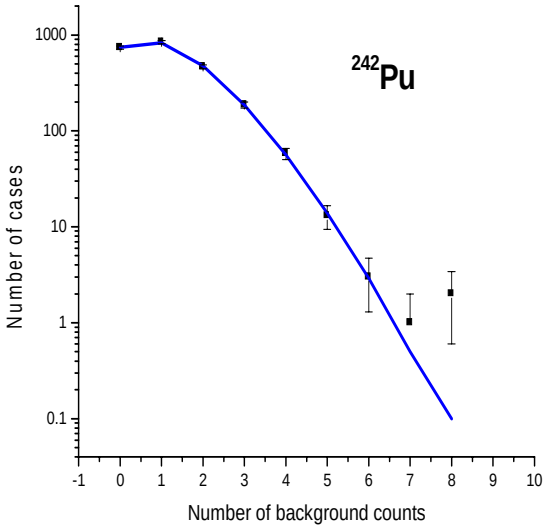

Figure 6: Distribution of ^{242}Pu background counts together with fit.

The fit results are summarized in Table 2. The quantity χ^2/N_{DF} is meant as a

Table 2: Results of fits to background count data.

Nuclide	Number cases*	$\frac{(\alpha-1)}{\beta}t_b$ †	β/t_b	$\chi^2/N_{DF}^{\ddagger}$
^{239}Pu	2307	3.64	2.71	0.98
^{238}Pu	2307	2.58	2.20	0.89
^{242}Pu	2307	6.72	5.35	0.57

* counting period t

† most probable number of background counts in count period $t_b = 6t$

‡ this quantity should be less than about 1 for a satisfactory fit–see text

measure of goodness of fit, where N_{DF} ("number of degrees of freedom") is the number of data minus the number of fit parameters. This quantity should be about 1 or less for a satisfactory fit, as is true for all these cases.

The quantities α and β have a simple-to-understand qualitative interpretation. If, instead of using an informative prior defined by given, postulated, values of α and β, the measured background counts had been $N_b + \alpha - 1$ for a counting period $t_b + \beta$, the result, using a flat prior, would be the same. Thus $\alpha - 1$ behaves like background counts, and β behaves like background count time. The flat prior ($\alpha = 1$, $\beta = 0$) is then like the limit of measuring 0 background counts in 0 time. The quantity β being large is equivalent to counting the background for a long time, and the background count rate is then well known.

If β is large compared to t_b, the prior overwhelms the background count data. Note that if the background count rate is truly variable in the population, the value of β will be small, and the background measurement will tend to be more important than the background prior.

A software package (BKGRND)(Miller & Little 2001) is available from the authors to enable the reader to carry out these and similar calculations using a supplied database of actual plutonium RAS background count data. By varying the size of the dataset and combining background count data over multiple counting periods one finds that $(\alpha - 1)/\beta$, the peak (or mode) of the distribution, is relatively well determined but differing values of β are obtained. Thus the Empirical Bayes method is not entirely straightforward. We recommend conservatism in not allowing the fitted value of β/t_b to be larger than some limit (for example $\beta/t_b = 2$). This means that we do not allow very narrow priors, such that the background count data is completely unimportant. The ratio of $(\alpha - 1)/\beta$ is kept constant if β is decreased in this way.

6 Discussion

A method for using the exact marginalized (integrated) likelihood function for Poisson measurement processes involving a background subtraction has been

described. The exact marginalized likelihood function may be thought of as summarizing (within a Bayesian context) the full information content of the measurement. This exact likelihood function may be substituted for the commonly used Gaussian approximation in either a Bayesian or hybrid Classical subsequent analysis in order to make inferences from the data.

The Gaussian approximation is surprisingly accurate, but differences occur in some cases of small numbers of measured counts. We have discussed such cases involving internal dosimetry in Sec. 4 where there are significant difference in the committed effective dose equivalents inferred from single bioassay measurements using the two methods.

The mathematics seems rather complex, and one might ask how such difficult-to-evaluate expressions could be useful in practice. It is important to realize that the likelihood function for each measurement can be summarized once and for all as a numerical table with a fairly small investment of computation time. Further analysis can then be done by interpolating values from these tables, which can be done about as fast as the Gaussian approximation can be evaluated. We have taken the interpolation-table approach in incorporating use of the exact likelihood function as options in our Bayesian internal dosimetry codes UF(Miller et al. 1999) and ID(Miller, Martz, Little & Guilmette 2001). Computation time is a limiting factor with the ID code for a single case and with the UF code for thousands of cases, and much of that time is taken up evaluating marginalized log likelihood functions. Having to repeatedly evaluate the multiple integrals of Eq. 9 without using interpolation tables would be out of the question (tens of thousands of times slower). The setup time for the interpolation tables is typically on the order of a minute (for a case involving tens of data points and a 1 GHz Pentium processor). That time is long enough that the Gaussian approximation is still handy to use. It has been proposed that analysis laboratories report exact marginalized likelihood functions in tabular form as part of their "product delivery",(Miller & Little 2001) in which case we could dispense with the Gaussian approximation entirely.

We have chosen not to attempt to refine the definition of decision levels and MDA(Health Physics Society 1996) based on this work. The use of a decision level itself seems problematical in that it may in practice lead to loss of measurement information. For example, a number of sequential measurement results might be reported as "below decision level" with no action or further reporting of data, while the posterior probability based on all the measurements (combined by adding their log-likelihood functions) might indicate a need for action. It actually seems conceptually and operationally simpler to us (in this era of desk-top computers) to just calculate the posterior distribution of the quantities of interest using all the data (by an appropriate mouse click) and to make decisions (for example to declare "positive" or to resample) based on examining the posterior probability distribution.

The Empirical Bayes methodology for count background subtraction (determination of true background count rate) proposed here has important practical advantages: 1) more information is utilized (the historical distribution of background measurements) producing a higher quality measurement result, 2) ad-

hoc methods for dealing with the zero count problems are avoided, and 3) the quantities α and β obtained in the empirical Bayes analysis of the background count distributions are useful in quality control (small β indicating variable background count rates and possible problems).

Finally, it is very important that count measurement data be recorded in such a way that the full information content is not diminished. This is an obvious and simple requirement, often not satisfied. One way of doing this is to record counts N, background counts N_b, the ratio of background count time to count time R, calibration factor f, and σ_f, treating y and σ_y of Eq. 13 as secondary quantities.

7 Acknowledgments

We thank Dan Strom for a discussion of references to previous work in this area. We are grateful to Joe Graf for suggesting the shaded contour plot representation of the posterior distribution.

References

Brodsky, A. (1992), 'Exact Calculation of Probabilities of False Positives and False Negatives for Low Background Counting', Health Physics 63, 198–204.

Health Physics Society (1996), An American National Standard–Performance Criteria for Radiobioassay, Technical Report HPS N13.30-1996, Health Physics Society, 1313 Dolley Madison Blvd, Suite 402, McLean, VA, 22101.

Little, R. J. A. (1982), 'The Statistical Analysis of Low-Level Radioactivity in the Presence of Background Counts', Health Physics 43, 693.

Martz, H. F. (2000), An Introduction to Bayes, Hierarchical Bayes, and Empirical Bayes Statistical Methods in Health Physics, in T. B. Borak, ed., 'Applications of Probability and Statistics in Health Physics', Medical Physics Publishing, Madison, Wisconsin, chapter 2, pp. 55–84.

Miller, G., Inkret, W. C., Little, T. T., Martz, H. F. & Schillaci, M. E. (2001), 'Bayesian Prior Probability Distributions for Internal Dosimetry', Radiation Protection Dosimetry 94, 347–352.

Miller, G., Inkret, W. C. & Martz, H. F. (1999), 'Internal Dosimetry Intake Estimation Using Bayesian Methods', Radiation Protection Dosimetry 82, 5–17.

Miller, G. & Little, T. (2001), New Ways to Report Counting Measurements, in '47*th* Annual Radiochemical Measurements Conference, Honolulu, 4-8 November, 2001'.

Miller, G., Martz, H. F., Little, T. T. & Guilmette, R. (2001), 'Bayesian Internal Dosimetry Calculations Using Markov Chain Monte Carlo', Radiation Protection Dosimetry . (to appear, LA-UR-01-1573).

Moss, W. D., Campbell, E. E., Schulte, H. F. & Tietjen, G. L. (1969), 'A Study of the Variations Found in Plutonium Urinary Data', Health Physics 17, 571.

Potter, W. E. (1999), 'Neyman-Pearson Confidence Intervals for Extreme Low-Level, Paired Counting', Health Physics 76, 186–187.

Press, W. H., Flannery, B. P., Teukolsky, S. A. & Vetterling, W. T. (1986), Numerical Recipes–The Art of Scientific Computing, Cambridge University Press.

Robert, C. P. (1997), The Bayesian Choice, Springer.

Rothschild, V. & Logothetis, N. (1986), Probability Distributions, John Wiley and Sons, New York.

Strom, D. J. & MacLellan, J. A. (2001), 'Evaluation of Eight Decision Rules for Low-Level Radioactivity Counting', Health Physics 81, 27–34.