The Astrophysical Journal, 374:344–355, 1991 June 10

DETERMINATION OF CONFIDENCE LIMITS FOR EXPERIMENTS WITH LOW NUMBERS OF COUNTS

Ralph P. Kraft, David N. Burrows, and John A. Nousek
Department of Astronomy and Astrophysics, The Pennsylvania State University, 525 Davey Laboratory, University Park, PA 16802
Received 1990 June 27; accepted 1990 December 5

ABSTRACT

We compare two different methods, classical and Bayesian, for determining confidence intervals involving Poisson-distributed data. We are specifically concerned with cases where the number of counts observed is small and is comparable to the mean number of background counts. We give our reasons for preferring the Bayesian over the classical method. Tables of confidence limits calculated by the Bayesian method are provided for quick reference.

Subject headings: gamma rays: general — numerical methods — X-rays: general

1. INTRODUCTION

We compare two different methods for calculating confidence intervals for counting experiments in the presence of a nonzero background, when the total number of counts detected is low enough so that the appropriate probability distribution is Poissonian rather than Gaussian. The problem of obtaining correct upper limits and confidence regions for experiments with no background has been discussed recently by Gehrels (1986). Helene (1983, 1984) discusses the case of nonzero background.

These authors use fundamentally different approaches. The method described by the Gehrels is called the classical method, and the confidence intervals calculated using this technique are often called Neyman-Pearson confidence intervals. Helene used a different approach based on Bayes's theorem. Bayes's theorem is used to compute the probability distribution function describing the relative probability that each value of the known parameter could lead to the observed experimental outcome. This probability distribution is integrated to determine the confidence intervals. The two methods are both statistically valid, and they give similar or identical results in most instances. However, we find significant differences in some cases.

The differences between the results obtained with these two methods arise from the underlying philosophies of these statistical approaches. In classical statistics, one assumes that there is a "true," but unknown, value which we attempt to measure (in our case, a source flux), and that our measurements represents a small sample of possible measurements taken from a "parent" probability distribution whose mean value is the true source flux we are attempting to determine (Bevington 1969). If we could make an infinite number of measurements, we could determine the parent distribution exactly and therefore obtain the desired value of our measured quantity. Given a finite set of measurements, we only obtain a sample of the parent population. This allows us to estimate the true flux with some uncertainty which can be expressed in terms of a confidence interval. Another observer making the same set of observations would obtain a different random sample of the parent population and would therefore make a different estimate of the true flux and confidence interval. The classical definition of confidence levels is the fraction of observers who obtain confidence intervals containing the true flux at the given confidence level.

The Bayesian approach, on the other hand, assumes that the observer can constrain the experimental hypotheses by using both the measured data and a priori knowledge about the physical system being measured. Rather than asking what fraction of *observers* detecting different numbers of counts from a given source would obtain confidence intervals that include the true rate, the Bayesian approach inverts this question to determine the probabilities that *sources* of different flux could produce the observed rate. Thus, the Bayesian confidence interval is making a statement about the source population rather than one about the population of observations.

It is not our intent to enter the philosophical debate between advocates of classical statistics and advocates of Bayesian statistics. We will, however, argue that the Bayesian definition of confidence intervals reflects common astronomical usage better than the classical definition does, and that the Bayesian method provides a more intuitively satisfying result for the case of interest to our discussion than does the classical method. Our motivation is twofold: to bring this problem and its correct resolution to the attention of the astronomical community and to provide more complete tables of Bayesian confidence limits than those given by Helene. In § 2 of this paper, we discuss both methods as they apply to the problem of interest (low numbers of counts in the presence of nonzero background) and give our reasons for preferring the Bayesian method used by Helene over the classical method used by Gehrels. Our presentation is not meant to be mathematically rigorous, but it will give the reader the essence of each technique. We expand the results of Helene by calculating confidence limits over a more complete set of parameters. In § 3 we apply the Bayesian method to the problem of determining upper limits on the soft X-ray flux from SN 1987A. Section 4 contains a brief summary and conclusions.

2. CALCULATION OF UPPER AND LOWER LIMITS

2.1. *The Classical Method*

If the flux of photons from a source is Poisson-distributed, the probability of detecting N photons in the observation time t is

$$P(N) = \frac{e^{-S} S^N}{N!}\,, \tag{1}$$

where S is the mean number of counts from the source in time

t. Equation (1) is often called the distribution function or the conditional probability distribution function for the random variable N.

In any real observation there is always some background contribution. Assuming the background is also Poisson-distributed, the probability of detecting N counts in this case is

$$P(N) = \sum_{\substack{N_s, N_b \\ N_s + N_b = N}} \frac{e^{-S} S^{N_s}}{N_S!} \frac{e^{-B} B^{N_b}}{N_b!} = \frac{e^{-(S+B)}(S+B)^N}{N!}, \quad (2)$$

where B is the mean number of background counts observed in time t. N_s and N_b are the detected number of source counts and background counts, respectively. Although we cannot determine N_s and N_b from the observation, their sum must be the observed number of counts, N.

At this point, we assume that the mean number of background counts, B, is known to a high degree of precision (e.g., B could be well sampled in an off source spatial region or off-line spectral region), and we neglect any uncertainty in the measurement of B.[1] However, the particular number of background counts, N_b, in a given observation cannot be determined, because they are the result of Poisson fluctuations. Given the observationally determined values of N and B, what can we then say about S?

The technique used to constrain a model parameter for a given set of data is called the method of confidence intervals. In classical statistics a confidence interval is defined in the following manner. Two statistics, called $S_{\max}$ and $S_{\min}$, form a confidence interval at confidence level CL ($0 \le \text{CL} \le 1.0$) if the probability that $S_{\min} \le S \le S_{\max}$ is greater than or equal to CL, independent of S (Larson 1974). These two statistics, $S_{\max}$ and $S_{\min}$, will depend on N and B. The confidence interval is defined in such a manner that if many observations are made and confidence intervals are assigned to each measurement, $100 \times \text{CL}\%$ of the observations would generate lower limits less than the true value of the unknown parameter and upper limits greater than the true value of the parameter. Any given set of confidence limits may or may not actually contain the true value of the unknown parameter, since there is always a probability $1 - \text{CL}$ that the true value lies outside the interval.

Applying the formal classical definition of confidence limits to the problem of determining limits for counting experiments with negligible background ($B = 0$), the single-sided lower limit on S is given by

$$\sum_{n=0}^{N-1} \frac{e^{-S_{\min}} S_{\min}^n}{n!} = \text{CL}, \quad (3)$$

while the equivalent expression for an upper limit is

$$\sum_{n=0}^{N} \frac{e^{-S_{\max}} S_{\max}^n}{n!} = 1 - \text{CL} \quad (4)$$

(Cramér 1945; Gehrels 1986 and references therein). Here $S_{\max}$ and $S_{\min}$ are the single-sided upper and lower confidence limits, respectively, and CL is the confidence level of the limits. Equations (3) and (4) cannot, in general, be solved analytically for $S_{\max}$ and $S_{\min}$, given CL and N. Gehrels (1986) has solved them numerically for a wide range of CL and N, and provides analytic approximations for several limiting cases.

Figures 1a and 1b graphically illustrate this calculation. The Poisson distribution is plotted for mean values of 2.202 and 13.06. These are the single-sided lower and upper limits on the mean number of source counts at confidence level 0.975 for an experiment in which 6 counts were detected. The shaded area under each curve is equal to the confidence level, and corresponds to the sums of equations (3) and (4).

Double-sided confidence limits are commonly defined (e.g., Gehrels) by replacing CL with $\frac{1}{2}(1 + \text{CL})$ in equations (3) and (4) and solving for $S_{\max}$ and $S_{\min}$ to obtain limits with confidence level CL′. This choice of confidence limits is called the central confidence interval (Barnett 1973). Graphically, the upper and lower limits are chosen to make the unshaded areas of Figures 1a and 1b equal to each other. Thus the 97.5% single-sided confidence limits given above for $N = 6$ become 95% double-sided confidence limits. Note that this is a purely arbitrary choice. For a given confidence level there is an infinite choice of possible confidence intervals, ranging from a single-sided upper limit ($S_{\min} = 0$) at one extreme, through the central interval, to a single-sided lower limit ($S_{\max} = \infty$) at the other extreme. Note also that the confidence level and upper and lower limits are not mutually independent: once any two are chosen, the third is uniquely defined.

An alternative choice of confidence limits is that which minimizes the size of the confidence interval. We refer to this as the minimal interval. This places the tightest constraints on model parameters that can be derived from a given data set, and is the preferred choice in our opinion. We have performed numerical simulations which demonstrate that the central choice of limits does not minimize the confidence interval for experiments with small numbers of counts. For example, we have calculated 95% confidence intervals for an observation detecting 3 counts with negligible background. Figure 2 shows the width of the confidence interval, $S_{\max} - S_{\min}$, as a function of CL_U/CL_L, where CL_U and CL_L are the confidence levels of the single-sided upper and lower limits, corresponding to $S_{\max}$ and $S_{\min}$, respectively. The central confidence interval corresponds to $\text{CL}_U/\text{CL}_L = 1.0$. This choice of limits does not minimize the confidence interval for small (but nonzero) N. In the Gaussian limit (i.e., as N goes to infinity) the central confidence interval becomes the confidence interval of minimum size.

Gehrels's tabulation of classical confidence limits does not consider the case of $B \neq 0$. It is desirable to extend these results to apply to situations in which the background cannot be neglected. There are several ways in which the method described above could be used naively to generate "confidence intervals" in the presence of nonzero background that do *not* satisfy the classical definition of confidence intervals. We describe several methods that are commonly used and explain why they are incorrect.

Perhaps the most common error is to derive an $n\ \sigma$ upper limit using the Poisson distribution for the mean number of background counts. For example, one might obtain a 3 σ upper limit by multiplying the standard deviation of the background by 3. While this is a correct procedure for determining the level of significance of a source *detection*, it does not correctly obtain upper or lower limits on the source *flux*. The $n\ \sigma$ uncertainty in the background only gives the probability that a statistical fluctuation in the background could have given the observed number of counts in the absence of a source.

[1] This assumption is very important below when we extend the classical method to cases of nonzero background. If the background rate is uncertain, the classical method cannot be readily extended to obtain confidence limits satisfying the classical definition.

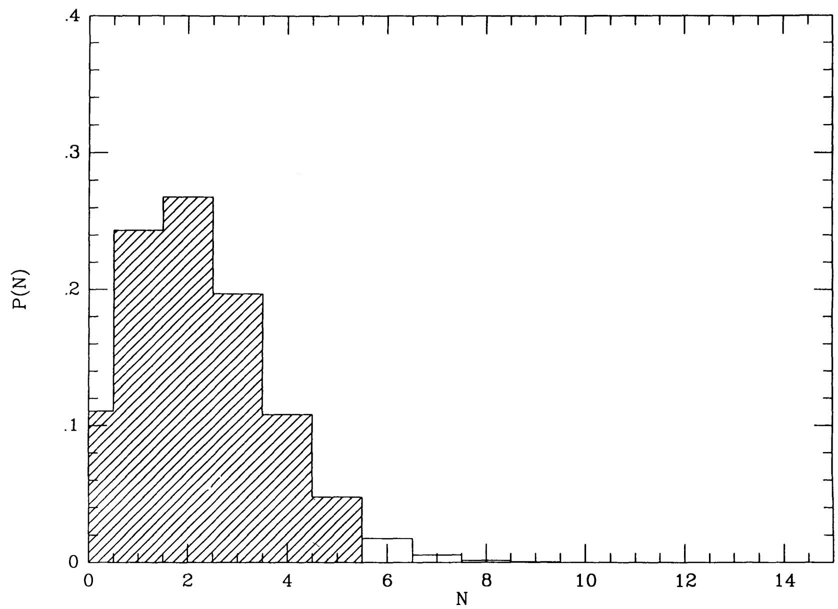

FIG. 1a

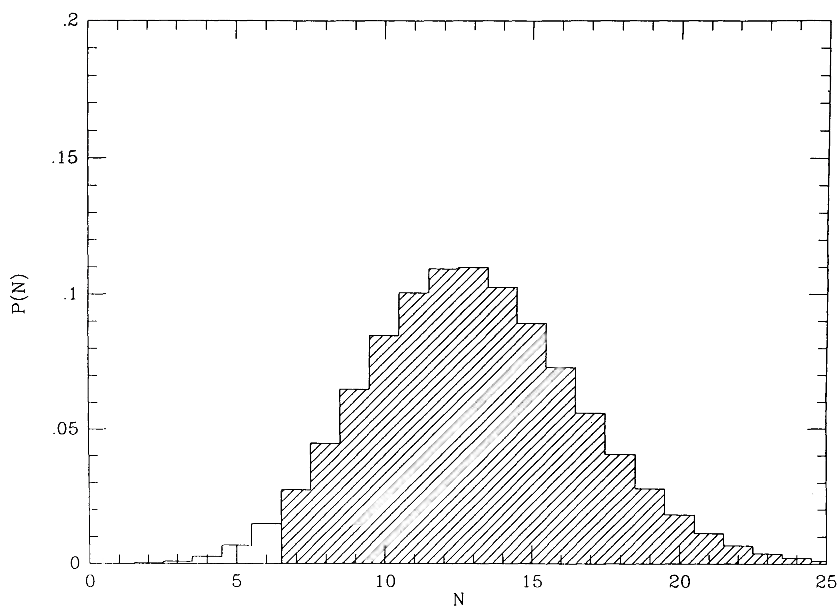

FIG. 1b

FIG. 1.—Conditional probability distribution function for Poisson-distributed data. The mean in (a) is 2.202, and the mean in (b) is 13.06. These correspond to the single-sided upper and lower limits calculated using the classical method with CL = 0.975 and $N = 6$.

Another incorrect method is to subtract the mean number of background counts from the observed number of counts and use this as the number of observed "source counts," $N' = N - B$. This method, in contrast to the previously considered case, completely ignores Poisson fluctuations in the number of background counts detected, confusing the precision with which the mean number of background counts is known with the determination of the number of background counts contributing to an individual observation. The problems with this technique are apparent in the case of $N < B$, where it requires negative source counts, but are also evident if B is not an integer, which would require the detection of fractional photons in this interpretation.

One way to extend Gehrels's method to cases of nonzero background that satisfies the classical definition of confidence intervals is to calculate confidence limits for the total number of observed counts, N, in the manner described above at a confidence level of CL, and then subtract the mean number of background counts, B, from the calculated limits. Negative limits are dealt with by setting the corresponding limits to zero.

That these upper and lower limits satisfy the classical definition of confidence limits can easily be seen in the following example. Suppose we make an observation of $N = N_S + N_B$ counts, where the mean background, B, is nonzero. From this observed number of counts, N, we can calculate upper and lower limits (called U and L to distinguish them from S_{max} and

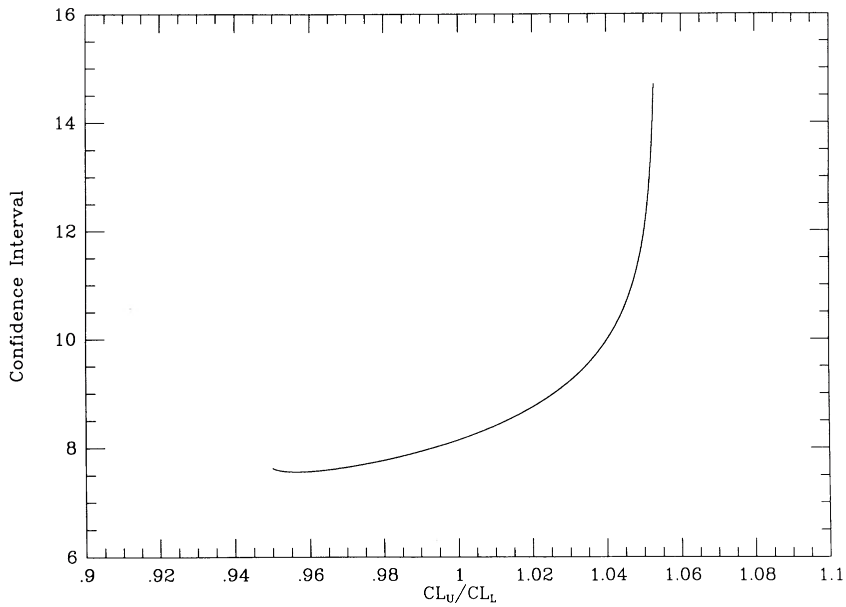

FIG. 2.—Confidence interval width $S_{max} - S_{min}$ as a function of the ratio of the single-sided confidence levels, CL_U/CL_L, for the case of $B = 0$ and $N = 3$

S_{min}) in the manner described above. These limits are on $S + B$, rather than on S. For purposes of this argument, we will consider only the upper limit, U, but the argument can be extended to the lower limit as well. We know that the limits on the distribution of $S + B$ satisfy the definition of confidence intervals given above, so that at least $100 \times$ CL% of the time $S + B$ will be less than or equal to U. Since we have assumed that B is known to relatively high precision from measurements in an off-source region, the inequality can only be satisfied by placing a limit on S. Rearranging the inequality, it must be true that $S \leq U - B$ for at least $100 \times$ CL% of the observations. Therefore, the statistic S_{max} satisfies the definition of an upper limit for S, where

$$S_{max} = \begin{cases} U - B\,, & U \geq B\,, \\ 0\,, & U < B\,. \end{cases} \tag{5}$$

From this we conclude that Gehrels's tables can be used for cases of nonzero background by subtracting the mean number of background counts from the limits in his tables. This will give confidence limits that satisfy the classical definition of confidence intervals. We have several objections to this technique however, which we discuss in § 2.3, where we compare the classical method and Bayesian method.

2.2. *The Bayesian Method*

An alternative approach to determining upper and lower limits on S is described by Helene (1983, 1984). This method makes use of Bayes's theorem, which is given by

$$f_{N,B}(S) \propto p(S)P_S(N) \tag{6}$$

(Eadie et al. 1971). Here $f_{N,B}(S)$ is called the posterior probability function for the parameter S as a function of the observables, N and B. The first term on the right-hand side of equation (6), $p(S)$, is the prior distribution function (often called "prior"). This function incorporates the observer's degree of belief or prior knowledge in the different possible values of S before an observation is made. The second term on the right-hand side of equation (6), $P_S(N)$, is the conditional distribution function, which is the Poisson distribution for S (or $S + B$) in this case. For a full derivation and explanation of Bayes's theorem see, for example, Larson (1974), Lindley (1980), or Loredo (1990).

All prior or subjective knowledge of the physical conditions applying to the experiment is taken into account in the prior distribution function. A commonly cited argument against using Bayesian statistics is that this prior distribution function must be estimated in order to apply Bayes's theorem. Indeed, the use of a prior distribution function is both a strength and a weakness of Bayesian statistics. On the one hand, inclusion of a priori information such as the nonnegativity of source flux allows one to constrain the problem based on the characteristics of the physical system being measured. Such information has proved to be of great utility in applications such as maximum-entropy image processing (Narayan & Nityananda 1986). On the other hand, the choice of a prior function introduces an element of subjectivity into the data analysis that is disturbing to many scientists.

In deriving the Bayesian expression for confidence intervals, we have chosen a prior function that minimizes the introduction of subjective information by imposing only the condition of nonnegativity. As we have no knowledge of the source flux before an observation is made, we will assume an initial distribution function which is constant (i.e., uniform) for source fluxes ranging from zero to infinity and which is zero for negative source fluxes. By making this choice we have assumed that all positive source fluxes are equally probable. This assumption is clearly unrealistic. For example, one would not expect the X-ray flux from a distant quasar to exceed the solar X-ray flux, and the prior probability distribution function should therefore become zero for large fluxes. At the end of this section we show that our results depend only weakly on this simplification.

The posterior probability distribution function for S is found by inserting equation (2) in Bayes's theorem with a constant $p(S)$, which gives

$$f_{N,B}(S) = C\,\frac{e^{-(S+B)}(S + B)^N}{N!} \tag{7}$$

for $S \geq 0$. The normalization constant, C, is given by[2]

$$C = \left[\int_0^\infty \frac{e^{-(S+B)}(S+B)^N}{N!}\, dS\right]^{-1} = \left(\sum_{n=0}^{N} \frac{e^{-B}B^n}{n!}\right)^{-1} . \quad (8)$$

The lower limit of integration in equation (8) is zero because we have chosen a prior that excludes negative values of the source flux. Helene (1983) considers the modifications required if B is uncertain.

Although equations (2) and (7) appear to be similar, they have very different interpretations. Equation (2) is the familiar Poisson probability distribution function that gives the absolute probabilities of obtaining different N for a given S and B. In equation (7) the roles of N and S are reversed, and we now regard this function as the continuous probability distribution for the source flux, S, given N and B. The interpretation in the Bayesian approach is that $f_{N,B}(S)$ gives the relative probabilities that the observed number of counts could have been produced by sources of different flux S. This distinction between equations (2) and (7) is the crucial difference between the Bayesian and classical methods.

Confidence limits for S in the Bayesian case are obtained by simply integrating the probability distribution function, $f_{N,B}$, over S and solving numerically for $S_{\min}$ and $S_{\max}$ such that

$$\int_{S_{\min}}^{S_{\max}} f_{N,B}(S) dS = \mathrm{CL} . \quad (9)$$

The interpretation of the confidence interval defined by $S_{\max}$ and $S_{\min}$ is that the probability of the source flux lying between these limits is CL. The Bayesian confidence interval places limits on the possible source fluxes that could produce the observation, and therefore makes a statement about the source population rather than one about the population of observations. We believe that this Bayesian definition most closely agrees with the intuitive concept of confidence intervals held by most astronomers, who wish to know the probability that their model parameters lie within their confidence interval rather than the probability that their data generated confidence intervals that contain these parameters. This distinction is subtle, but it has important implications for some special cases, as we show in § 2.3.

From equations (7) and (9), one can see that nonzero background rates are incorporated naturally into the Bayesian formulation of confidence intervals. This is in contrast to the classical case, where a correct method for including a nonzero background is not obvious. The imposition of nonnegative source flux in the Bayesian case occurs through the integration limits in equation (8) defining the normalization constant and through the use of a nonnegative lower limit of integration in equation (9).

As we discussed above, the choice of confidence limits satisfying equation (9) is somewhat arbitrary. We choose to select $S_{\min}$ and $S_{\max}$ such that the size of the confidence interval $S_{\max} - S_{\min}$ is minimized for a given confidence level. Using the definition of the confidence interval, this implies that

$$\frac{\partial S_{\min}}{\partial S_{\max}} = 1 . \quad (10)$$

By differentiating equation (9) with respect to $S_{\max}$ and treating CL as a constant, we obtain

$$\frac{\partial S_{\min}}{\partial S_{\max}} = \frac{f_{N,B}(S_{\max})}{f_{N,B}(S_{\min})} . \quad (11)$$

Combining equations (10) and (11), the condition for the minimal confidence interval is $f_{N,B}(S_{\max}) = f_{N,B}(S_{\min})$. Since $f_{N,B}(S)$ has a single local maximum at $S = B + N$, there is only one way to choose $S_{\max}$ and $S_{\min}$ so that equations (9) and (10) are satisfied. Upper and lower limits that satisfy equations (9) and (10) can be easily found if equation (9) is integrated

[2] The functional form of $f_{N,B}(S)$ is similar to the Poisson distribution function $P(N)$. The normalization constant, C, is required in eq. (7) because, while the sum of $P(N)$ over all N equals one, the integral of $P(N)$ over all S for a given N does not equal one if $B \neq 0$.

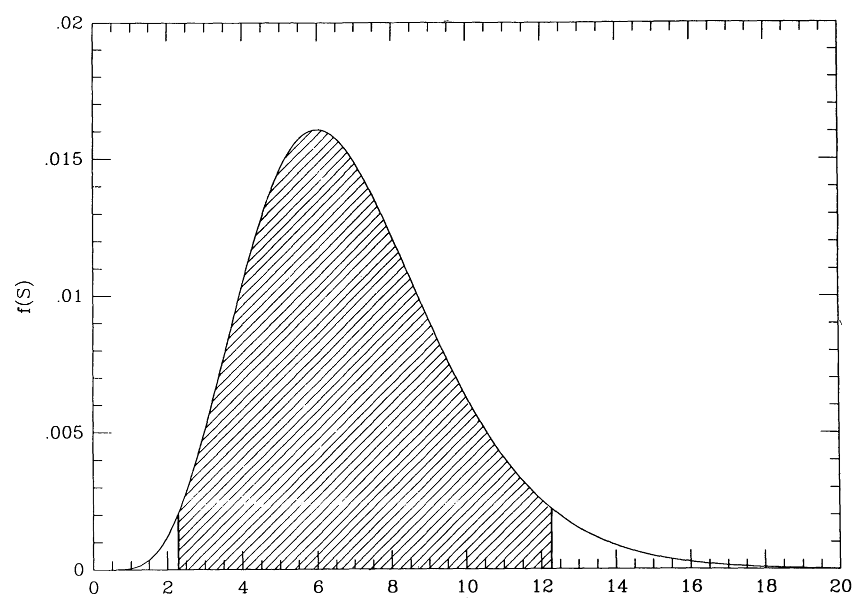

FIG. 3.—Bayesian posterior distribution function $f_{N,B}(S)$ as a function of S for $N = 6$ and $B = 0$. The miminal 95% confidence interval for S is indicated by the shaded region.

numerically by starting at the most probable value and integrating in both directions, always choosing to sum the side with the higher probability, until the desired confidence level is reached. This solution is illustrated in Figure 3, which shows the probability distribution function, $f(S)$, for $N = 6$, $B = 0$, the same case discussed in the example of the previous section. The shaded area is equal to a confidence level of 95%. Note that the area under the curve from $S = 0$ to $S_{\min}$ is less than that from $S_{\max}$ to infinity. This is true in general for Bayesian minimal confidence intervals with Poisson-distributed data.

Helene has solved equation (9) numerically for a range of B and N. However, many astrophysically important cases have been left out of his tables, including cases with $N < B$ (this case occurs when a statistical fluctuation in the background results in the number of counts observed being less than the mean number of background counts). We have tabulated values for $S_{\min}$ and $S_{\max}$ for a wide range of N and B for CLs of 0.90, 0.95, and 0.99 in Tables 1–3.

These tables were computed for a prior that is zero for negative source values and one otherwise. We have investigated the dependence of our results on the form of the prior distribution function, and have found that the confidence limits are insensitive to its exact functional form. Figure 4 shows upper limits calculated using Bayes's theorem for three different priors,

TABLE 1

BAYESIAN CONFIDENCE INTERVALS (CL = 0.90)[a]

	Number of Observed Counts, N										
B	0.0	1.0	2.0	3.0	4.0	5.0	6.0	7.0	8.0	9.0	10.0
0.0	0.00	0.08	0.44	0.94	1.51	2.13	2.79	3.47	4.17	4.89	5.63
	2.30	3.93	5.48	6.95	8.36	9.72	11.06	12.37	13.66	14.94	16.20
0.1	0.00	0.00	0.34	0.84	1.41	2.03	2.68	3.37	4.07	4.79	5.53
	2.30	3.80	5.38	6.85	8.26	9.62	10.96	12.27	13.56	14.84	16.10
0.5	0.00	0.00	0.00	0.44	1.01	1.63	2.29	2.97	3.67	4.39	5.13
	2.30	3.51	4.84	6.42	7.85	9.22	10.56	11.87	13.16	14.44	15.70
1.0	0.00	0.00	0.00	0.00	0.52	1.13	1.78	2.47	3.17	3.89	4.63
	2.30	3.27	4.44	5.71	7.30	8.72	10.06	11.37	12.66	13.94	15.20
1.5	0.00	0.00	0.00	0.00	0.09	0.65	1.29	1.97	2.67	3.39	4.13
	2.30	3.11	4.13	5.29	6.61	8.16	9.55	10.87	12.16	13.44	14.70
2.0	0.00	0.00	0.00	0.00	0.00	0.22	0.81	1.48	2.17	2.89	3.63
	2.30	2.99	3.88	4.93	6.09	7.49	8.99	10.35	11.66	12.94	14.20
2.5	0.00	0.00	0.00	0.00	0.00	0.00	0.37	1.00	1.68	2.40	3.13
	2.30	2.91	3.68	4.62	5.69	6.86	8.34	9.80	11.14	12.43	13.70
3.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.55	1.20	1.90	2.63
	2.30	2.84	3.52	4.36	5.34	6.44	7.60	9.18	10.60	11.92	13.19
3.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.15	0.75	1.42	2.14
	2.30	2.78	3.39	4.15	5.04	6.06	7.16	8.46	9.99	11.38	12.68
4.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.33	0.96	1.66
	2.30	2.74	3.29	3.97	4.78	5.72	6.76	7.88	9.31	10.79	12.15
4.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.54	1.19
	2.30	2.70	3.20	3.81	4.55	5.42	6.39	7.46	8.58	10.14	11.57
5.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.15	0.75
	2.30	2.67	3.13	3.68	4.36	5.15	6.06	7.07	8.15	9.42	10.95
5.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.35
	2.30	2.64	3.06	3.57	4.19	4.92	5.76	6.70	7.74	8.83	10.26
6.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	2.30	2.62	3.01	3.48	4.04	4.71	5.49	6.37	7.35	8.40	9.51
6.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	2.30	2.60	2.96	3.39	3.91	4.53	5.25	6.07	6.99	8.00	9.07
7.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	2.30	2.58	2.92	3.32	3.80	4.37	5.03	5.80	6.67	7.62	8.65
7.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	2.30	2.57	2.88	3.25	3.70	4.22	4.84	5.56	6.37	7.27	8.25
8.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	2.30	2.55	2.85	3.20	3.61	4.10	4.67	5.34	6.09	6.95	7.88
8.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	2.30	2.54	2.82	3.14	3.53	3.99	4.52	5.14	5.84	6.64	7.53
9.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	2.30	2.53	2.79	3.10	3.46	3.88	4.38	4.96	5.62	6.37	7.21
9.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	2.30	2.52	2.77	3.06	3.40	3.79	4.26	4.79	5.41	6.12	6.91
10.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	2.30	2.51	2.74	3.02	3.34	3.71	4.14	4.65	5.22	5.89	6.63

[a] The top number in each pair is the lower limit, and the bottom number is the upper limit.

TABLE 2

BAYESIAN CONFIDENCE INTERVALS (CL = 0.95)[a]

	Number of Observed Counts, N										
B	0.00	1.0	2.0	3.0	4.0	5.0	6.0	7.0	8.0	9.0	10.0
0.0..........	0.00	0.04	0.30	0.71	1.21	1.76	2.35	2.97	3.62	4.29	4.98
	3.00	4.77	6.40	7.95	9.43	10.86	12.26	13.63	14.98	16.30	17.61
0.1..........	0.00	0.00	0.20	0.61	1.11	1.66	2.25	2.87	3.52	4.19	4.88
	3.00	4.65	6.30	7.85	9.33	10.76	12.16	13.53	14.88	16.20	17.51
0.5..........	0.00	0.00	0.00	0.22	0.71	1.26	1.85	2.47	3.12	3.79	4.48
	3.00	4.36	5.82	7.40	8.93	10.36	11.76	13.13	14.48	15.80	17.11
1.0..........	0.00	0.00	0.00	0.00	0.23	0.76	1.35	1.97	2.62	3.29	3.98
	3.00	4.11	5.41	6.78	8.33	9.85	11.26	12.63	13.98	15.30	16.61
1.5..........	0.00	0.00	0.00	0.00	0.00	0.30	0.86	1.48	2.12	2.79	3.48
	3.00	3.94	5.08	6.36	7.68	9.23	10.73	12.12	13.48	14.80	16.11
2.0..........	0.00	0.00	0.00	0.00	0.00	0.00	0.40	0.99	1.63	2.29	2.98
	3.00	3.82	4.82	5.98	7.24	8.54	10.12	11.60	12.97	14.30	15.61
2.5..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.52	1.14	1.80	2.48
	3.00	3.72	4.61	5.66	6.84	8.09	9.37	10.99	12.44	13.79	15.11
3.0..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.11	0.67	1.31	1.98
	3.00	3.64	4.44	5.39	6.48	7.66	8.90	10.27	11.85	13.26	14.60
3.5..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.25	0.84	1.50
	3.00	3.58	4.30	5.16	6.16	7.28	8.46	9.70	11.16	12.69	14.07
4.0..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.41	1.02
	3.00	3.53	4.18	4.97	5.89	6.92	8.05	9.25	10.48	12.03	13.51
4.5..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.02	0.58
	3.00	3.49	4.09	4.80	5.64	6.61	7.67	8.82	10.02	11.26	12.88
5.0..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.18
	3.00	3.45	4.00	4.66	5.43	6.33	7.33	8.42	9.57	10.77	12.15
5.5..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	3.00	3.42	3.93	4.53	5.25	6.08	7.01	8.05	9.15	10.32	11.52
6.0..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	3.00	3.39	3.86	4.42	5.09	5.85	6.73	7.70	8.76	9.89	11.05
6.5..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	3.00	3.37	3.81	4.33	4.94	5.66	6.47	7.39	8.39	9.47	10.61
7.0..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	3.00	3.35	3.76	4.24	4.81	5.48	6.24	7.10	8.05	9.08	10.18
7.5..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	3.00	3.33	3.71	4.17	4.70	5.32	6.03	6.84	7.74	8.72	9.78
8.0..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	3.00	3.31	3.67	4.10	4.60	5.18	5.84	6.60	7.45	8.38	9.39
8.5..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	3.00	3.29	3.64	4.04	4.51	5.05	5.67	6.38	7.18	8.06	9.03
9.0..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	3.00	3.28	3.61	3.99	4.43	4.93	5.52	6.18	6.93	7.77	8.69
9.5..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	3.00	3.27	3.58	3.94	4.35	4.83	5.38	6.00	6.71	7.50	8.38
10.0..........	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	3.00	3.26	3.55	3.89	4.28	4.73	5.25	5.84	6.50	7.25	8.08

[a] The top number in each pair is the lower limit, and the bottom number is the upper limit.

plotted as a function of the total number of counts observed. Posterior probability functions corresponding to equation (7) were derived for these cases from Bayes's theorem, and the resulting upper limits computed for the case $B = 3.0$ at the 95% confidence level. The curve labeled a is the solution to equation (9) for a constant prior, and corresponds to our Table 2 for $B = 3.0$. For comparison, we have considered two other priors that do not extend to infinity. The first of these is an exponential prior of the form $\exp(-\lambda S)$ with $\lambda = 0.05$. The upper limit calculated with this prior are shown as curve b in Figure 4. The second of these is a Lorentzian of the form $(\lambda^2 + S^2)^{-1}$ with $\lambda = 20$. Upper limits for this prior are shown as curve c in Figure 4. These curves differ only slightly from curve a, and are in both cases lower than the upper limits derived from curve a, so that the upper limits in our tables are not violated by these priors.

We conclude that the Bayesian technique produces limits that are insensitive to the exact form of the prior as long as it does not exclude the region of parameter space that is of interest. The requirement of an assumed prior is therefore not a practical liability for this case.

2.3. *Comparison of the Two Methods*

We have described the classical and Bayesian methods and have outlined strengths and weakness for both. We now discuss our reasons for preferring the Bayesian approach in

TABLE 3

BAYESIAN CONFIDENCE INTERVALS (CL = 0.99)[a]

	Number of Observed Counts, N										
B	0.0	1.0	2.0	3.0	4.0	5.0	6.0	7.0	8.0	9.0	10.0
0.0	0.00	0.01	0.13	0.39	0.75	1.17	1.65	2.16	2.70	3.27	3.87
	4.60	6.64	8.45	10.15	11.76	13.32	14.84	16.32	17.77	19.19	20.59
0.1	0.00	0.00	0.03	0.29	0.65	1.07	1.55	2.06	2.60	3.17	3.77
	4.60	6.54	8.33	10.05	11.66	13.22	14.74	16.22	17.67	19.09	20.49
0.5	0.00	0.00	0.00	0.00	0.25	0.67	1.15	1.66	2.20	2.77	3.37
	4.60	6.24	7.92	9.55	11.24	12.82	14.34	15.82	17.27	18.69	20.09
1.0	0.00	0.00	0.00	0.00	0.00	0.19	0.65	1.16	1.70	2.27	2.87
	4.60	5.99	7.51	9.07	10.61	12.24	13.82	15.32	16.77	18.19	19.59
1.5	0.00	0.00	0.00	0.00	0.00	0.00	0.18	0.67	1.21	1.77	2.37
	4.60	5.80	7.17	8.64	10.13	11.61	13.20	14.79	16.26	17.69	19.09
2.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.20	0.71	1.28	1.87
	4.60	5.66	6.89	8.25	9.68	11.13	12.57	14.15	15.73	17.18	18.59
2.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.25	0.79	1.37
	4.60	5.55	6.66	7.92	9.27	10.67	12.09	13.50	15.09	16.64	18.08
3.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.33	0.88
	4.60	5.46	6.48	7.63	8.90	10.24	11.62	13.02	14.40	16.01	17.54
3.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.42
	4.60	5.39	6.32	7.38	8.56	9.84	11.17	12.54	13.91	15.28	16.92
4.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01
	4.60	5.33	6.18	7.16	8.27	9.47	10.75	12.08	13.44	14.79	16.16
4.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	4.60	5.28	6.06	6.97	8.00	9.14	10.36	11.65	12.97	14.31	15.65
5.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	4.60	5.23	5.96	6.81	7.77	8.83	10.00	11.23	12.52	13.84	15.16
5.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	4.60	5.19	5.87	6.66	7.56	8.56	9.66	10.85	12.09	13.38	14.69
6.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	4.60	5.15	5.80	6.53	7.37	8.32	9.35	10.48	11.68	12.93	14.22
6.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	4.60	5.12	5.72	6.42	7.20	8.09	9.07	10.15	11.30	12.51	13.77
7.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	4.60	5.10	5.66	6.31	7.05	7.89	8.82	9.84	10.94	12.11	13.33
7.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	4.60	5.07	5.61	6.22	6.92	7.71	8.58	9.55	10.60	11.73	12.91
8.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	4.60	5.05	5.56	6.14	6.80	7.54	8.37	9.29	10.29	11.37	12.51
8.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	4.60	5.03	5.51	6.06	6.68	7.39	8.17	9.04	9.99	11.03	12.13
9.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	4.60	5.01	5.47	5.99	6.58	7.25	7.99	8.82	9.73	10.71	11.77
9.5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	4.60	4.99	5.43	5.93	6.49	7.12	7.82	8.61	9.48	10.41	11.44
10.0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	4.60	4.98	5.40	5.87	6.41	7.00	7.67	8.24	9.24	10.14	11.12

[a] The top number in each pair is the lower limit, and the bottom number is the upper limit.

certain cases. We will give two examples and compare the limits calculated by each method in both cases.

We first consider the 90% confidence upper limits obtained by each method for the case $B = 0.5$. In Figure 5*a* we have plotted the upper limits calculated using the Bayesian method assuming a constant prior (curve *a*) as a function of the observed number of counts, N. For comparison, upper limits calculated using the classical method with the minimal confidence interval are plotted as curve *b*. Note that, for $N \geq 2$, the two methods give nearly identical results, but that the results diverge for $N = 1$ or $N = 0$. If the background rate is small, there is little difference between the two methods.

We have argued before that the smallest confidence interval is to be preferred, since it sets the tightest constraints on model parameters. From Figure 5*a*, it would appear that the classical limit is preferable to the Bayesian limits for small N by this reasoning. To see why we come to the opposite conclusion, consider the 90% confidence intervals for the case of $B = 4.0$ (Fig. 5*b*). The Bayesian upper limits are plotted in curve *a*. The upper limits in curve *b* are the minimal classical upper limits. (Again, note that curves *a* and *b* are identical for large N but diverge for small N.) For comparison, the upper limits computed using the method described in § 2.1 from the tables in Gehrels (1986) (central limits) are plotted as curve *c*. (The dif-

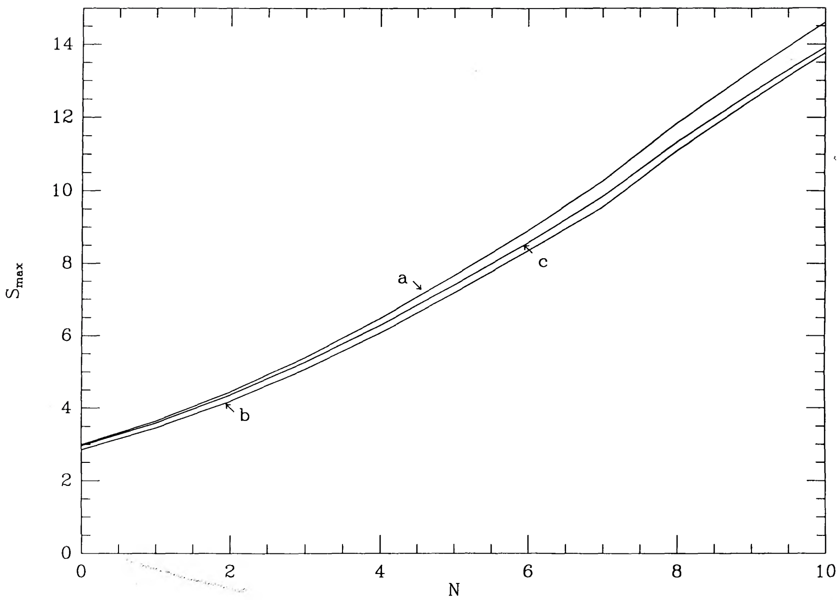

FIG. 4.—Upper limits computed using the Bayesian technique for different priors. See text for a description of the different curves.

ference between curves *b* and *c* reflects the difference between minimal and central confidence limits). Note that for the classical method (curves *b* and *c*) the upper limit to the mean source rate is zero if $N = 0$ or $N = 1$, whereas the Bayesian method always has a nonzero upper limit. The classical method does not continue smoothly into the regime where the number of counts observed is well below the mean background. The probability of detecting 0 and 1 counts in a given observation is not unreasonably small for this case (about 0.09 for $B = 4.0$ and $S = 0$), and merely requires a downward fluctuation in both the source and the background counts for that time interval. In the classical case, this fluctuation has the effect of producing a source upper limit of *exactly* zero, an intuitively unreasonable result. However, because the probability of obtaining this observation is less than 10%, this vanishingly small confidence interval satisfies the classical definition of confidence intervals. We argue that a downward fluctuation in the number of *background* counts detected should not have the effect of decreasing the uncertainty with which the *source* flux is determined, regardless of what the true source flux is.

This property of the classical upper limits has practical repercussions when one wishes to use these data in conjunction with data from other wavelength bands to constrain model parameters. Because this data point is assigned zero uncertainty, the fitting process will be driven by this point and will heavily weight the model toward a data point that should be only an upper limit. In our opinion, this is a serious weakness of the classical method.

A related objection to the results obtained with the classical method is that the limits calculated for the case $N = 0$ depend on B. The case $N = 0$ is the only case where we know exactly how many background counts were observed. It should not make any difference whether the mean background rate is $B = 1$ or $B = 20$: the limits on the source flux should be the same, since we know exactly how many source counts and background counts were observed. The limits calculated using the Bayesian method have this property.

Finally, if there is any uncertainty in the measurement of B, there is no simple way in which to extend the classical method described by Gehrels. In some cases the background uncertainty is not negligible. The Bayesian method can easily be extended to cases where B is uncertain (Helene 1983).

These examples illustrate our reasons for preferring Bayesian confidence intervals over classical confidence intervals for this problem. The classical confidence intervals are poorly behaved in this regime, in that they produce counterintuitive results. By contrast, the Bayesian confidence intervals can be naturally formulated to include the case of nonnegligible background and merge gracefully into physically reasonable limits when the number of counts detected in an observation is less than the number expected from the mean background rate. The Bayesian method can also be extended to cases where there is some uncertainty in the measurement of B; the classical method cannot. We have shown that the main argument against the general use of Bayesian statistics, the choice of the prior, is only of minor practical importance, since the limits calculated are insensitive to its exact functional form. We therefore recommend the use of Bayesian confidence intervals for observations with very small numbers of counts in the presence of background.

3. ASTROPHYSICAL APPLICATION

In order to illustrate the practical use of this method, we will walk through an actual application, contrasting the classical and Bayesian results, and mentioning possible erroneous conclusions that might be drawn from using incorrect methods.

Burrows et al. (1989) report the results of two observations of SN 1987A using an X-ray CCD camera. The data from this experiment consist of an image of the sky in X-rays, where the location of each photon detected corresponds to a position of origin on the sky. The observed images were not completely dark, but they contained no sign of a concentration of X-rays which could be attributed to SN 1987A. (The X-ray sky is not completely dark but glows with the so called "diffuse X-ray background"; some apparent photons were no doubt also the result of charged particles mimicking X-rays striking the detector. Neither of these sources will be concentrated on the

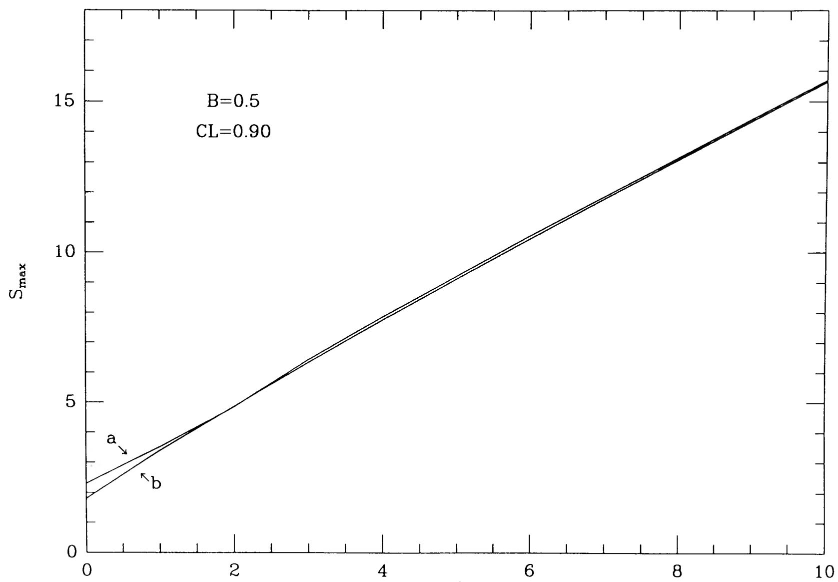

FIG. 5*a*

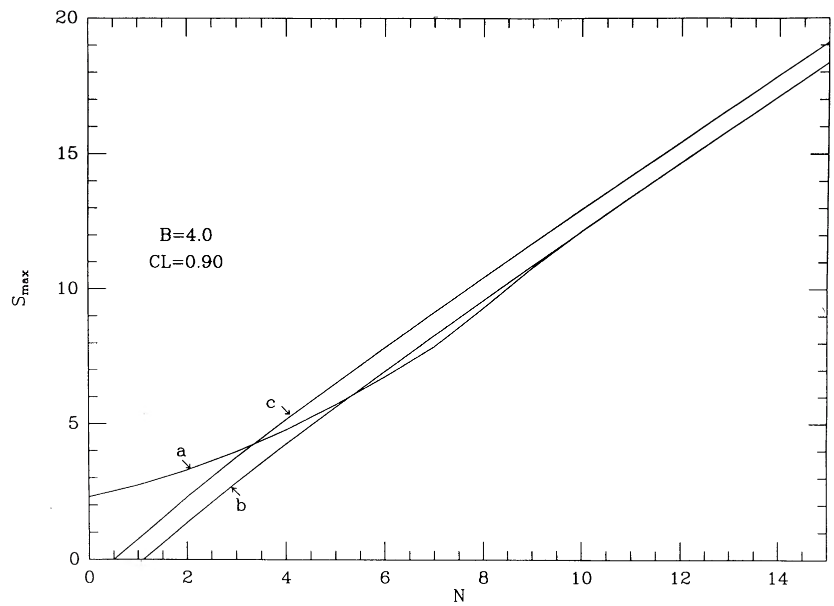

FIG. 5*b*

FIG. 5.—(*a*) Upper limits computed using the Bayesian technique (curve *a*) and the classical technique (curve *b*) for $B = 0.5$ and $CL = 0.90$. (*b*) Upper limits computed using the Bayesian technique (curve *a*) and the two classical techniques described in the text (curves *b* and *c*) for $B = 4.0$ and $CL = 0.90$.

position of the image of SN 1987A, so they represent the background, *B*, for this experiment.)

From other observations of SN 1987A there exist measurements of the flux at other energies (wavelengths). The task at hand, then, is to apply these data—the lack of detection of soft X-rays—to the other positive detections to produce a joint spectral model which is mutually consistent. The general procedure is to establish a value and uncertainty at each energy which has been measured, and fit to produce the model which most closely reproduces those values within the uncertainty allowed by the measurement.

Techniques for fitting models are outside our scope; this paper focuses on how one establishes the uncertainties when there are few counts and a nonzero background.

Returning to the X-ray CCD example, in the first observation one apparent photon was detected in the position expected for SN 1987A. The presence of additional photons scattered over the image prevents us from concluding that the single event truly came from the SN 1987A, leaving open the possibility that it was a background event.

We can, with relatively good precision, determine the mean background rate by combining the rate seen in a large number of resolution elements not including SN 1987A and averaging the results. In this case we obtain $B = 0.73 \pm 0.19$ counts.

An erroneous chain of reasoning at this point is to say, "We have determined the background in our device. The 3 σ upper limit to the background is 1.3 counts, which is more than we detected, so we conclude that we would have detected any source brighter than this. Therefore, the 3 σ upper limit to the source flux is 1.3." This argument fails to consider that the uncertainty in the mean background is much less than the fluctuations in the background, and what is relevant is the particular fluctuation which occurred in the SN 1987A resolution element, not the average over the image.

An argument which considers the background fluctuations, but also leads to an incorrect conclusion, might be: "The mean background is 0.73. The fluctuations we see follow a Poisson distribution. The probability that a result greater than 2.0 would occur is less than 5%, so we can conclude that we would have recognized a source as being real and not a background fluctuation at the 95% confidence level for an upper limit of 2.0." The error here is a confusion of the criterion for judging the existence of a source with the measurement of the intensity of the source. So, while the argument provides a valid measure of the source detection threshold for the experiment, it is incorrect for the source intensity because it fails to consider the fluctuations in the number of counts produced by the source.

The correct value for the upper limit, which considers both source and background fluctuations, can be found in Table 2 (for the 95% confidence level), by interpolating the $N = 1$ column to a background value of $B = 0.73$, for a result of 4.2. This value, based on the Bayesian approach, ignores the small uncertainty in the mean background (B). Helene (1983) provides the algorithms for incorporating this uncertainty if it is necessary for a particular case.

A better case for contrasting the classical and Bayesian methods results from the second CCD observation. Here no counts were detected in the SN 1987A region, with a mean background of 1.03 ± 0.19. Table 2 again provides the Bayesian 95% upper limit, which is 3.0. The classical result for this case is 1.96. The classical result is significantly lower than the Bayesian. The reason for this is that the classical approach is only concerned with ensuring that 95% of all observers in potential future measurements will obtain confidence intervals that include the true value. Because the observed $N = 0$ is less than B, the classicist can be confident that few later observations will have the low value measured in this case. Thus he need not force the value of S as high as the Bayesian must to keep the "failure rate" of his prediction down. In the limit of N substantially less than B, this leads to a classical upper limit of exactly zero.

As a final example, suppose 5 counts are detected in an observation with a mean background of 1.03 ± 0.19. It is not known how many of these 5 counts are background counts and how many are source counts. In the absence of a source, a 5 σ fluctuation in the background counts would be required to produce the 5 counts detected (from a Poisson distribution with a mean of 1.03). On the other hand, the 95% Bayesian confidence interval for the flux extends from 0.76 to 9.85 counts, covering over an order of magnitude in the source flux. Thus this observation would constitute a source detection with a high degree of confidence, but puts only weak constraints on the source flux.

There are occasions when the classical limits are appropriate. For example, if a large number of observations were made of SN 1987A the self-consistency of the measurements could be checked by comparing the classical limits. However, in the situation we have described, fitting the spectrum of the SN 1987A target, it is the uncertainty in the source flux which is critical, not the uncertainty in the observation. If the uncertainties are too small, and certainly if they are zero, the results of the later fits will be unreasonably dominated by this null result.

4. CONCLUSION

We have examined the problem of error analysis for photon-counting experiments with low numbers of counts and have compared two different solutions to the problem. We prefer the Bayesian method over the classical method, for several reasons. The Bayesian method can be applied in a straightforward manner for cases of nonzero background, whereas there are some problems in extending the classical method to cases of nonzero background. First, the classical method gives no information about the source for many cases when the mean background, B, is higher than the number of counts observed, N. Second, the upper limits for the case of $N = 0$ depend on B. Since this is the only case where we know exactly how many counts from the source were detected (zero), the dependence of S_{max} on B is not reasonable. Third, there is no straightforward way to extend the classical method to cases in which there is any uncertainty in the background. Although the Bayesian method requires that an assumption be made about the a priori relative probability of the possible experimental hypotheses, we have shown that this requirement poses no practical difficulty in our case and that the use of a constant prior for source fluxes from zero to infinity is a good approximation to any reasonable prior. Tables of confidence intervals for several confidence levels were calculated in order to make these results available for cases not included in the tables of Helene (1983, 1984).

This method is not restricted to the estimation of source fluxes. Although our discussion has been formulated in terms of finding confidence intervals for source fluxes in imaging applications, the same results can be used to apply to measurement of spectral line intensities in the presence of a continuum. The ratio of small numbers can also be treated by this method. This is relevant to such problems as determination of line ratios to calculate temperature, abundances, etc. A detailed treatment of confidence intervals for ratios in the case of zero background has been given by Helene (1984). If the background is nonzero, the upper and lower limits can still be calculated, but the functional form is very complex, even if the mean backgrounds for the numerator and denominator are the same.

A program for calculating upper and lower limits using the Bayesian method (eq. [9] with minimal confidence intervals) is available from the authors in the form of FORTRAN source code or as an IBM PC executable file.

We would like to thank Professors M. Akritas and G. J. Babu of the Department of Statistics at Penn State for their careful reading of the paper and for pointing out an error in an early version of the manuscript. We would also like to thank Professor E. D. Feigelson of the Department of Astronomy and Astrophysics for useful discussions. Finally, we acknowledge the very thoughtful reviews of this paper by Neil Gehrels and a second, anonymous, referee. This work was supported under NASA grants NAG5-100 and NAG5-1211.

REFERENCES

Barnett, V. 1973, Comparative Statistical Inference (New York: Wiley)
Bevington, P. R. 1969, Data Reduction and Error Analysis for the Physical Sciences (New York: McGraw-Hill)
Burrows, D. N., Nousek, J. A., Berthiaume, G. D., & Garmire, G. P. 1989, ApJ, 347, 1114
Cramér, H. 1945, Mathematical Methods of Statistics (Princeton: Princeton Univ. Press)
Eadie, W. T., Drijard, D., James, R. E., Roos, M., & Sadoulet, B. 1971, Statistical Methods in Experimental Physics (Amsterdam: North-Holland)
Gehrels, N. 1986, ApJ, 303, 336
Helene, O. 1983, Nucl. Instr. Meth., 212, 319
———. 1984, Nucl. Instr. Meth., 228, 120
Larson, H. J. 1974, Introduction to Probability Theory and Statistical Inference (London: Wiley)
Lindley, D. V. 1980, Introduction to Probability and Statistics from a Bayesian Viewpoint, Part 2, Inference (Cambridge: Cambridge Univ. Press)
Loredo, T. J. 1990, in Maximum Entropy and Bayesian Methods, ed. P. F. Fougere (Dordrecht: Kluwer), 81
Narayan, R., & Nityanada, R. 1986, ARA&A, 24, 127