Uncertainty intervals for polarized beam scattering asymmetry statistics

Kevin J. Coakley, Jabez J. McClelland, Michael H. Kelley, and Robert J. Celotta
National Institute of Standards and Technology, Gaithersburg, Maryland 20899

(Received 4 June 1992; accepted for publication 17 March 1993)

In many scattering experiments, the quantity of most direct physical interest is a measure of the difference between two closely related scattering signals, each generated by a Poisson scattering process. This difference is often expressed in terms of an asymmetry statistic, that is, the difference normalized to the sum of the two signals, corrected for an additive background contribution. Typically, a propagation of errors approach is used to compute confidence intervals for asymmetry. However, these confidence intervals are not reliable in general. In this work, generally accurate confidence intervals for asymmetry are obtained using a parametric bootstrap approach. Based on the observed data, data are simulated using a Monte Carlo resampling scheme. The resampled data sets satisfy a constraint that ensures that background-corrected count rates are not negative.

I. INTRODUCTION

In many areas of research, asymmetry statistics, that is, the normalized difference between two quantities, are of direct physical interest. For example, in atomic collision physics, asymmetry statistics computed from the scattering of spin-polarized electrons from atoms carry information about atomic structure.[1] In materials science studies, maps of magnetic microstructure are based on the polarization of secondary electrons emitted from the material after it is bombarded by an energetic beam of electrons.[2] To estimate these polarizations, asymmetry statistics are computed.

Many of the experiments in which asymmetries are of interest involve the counting of electrons or other particles. Generally, streams of pulses (assumed to be Poisson distributed) are counted in two experiments, one for each orientation of the spins in the system relative to some quantization axis. The number of counts measured in each experiment is associated with an intensity for each of the two spin orientations. The asymmetry is estimated by taking the ratio of the difference and the sum of the two intensities.

The estimation of the asymmetry from observed data is often complicated by an additive background signal. To correct for background, the number of background counts is measured in a third experiment and then subtracted from each of the other two principal measurements. In this paper, the case where measured background is less than or equal to the other measurements is considered. That is, the background-corrected observations are not negative. In a later paper, the case where background corrected counts are negative is treated.

In order to compute confidence intervals for the asymmetry statistic, a propagation of errors (POE) approach[3] is typically used. The intensities are assumed to be normally distributed with a standard deviation equal to the square root of the number of total counts detected. These standard deviations are propagated through the expression for the asymmetry in the standard way. Though used extensively, this POE approach can fail dramatically in some situations. For example, when the background signal is too close to the other signals, or if one of the two measured signals is very small, unreasonable results are obtained.

In this work, uncertainty intervals for the asymmetry statistic are computed using a variation of the parametric bootstrap[4] approach. In this approach, data are simulated using a Monte Carlo resampling scheme. The counts measured are modeled as independent realizations of Poisson processes with different rate parameters. Based on this parametric assumption, bootstrapped replications of the observed data are simulated. To ensure physically meaningful results, the simulated data satisfy a constraint that the background not exceed either of the two principal signals. For each simulated data set, an asymmetry statistic is computed. Based on the histogram of asymmetry statistics computed from the simulated data, lower and upper end points of an approximate 95% confidence interval are computed. For a variety of cases, the computed confidence intervals were generally accurate. The disagreement between the bootstrap and POE approaches becomes more dramatic as the background signal gets close to one of the other signals.

In this paper, the estimator of scattering asymmetry is first defined. For a high-background example, it is shown that the POE approach fails to give a meaningful uncertainty interval. A bootstrap uncertainty interval is defined and compared to the POE approach for various examples.

II. ASYMMETRY STATISTIC

Suppose that the number of scattering events for two different scenarios are measured in two independent experiments. Further, assume that each experiment lasts the same amount of time t. This assumption can be relaxed without loss of generality. The first observation N_1 can be expressed as the sum of two unobservable quantities as follows:

$$N_1 = N_1^* + N_{BG,1}^* . \tag{1}$$

Above, N_1^* represents what would have been observed if there had been no background. The number of counts due to the background is $N_{BG,1}^*$. The superscript star indicates that these quantities are unobservable. The terms on the right-hand side of Eq. (1) are realizations of Poisson processes with parameters $\lambda_1 t$ and $\lambda_{BG} t$. The second measurement is expressed as

$$N_2 = N_2^* + N_{BG,2}^*, \tag{2}$$

where two terms on the right side of Eq. (2) are independent realizations of Poisson processes with parameters $\lambda_2 t$ and $\lambda_{BG} t$. The goal is to estimate the asymmetry term

$$R = \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2}. \tag{3}$$

Note that the asymmetry R lies between -1 and $+1$. Hence, the end points of any confidence interval for asymmetry should also lie between -1 and $+1$, inclusive.

Assume that the background is measured in a third independent experiment which also lasts time t. Further, assume that the experimental conditions for the background measurement are the same as for the other experiments. Let the number of detected background counts be $N_{BG,3}$. This measurement is modeled as a realization of a Poisson process with parameter $\lambda_{BG} t$. With this third measurement, a background-corrected estimate of the asymmetry is

$$\hat{R} = \frac{(N_1 - N_{BG,3}) - (N_2 - N_{BG,3})}{(N_1 - N_{BG,3}) + (N_2 - N_{BG,3})} = \frac{N_1 - N_2}{N_1 + N_2 - 2N_{BG,3}}. \tag{4}$$

The focus of this paper is on confidence intervals for asymmetry.

Here, it is assumed that N_1 and N_2 are not less than $N_{BG,3}$; that is, background-corrected counts are not negative. In a following paper, the case where background-corrected counts are negative is treated.

As a caveat, because of dead-time effects, real data are not exactly Poisson. In this work, dead-time effects are assumed to be negligible. However, if count rates were high enough so that that dead-time effects were significant, and the duration of the dead time was known, the bootstrap approach could be modified to simulate non-Poisson data similar to the observed data. This is not pursued here.

Next, the failure of the POE method is demonstrated for a particular case.

III. PROPAGATION OF ERRORS

A. Method

In the POE approach, the background-corrected ratio is first written as

$$\hat{R} = (x-y)/(x+y-2z), \tag{5}$$

where x,y,z represent $N_1, N_2, N_{BG,3}$. Based on a first-order Taylor Series expansion of $\hat{R}$ in terms of x, y, and z, the variance of $\hat{R}$ is approximated as

$$\sigma_{\hat{R}}^2 \approx \left(\frac{\partial \hat{R}}{\partial x}\sigma_x\right)^2 + \left(\frac{\partial \hat{R}}{\partial y}\sigma_y\right)^2 + \left(\frac{\partial \hat{R}}{\partial z}\sigma_z\right)^2. \tag{6}$$

To get this approximation for the variance, the statistical fluctuations in x, y, and z are assumed to be independent. Because the three measured quantities are Poisson random variables, the variance of each is estimated by its observed value. For instance, σ_x^2 is equated to x. Thus,

$$\sigma_{\hat{R}}^2 \approx \frac{4}{(x+y-2z)^4}\,(x(y-z)^2 + y(z-x)^2 + z(x-y)^2). \tag{7}$$

The 95% confidence interval is then obtained by treating $\hat{R}$ as an approximate Gaussian random variable. Hence, the confidence interval for R is $(\hat{R} - 1.96\sigma_{\hat{R}}, \hat{R} + 1.96\sigma_{\hat{R}})$.

For many experiments, the above scheme for computing a confidence interval may be a reasonable approximation. However, for experiments where statistical fluctuations in the additive background signal are proportionally large compared to the other signals, the above procedure is not appropriate, in general. Next, a 95% POE confidence interval is computed for such an example.

B. Example

Suppose that $(N_1, N_2, N_{BG,3}) = (460,420,400)$. For this example, $\hat{R} = 0.5$, $\sigma_{\hat{R}} = 0.48$, and the POE 95% confidence interval is $(-0.44, 1.44)$. Since the actual asymmetry cannot exceed unity, this confidence interval is nonsensical.

The POE method failed for at least two reasons. First, the assumption that the random variable is Gaussian is inaccurate. Second, in the POE approach, statistical fluctuations in the background which lead to negative background-corrected counts are allowed. Hence, the end points can be outside the allowed interval for true asymmetry, which is $(-1,+1)$.

To clarify this point, note that according to Eq. (1), the measurement N_1 is the sum of the unobserved quantities N_1^* and $N_{BG,1}^*$. Thus, the observed quantity N_1 is correlated with the unobserved quantity $N_{BG,1}^*$. For example, if $N_{BG,1}^*$ is large, then so is N_1 because N_1 is never less than $N_{BG,1}^*$. In the POE method described here, this correlation is not accounted for. Because of this oversight, the end points of the POE confidence interval can lie outside the region $(-1,1)$.

For the same example, $(N_1, N_2, N_{BG,3}) = (460, 420,400)$, we next show how to get a more reasonable confidence interval using a parametric bootstrap method. In the bootstrap method, the asymmetry statistic is not modeled as a Gaussian random variable. Further, in the bootstrap approach, simulated background-corrected counts are not allowed to be negative.

IV. BOOTSTRAP

A. Method

In the parametric bootstrap method[4] artificial replications of the observed data are simulated. The observations are assumed to be realizations of Poisson processes. The kth bootstrap replication of N_1, N_2, and $N_{BG,3}$ are drawn from Poisson random number generators with means N_1, N_2, and $N_{BG,3}$. Suppose that in all, N_{BOOT} bootstrap replications of the data are simulated. Each of the N_{BOOT} sets of bootstrapped observations are simulated independently of one another. However, each of the N_{BOOT} sets satisfies a constraint that ensures that the statistical fluctuations in the background and the other signals are physically meaningful. The constraint is, for $k=1$, N_{BOOT},

$$N_{BG,3}^k \leqslant N_1^k$$

and

$$N_{BG,3}^k \leqslant N_2^k$$

and

$$2N_{BG,3}^k < N_1^k + N_2^k. \quad (8)$$

Simulated data which do not obey the above constraint are discarded. Because of this constraint, the three simulated signals are correlated with one another. For each simulated data set that satisfies the above constraint, asymmetry is computed using Eq. (4). Based on the histogram of the asymmetry statistics computed from the simulated data, a confidence interval is computed. Note that the above constraint is more powerful than the requirement that $\hat{R}$ be between -1 and 1. For instance, if simulated background is larger than both of the other signals, the denominator term in Eq. (4) will be negative even though computed asymmetry is between -1 and 1. The Eq. (8) constraint is identically equivalent to jointly requiring that computed asymmetry lies between -1 and $+1$, and that the denominator term in the asymmetry estimate be positive.

According to the Percentile method,[4] the end points of the 95% confidence interval are the 2.5% and 97.5% percentiles of the distribution of bootstrapped asymmetry statistics. However, in general, this approach does not yield the most accurate confidence interval. There are standard procedures for improving Percentile Method confidence intervals. For instance, in the Bias Corrected (BC) method of Efron[5] the discrepancy between the median of the distribution of the bootstrapped statistic, relative to the value of the observed statistic $\hat{R}$, is used to correct the confidence interval. In this paper, we compute a confidence interval using this discrepancy. However, the way we use the discrepancy to get the interval differs from the BC technique.

Some definitions must first be given before giving the confidence interval. Let $\hat{G}$ be the empirical cumulative distribution function for the bootstrap replications of the statistic; $\hat{G}(y)$ is the fraction of bootstrapped statistics which are less than or equal to y. Let Φ be the cumulative distribution function for the normal density, i.e.,

$$\Phi(z)=\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{z} e^{-x^2/2}dx. \quad (9)$$

We offer the following 95% confidence interval

$$(\hat{G}^{-1}[\Phi(-1.96+z_0)],\hat{G}^{-1}[\Phi(1.96+z_0)]), \quad (10)$$

where $\Phi(z_0)=\hat{G}(\hat{R})$. In general, to get a $1-2\alpha$ level confidence interval, -1.96 and 1.96 would be replaced with the appropriate z_α and $z_{1-\alpha}$ quantiles of the normal distribution. However, in this work, only 95% confidence intervals are computed.

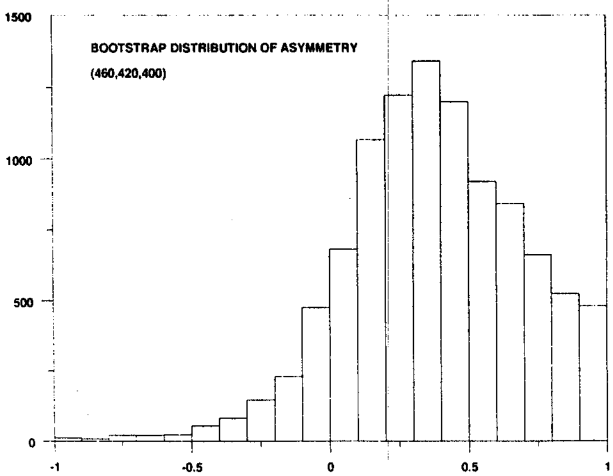

FIG. 1. Bootstrap distribution of asymmetry statistic for $(N_1,N_2,N_{BG,3})=(460,420,400)$.

We call the interval defined in Eq. (10) a constrained bias-corrected (CBC) bootstrap confidence interval. The bias correction term z_0 measures the difference between the median of the bootstrapped values of the statistic and the observed value of the statistic $\hat{R}$. If the median of the distribution $\hat{G}$ is $\hat{R}$, then $z_0=0$, and the 95% confidence interval end points are the 2.5% and 97.5% percentiles of $\hat{G}$. For the examples considered in this work, both the mean and median of the bootstrap distribution were sometimes far away from $\hat{R}$ when the background signal was high. For the example (460,420,400), the median and mean of the bootstrap distribution were 0.38 even though $\hat{R}$ was 0.5.

We note that in Efron's BC method, the multiplicative factor in front of z_0 is 2 rather than unity as in the CBC interval. For the examples studied in this work, coverage studies of the type which are described in a later section showed that the $2z_0$ correction yielded confidence intervals which were too narrow. The method developed here is not offered as the only way of proceeding. Other bootstrap strategies such as the accelerated bias-correction (BC_a) method of Efron,[5] or iterative bootstrap corrections schemes[6,7] might yield better results. DiCiccio and Romano[8] review alternative ways of computing bootstrap confidence intervals. Next, we study the accuracy of the CBC method for the example for which the POE method failed.

B. Example

Consider the same data as before, i.e., $(N_1,N_2,N_{BG,3})=(460,420,400)$. In Fig. 1, a histogram of the bootstrapped asymmetry statistics is plotted. This distribution is clearly not Gaussian. According to the CBC method, the 95% confidence interval is $(-0.11,1.00)$. This confidence interval is far more reasonable than the nonsensical $(-0.44,1.44)$ interval obtained by POE. This CBC interval is computed from a bootstrap simulation with 10 000 realizations. We note that for the example above, over 13 000 realizations were simulated in order to get 10 000

TABLE I. 95% confidence intervals provided by bootstrap (CBC) and POE methods. Data simulated with $(\lambda_1+\lambda_{BG},\lambda_2+\lambda_{BG},\lambda_{BG}) = (460,420,400)$.

Realization	N_1	N_2	$N_{BG,3}$	Bootstrap	POE
a	449	424	395	(−0.37,0.95)	(−0.47,1.08)
b	467	403	399	(0.35,1.00)	(−0.53,2.31)
c	466	402	395	(0.27,1.00)	(−0.41,2.05)
d	460	418	386	(−0.12,0.95)	(−0.25,1.04)
e	492	429	386	(0.02,0.92)	(−0.05,0.90)
f	449	401	400	(0.40,1.00)	(−1.19,3.11)
g	453	425	422	(0.07,1.00)	(−2.10,3.75)
h	469	450	407	(−0.44,0.86)	(−0.41,0.77)
i	474	430	387	(−0.10,0.91)	(−0.17,0.85)
j	501	397	394	(0.56,1.00)	(−0.02,1.91)

bootstrapped replications of the data which satisfied the constraints in Eq. (8). It took 0.4 s of CPU time on a SUN SPARCstation 1 to simulate these realizations.

To assess how well determined the CBC interval is, 30 different CBC intervals were obtained. Different intervals were obtained by simulating data using different random number seeds. The average value of the lower end point of the 95% confidence intervals was −0.11. The standard deviation of the 30 lower end points was 0.009. The upper end point was 1.0 in all 30 cases. Next, how well the CBC interval "covers" the true value of the asymmetry is investigated.

C. Coverage

In the coverage study, the true Poisson parameters were assumed to be (460,420,400). Based on this assumption, 1000 plausible data sets were simulated. Each data set satisfied the constraint that background not exceed either of the other two simulated counts. Simulated data which did not satisfy this constraint were discarded. For each of the 1000 data sets, a 95% confidence interval was computed using the CBC method. For each CBC confidence interval, 10 000 bootstrap replications were simulated. In Table I, CBC and POE confidence intervals are listed for 10 of the simulated data sets.

The true value of the asymmetry, which is 0.5, fell within the bootstrap confidence intervals 919 out of 1000 possible times. The true value was less than the lower end point of the bootstrap confidence interval 17 times. The true value exceeded the upper end point of the bootstrap confidence interval 64 times. Ideally, the true value of asymmetry should be less than the lower end point 25 out of 1000 times (2.5%). Similarly, the true value would exceed the upper end point 2.5% of the time. The observed fractions are 1.7% and 6.4%. The difference between the fraction of the time for which the true asymmetry was greater than the upper end point (6.4%) and 2.5% is statistically significant since the statistical error is 0.7%. Thus, it may be possible to improve upon the CBC method. Such improvements are not explored here.

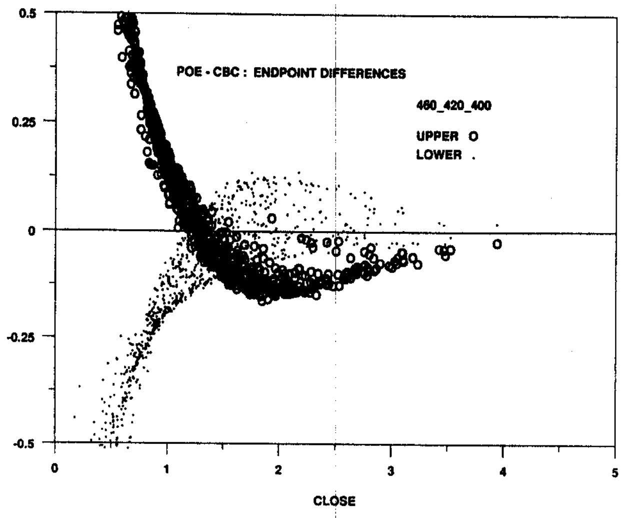

FIG. 2. End-point differences for POE and CBC methods for $(N_1,N_2,N_{BG,3}) = (460,420,400)$.

D. CBC and POE

According to Table I, the disagreement between the POE and CBC confidence intervals is greatest when the background $N_{BG,3}$ is close to either N_1 or N_2. One measure of this closeness is

$$\text{close} = \frac{\min(N_1,N_2) - N_{BG,3}}{\sqrt{\min(N_1,N_2) + N_{BG,3}}}. \tag{10}$$

The closeness measure is a normalized difference. Ideally, the denominator term should be the square root of the sum of the variance of the background counts plus the variance of the minimum of other two observed counts. We have approximated the variance of the minimum of two Poisson random variables as the minimum of the two random variables.

In Fig. 2 the difference between the end points of the POE and CBC 95% confidence interval are plotted against the above measure of closeness for the 1000 data sets in the coverage study for the case (460,420,400). The difference between the POE and CBC upper end point is plotted as a circle. The lower end-point difference is plotted as a dot. Only differences in the range (−0.5,+0.5) are shown. Note that the disagreement is worst for low closeness values. For closeness values less than unity, the POE intervals are inflated; the upper POE end point is greater than the upper CBC end point and the lower POE end point is less than the lower CBC end point. At larger closeness values, the upper end-point difference changes sign. For the examples studied here, the POE intervals are biased relative to the CBC interval even for large closeness values. The bias decreases as the closeness statistic increases.

E. Other examples

The accuracy of the bootstrap confidence interval is studied for other choices of the Poisson parameters. For each choice of $(\lambda_1+\lambda_{BG},\lambda_2+\lambda_{BG},\lambda_{BG})$, 1000 data sets are simulated as described earlier. For each simulated data set, a bootstrap confidence interval (r_l,r_u) is computed. The

TABLE II. Estimated coverage probabilities for Bootstrap confidence intervals (r_l,r_u). Sample size is 1000.

R	$\lambda_1+\lambda_{\mathrm{BG}}$	$\lambda_2+\lambda_{\mathrm{BG}}$	λ_{BG}	$\hat{p}(r_l>R)$	$\hat{p}(r_u<R)$	$\hat{p}(r_l\leqslant R\leqslant r_u)$
0.1	21	19	10	0.045	0.042	0.913
0.1	160	140	50	0.030	0.030	0.940
0.2	15	10	0	0.025	0.011	0.964
0.2	30	20	0	0.026	0.023	0.951
0.5	15	5	0	0.039	0.017	0.944
0.5	60	20	0	0.032	0.017	0.951
0.5	80	40	20	0.039	0.023	0.938
0.5	160	120	100	0.032	0.029	0.939
0.5	460	420	400	0.017	0.064	0.919
0.5	4600	4200	4000	0.030	0.018	0.952
0.5	1060	1020	1000	0.013	0.136	0.851
0.9	95	5	0	0.051	0.007	0.942
0.9	190	10	0	0.029	0.060	0.965
0.9	380	20	0	0.025	0.021	0.954
0.9	210	30	20	0.033	0.024	0.943
0.9	240	60	50	0.008	0.029	0.963
0.9	2400	600	500	0.030	0.020	0.950

fraction of cases where the lower end point is greater than true asymmetry is $\hat{p}(r_l>R)$. The fraction of cases where the upper end point is less than true asymmetry is $\hat{p}(r_u<R)$. The fraction of cases where true asymmetry falls in the confidence interval is $\hat{p}(r_l\leqslant R\leqslant r_u)$. These coverage statistics are listed in Table II.

In general, coverage is best when the background is farthest from the other signals. Note, that the coverage is worst for the case (1060,1040,1000). For this case, all three signals are very close. In Fig. 3, end-point differences are plotted vs closeness statistics for some of the cases listed in Table II. The curves have somewhat similar shapes; at low closeness values, the POE intervals are inflated relative to the CBC interval. The magnitude of this inflation depends on the values of the three signals $N_1,N_2,N_{\mathrm{BG},3}$. It is expected that for fixed N_2 and fixed $N_{\mathrm{BG},3}$, the magnitude of the difference between the upper end points diminishes as N_1 increases. At large closeness values, the POE end points are slightly biased relative to the CBC end points.

F. Asymptotic agreement

For any fixed set of Poisson rate parameters $(\lambda_1+\lambda_{\mathrm{BG}},\lambda_2+\lambda_{\mathrm{BG}},\lambda_{\mathrm{BG}})$ we expect that the agreement between the POE and CBC method can be improved if the duration of the experiment t is increased. This is because the numerator in the closeness statistic is $O(t)$ and denominator is $O(\sqrt{t})$. Hence, the closeness statistic is $O(\sqrt{t})$. Thus, for fixed rate parameters, the closeness statistic increases without limit as t increases. Further, as t gets large, the distribution of the asymmetry is well approximated by a Gaussian. This is so because of two reasons. First, for large t one can approximate the asymmetry statistic as the ratio of two correlated Gaussians. Second, the asymptotic distribution of the ratio of correlated Gaussians is also a Gaussian.[9,10] Hence, one of the assumptions in the POE method is asymptotically valid. Thus, according to this logic, if one does a long enough experiment, the POE method will be a valid approach. Next, we demonstrate the asymptotic agreement for a particular example.

For the case where the data was (460 000, 420 000, 400 000), the CBC and POE 95% confidence intervals

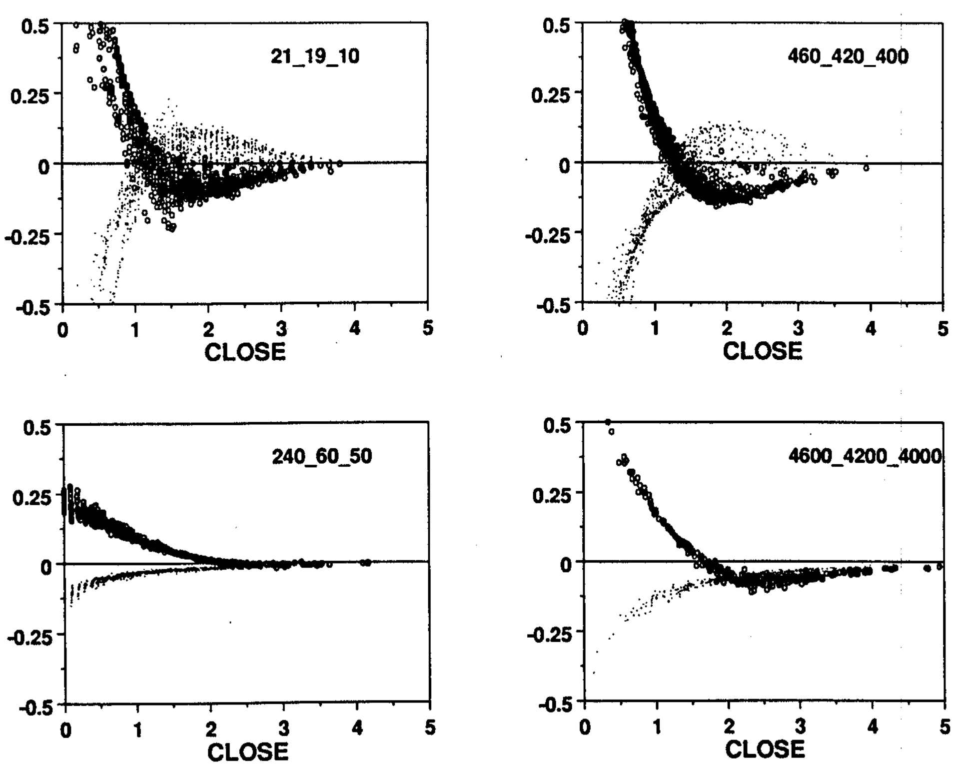

FIG. 3. End-point differences for POE and CBC methods vs closeness statistic.

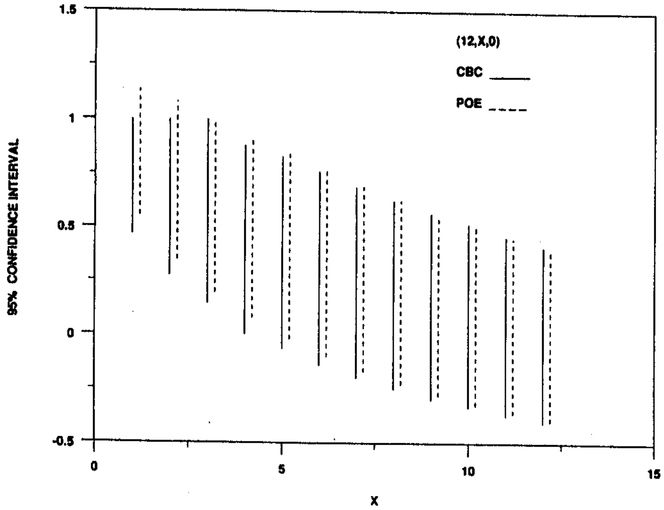

FIG. 4. Confidence intervals for POE and CBC methods for $(N_1,N_2,N_{BG,3})=(12+x,x,0)$ $x=1,2,...,12$.

were almost the same. For this case, the closeness statistic [Eq. (10)] equals 22.1. The POE 95% confidence interval was (0.470,0.530). The CBC confidence interval was (0.471,0.530). The statistical errors in determining the end points of the CBC interval were small; the standard deviations were 0.0006 and 0.0005. Also, a quantile-quantile plot indicated that the bootstrap distribution was well approximated by a Gaussian distribution.

G. No background

Now, we consider the case where there is no background signal. According to Table II, for cases where there was no background signal, the coverage of the CBC intervals was better for cases where N_1 and N_2 are larger. For instance, coverage for (380,20,0) was better than coverage for (95,5,0). In Fig. 4, CBC (solid) and POE (dashed) confidence intervals are shown for the cases where $(N_1,N_2,N_{BG,3})=(12,x,0)$, where $x=1,2,3,...,12$. In each bootstrap computation, 50 000 realizations were simulated. For values of x which are less than 3, the upper end point of the POE confidence interval exceeded unity. Note that at the lowest values of x, the POE intervals differ most from the CBC intervals. However, at larger values, the differences are not dramatic. Similar results were observed for the cases $(24,x,0)$, $(48,x,0)$, and $(96,x,0)$. In Fig. 5, the differences in the end points are plotted for these cases. All the curves have similar shapes although the magnitude of the difference in end points depends on N_1; as N_1 increases, the difference between the two approaches decreases. In closing, for the case of no background, the difference between the bootstrap and POE approach is most dramatic when both N_1 and N_2 are small.

V. DISCUSSION

Uncertainty intervals for background-corrected asymmetry statistics were computed using both a Bootstrap and a POE method. The bootstrapped replications of the observed data satisfied the constraint that background not exceed either of the two measured principal signals. The bootstrap confidence intervals were corrected for bias using a scheme based on the discrepancy between the observed asymmetry statistic and the median of the bootstrapped asymmetry statistics.

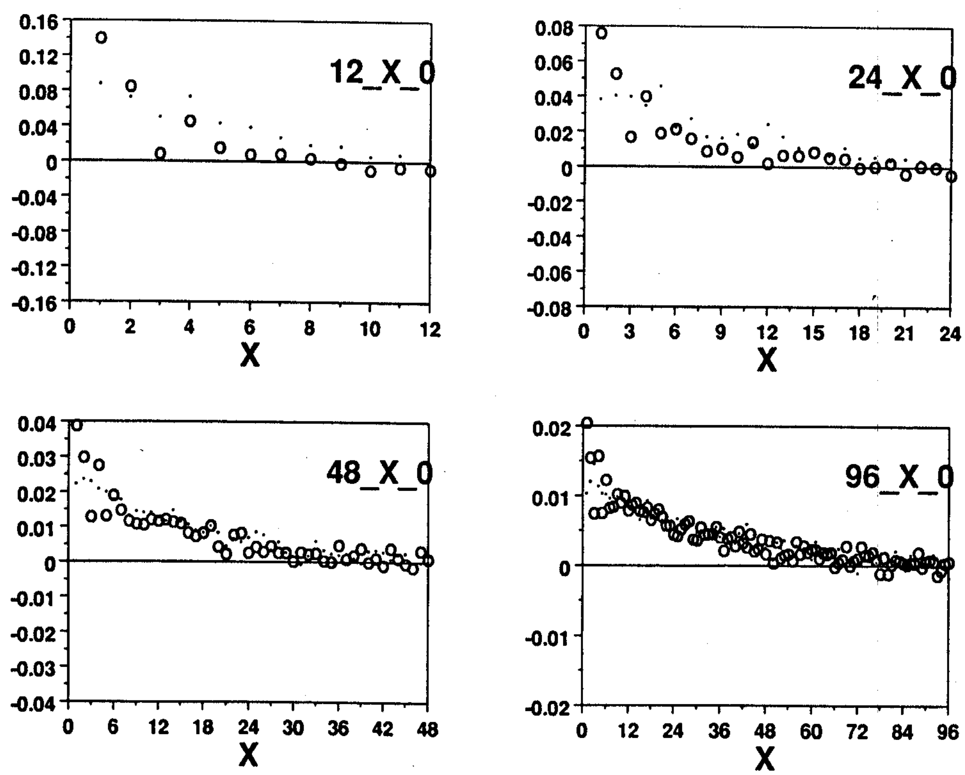

FIG. 5. End-point differences for POE and CBC methods for various no background signal cases.

For cases where the background is too close to either principal signal, the POE method failed to give reasonable results. The POE method failed for at least two reasons. First, the assumption that the random variable is Gaussian is inaccurate. Second, nonsensical fluctuations of the background were allowed. In contrast, the bootstrap method yielded generally accurate results. In general, the accuracy of the bootstrap method was best when the difference between the background and the other signals was greatest.

For cases where the background signal is close to the principal signals, the POE approach differs most dramatically from the bootstrap approach. A measure of this closeness was given. When there is no background signal, the two approaches disagree most when one of the detected principal signals is very small.

In the work considered here, the times of the three experiments to measure N_1, N_2, and N_3 are the same. For experiments where the observation times for the three experiments are not the same, all of the techniques developed here still apply. Each of the relevant terms, e.g., N_1, would be divided by the relevant observation times, e.g., t_1, in the appropriate equations. Additionally, the methods developed here can also be applied to the case where measured background exceeds one of other signals. For such cases, the asymmetry computed from the data is nonsensical. This case will be treated in a following paper.

ACKNOWLEDGMENTS

Conversations with Stefan Leigh, Keith Eberhardt, and Charles Hagwood were useful.

[1] J. J. McClelland, M. H. Kelley, and R. J. Celotta, Phys. Rev. A. **40**, 2321 (1989).

[2] M. R. Scheinfein, J. Unguris, M. H. Kelley, D. T. Pierce, and R. J. Celotta, Rev. Sci. Instrum. **61**, 2501 (1990).

[3] P. R. Bevington, *Data Reduction and Error Analysis for the Physical Sciences* (McGraw-Hill, New York, 1969).

[4] B. Efron, CBMS, SIAM-NSF, 1982.

[5] B. Efron, J. Am. Stat. Soc. **82**, 171 (1987).

[6] P. Hall and M. A. Martin, Biometrika **75**, 661 (1988).

[7] M. A. Martin, J. Am. Stat. Soc. **85**, 1105 (1991).

[8] T. DiCiccio and J. Romano, J. R. Stat. Soc. **50**, 338 (1988).

[9] D. V. Hinkley, Biometrika **56**, 635 (1969).

[10] S. Shanmugalingam, Statistician (London) **31**, 251 (1982).