Bayesian Identification of Hazardous Locations

JULIA L. HIGLE AND JAMES M. WITKOWSKI

A Bayesian analysis of accident data is used in the identification of hazardous locations. The Bayesian model used in the analysis is developed and discussed. Empirical comparisons of the results from the Bayesian analysis and from classical statistical analyses are also included. These comparisons suggest that there is an appreciable difference among the various identification techniques and that some classically based statistical techniques may be prone to err in the direction of false negatives.

One problem of ongoing interest in highway safety analysis is the identification of hazardous locations on the basis of historical data. Typically, a site is deemed hazardous if its recent accident history exceeds some specified level. One of the most common methods used in practice is to identify a site as hazardous if its accident rate over some period of time exceeds the mean accident rate over all sites in the region plus a multiple of the standard deviation of the site accident rates within that region over the same period of time. Such methods are based on the concept of confidence intervals within the context of classical statistics. The multiple used depends on the degree of confidence desired. Another commonly used technique is the rate-quality method (*1, 2*), which is based on statistical quality control procedures. This technique is used to calculate a critical accident rate, which depends on the degree of confidence desired, for each location. With the rate-quality method, a site is identified as hazardous if its observed accident rate exceeds its critical rate.

It is commonly acknowledged that because of the random variations that are inherent in accident phenomena, historical accident data do not always reflect long-term accident characteristics accurately. A site with a low accident rate (i.e., in the long run) may still have a high accident rate over a short period of time, and vice versa. Thus, the identification of hazardous locations is an inexact science at best. Regardless of the identification method used, traffic analysts will generally agree that the accident rate associated with a particular site is a random variable, a quantity that cannot be predicted with absolute certainty. Moreover, although regional accident characteristics may provide some useful information regarding the accident rate at a particular site, each site must be evaluated separately and should only be compared with sites that have similar underlying characteristics. The vast differences in accident histories that one finds among various sites suggest that the random variables used to describe the accident rates should differ from site to site.

To overcome some of the difficulties associated with the identification of hazardous locations, researchers have increasingly advocated the use of Bayesian analysis in this identification process (*3–7*). Bayesian analysis provides a framework wherein regional accident characteristics can be combined with site-specific accident histories, which results in a coherent method by which the random variables representing the accident rates at the various sites can be mathematically defined. Moreover, by using a Bayesian identification technique, one can identify hazardous sites on the basis of the probability that the accident rate exceeds some level. Such probabilistic identification methods differ both qualitatively and quantitatively from the confidence-based identification methods.

The research reported in this paper can be viewed as a complement to the research presented by Hauer and Persaud (*4–7*), although the techniques used differ substantially. These earlier papers are concerned with predicting the number of accidents that will occur at a particular location, and our research revolves around the accident rate at a particular location. An accurate prediction of the number of accidents at a particular site is invaluable in the assessment of the effectiveness of an improvement program, especially when one considers the phenomenon of regression to the mean. However, before an improvement program is implemented, one must first decide which sites require improvement. The contribution of this paper lies in the identification phase of the improvement process. Specifically, we develop a method for identifying hazardous locations on the basis of a Bayesian analysis of the accident data.

In this paper, we present the results of a Bayesian analysis of accident data from the jurisdiction of the Pima County Department of Transportation in Tucson, Arizona. The paper is divided into a discussion of the Bayesian methodology used in the study, a description of the data used, a comparison of the results of our Bayesian analysis with the results of an analysis based on classical statistical techniques, and our conclusions.

J.L. Higle, Department of Systems and Industrial Engineering, University of Arizona, Tucson, Ariz. 85721. J.M. Witkowski, JHK and Associates, 111 S. Church Street, Suite 470, Tucson, Ariz. 85701.

BAYESIAN METHODOLOGY

Bayesian analysis differs significantly from the classical statistical analysis of accident data. The motivation for the use of the Bayesian analysis is the desire to treat the actual accident rate (i.e., the number of accidents per million vehicles entering an intersection) at a particular location as a random variable and to use a combination of the regional accident characteristics and the accident history at that location to determine the probability that the location is hazardous. In this way, we hope to better utilize the available information throughout the identification process.

Our Bayesian analysis uses a two-step procedure. In the first step, we aggregate the accident histories across a number of sites (i.e., across all sites within an appropriately defined region). The result of this step is a gross estimation of the probability distribution of the accident rates across the region. We then use this regional distribution and the accident history at a particular site to obtain a refined estimation of the probability distribution associated with the accident rate at that particular site. Naturally, we obtain this refined estimation for all sites within the region, and two sites with equivalent histories will have identically refined distributions. This essentially concludes the Bayesian portion of the analysis. With the collection of refined distributions, one can now assess the probability that any given site is hazardous.

To formally describe the Bayesian identification process, we require the following notation:

$\tilde{\lambda}_i$ = accident rate at location i (note that $\tilde{\lambda}_i$ is treated as a random variable);
N_i = number of accidents at location i during the period of time in question;
V_i = number of vehicles passing through location i during the period of time in question;
$f_i(\lambda \mid N_i, V_i)$ = probability density function associated with the accident rate at location i, given the observations N_i and V_i; and
$f_R(\lambda)$ = probability density function associated with the accident rate across the region.

Thus, $f_R(\lambda)$ represents the gross estimation of the probability distribution of the accident rate across the region, and $f_i(\lambda \mid N_i, V_i)$ represents the refined estimation of the probability distribution at site i, as previously discussed. Moreover, the cumulative distribution function associated with the accident rate, $\tilde{\lambda}_i$, is given by

$$P\{\tilde{\lambda}_i \leq \hat{\lambda}\} = \int_0^{\hat{\lambda}} f_i(\lambda \mid N_i, V_i)d\lambda$$

In performing this analysis, we make the following assumptions, which are similar to those of Morin, Norden et al., Hauer and Persaud, and Glauz et al. (*1, 2, 4, 6–8*), to name but a few.

A1. At any given location, when the accident rate is known (i.e., if $\tilde{\lambda}_i = \lambda$), the actual number of accidents follows a Poisson distribution with expected value λV_i. That is,

$$P\{N_i = n \mid \tilde{\lambda}_i = \lambda, V_i\} = \frac{(\lambda V_i)^n}{n!} e^{-\lambda V_i}$$

A2. The probability distribution of the regional accident rate, $f_R(\lambda)$, is the gamma distribution.

The first assumption indicates that because the actual accident rate is explicitly treated as a random variable, the conditional distribution of the number of accidents (given the accident rate) is the Poisson distribution. The second assumption implies that

$$f_R(\lambda) = \frac{\beta^\alpha}{\Gamma(\alpha)} \lambda^{\alpha-1} e^{-\beta\lambda}$$

for some α and β. Thus, the first step associated with the Bayesian analysis, that of determining $f_R(\lambda)$, is equivalent to determining the values of α and β. There are a number of possibilities.

The most commonly used estimates are the method of moments estimates (MME), where α and β are chosen so that the mean and variance associated with the gamma distribution are equal to the mean and variance of the sample. That is, let $\bar{x}$ be the sample mean of the observed accident rates, s^2 be the sample variance of the observed accident rates, and m be the number of sites in the region. Then

$$\bar{x} = \frac{1}{m} \sum_{i=1}^{m} \frac{N_i}{V_i}$$

$$s^2 = \frac{1}{m-1} \sum_{i=1}^{m} \left(\frac{N_i}{V_i} - \bar{x}\right)^2$$

Using the MME, one selects α and β so that $\bar{x} = \alpha/\beta$ and $s^2 = \alpha/\beta^2$, or equivalently, $\beta = \bar{x}/s^2$ and $\alpha = \beta\bar{x}$.

Other commonly used estimates are the maximum likelihood estimates (MLE), where α and β are chosen so that they represent the values that are most likely to have generated the observed data. That is, if $\hat{\lambda}_i$ is the observed accident rate at site i (i.e., $\hat{\lambda}_i = N_i/V_i$), then α and β are chosen to maximize

$$\mathscr{L}(\hat{\lambda}_1, \hat{\lambda}_2, \ldots, \hat{\lambda}_m \mid \alpha, \beta) = \prod_{i=1}^{m} \frac{\beta^\alpha}{\Gamma(\alpha)} \hat{\lambda}_i^{\alpha-1} e^{-\beta\hat{\lambda}_i}$$

$$= \left\{\frac{\beta^\alpha}{\Gamma(\alpha)}\right\}^m \left[\prod_{i=1}^{m} \hat{\lambda}_i\right]^{\alpha-1} e^{-\beta\sum_{i=1}^{m}\hat{\lambda}_i}$$

The function $\mathscr{L}$ represents the likelihood function associated with the observed data when the parameters α and β are assumed. The MLE values for α and β may be obtained by solving the equations

$$\frac{\partial \mathcal{L}(\hat{\lambda}_1, \hat{\lambda}_2, \ldots, \hat{\lambda}_m \mid \alpha, \beta)}{\partial \alpha} = 0$$

$$\frac{\partial \mathcal{L}(\hat{\lambda}_1, \hat{\lambda}_2, \ldots, \hat{\lambda}_m \mid \alpha, \beta)}{\partial \beta} = 0$$

Although the MME and the MLE are among the most commonly used methods of parameter estimation, other methods for estimating α and β exist and are discussed at length by Berger (*9*).

Once values for α and β have been determined, the first step of the analysis has been completed. In the second step, the observed accident rate at each site is used in combination with the gross estimate of the regional probability distribution to obtain the site-specific probability density functions, $f_i(\lambda \mid N_i, V_i)$. These density functions are obtained using Bayes's theorem. That is,

$$f_i(\lambda \mid N_i, V_i) \propto f(N_i \mid \lambda, V_i) f_R(\lambda)$$

Within the framework of Bayesian analysis, it is well known that under Assumptions A1 and A2, the resulting probability distribution $f_i(\lambda \mid N_i, V_i)$ is a gamma distribution (*9*, *10*). Moreover, the parameters associated with this distribution, α_i and β_i, are easily obtained from the original choices of α and β and the observed data, N_i and V_i, as follows:

$$\alpha_i = \alpha + N_i$$

$$\beta_i = \beta + V_i$$

Thus, the probability density function associated with the accident rate at location $i(\tilde{\lambda}_i)$ is given by

$$f_i(\lambda \mid N_i, V_i) = \frac{\beta_i^{\alpha_i}}{\Gamma(\alpha_i)} \lambda^{\alpha_i - 1} e^{-\beta_i \lambda}$$

Note that as N_i and V_i increase, the site-specific parameters (α_i and β_i) will be largely determined by the observed data (N_i and V_i) and will become insensitive to the initial choice of α and β. As such, for each computation, it may be preferable to use the MME values rather than the MLE values, because they are substantially easier to calculate. All computations within this paper were based on the MME values of α and β.

With this collection of probability density functions, the identification of hazardous locations is now a straightforward matter. If $\bar{\lambda}$ is an upper limit on the "acceptable" accident rates, then we wish to identify a site i as hazardous if the probability is significant that $\tilde{\lambda}_i$ exceeds $\bar{\lambda}$. That is, if

$$P(\tilde{\lambda}_i > \bar{\lambda} \mid N_i, V_i) > \delta$$

where δ is some predetermined tolerance level, then site i is recognized as a hazardous location. Naturally, the appropriate values for $\bar{\lambda}$ and δ must be determined. For example, in the results section of this paper, various values of $\bar{\lambda}$ and δ are used to develop criteria for the identification of hazardous locations that are analogous to the criteria used in classically based statistical procedures. This allows a direct comparison between the results obtained from the Bayesian procedure presented in this paper and the results obtained from the classical techniques.

DATA DESCRIPTION

For the purposes of this study, 5-year (July 1981–June 1986) accident histories for signalized intersections under the jurisdiction of the Pima County Department of Transportation, in Tucson, Arizona, were used. Because significant improvement plans were undertaken during the third year (July 1983–June 1984), the data were broken into two separate sets. The first set corresponds to July 1981–June 1983, whereas the second set corresponds to July 1984–June 1986. Between July 1981 and June 1984, four intersections were signalized. Thus, the first data set includes 33 intersections, and the second data set includes 37 intersections. The two data sets were analyzed independently. For each intersection, the observed accident rate was calculated as the ratio of the total number of accidents to the total traffic volume over the 2-year period. The data are summarized in Tables 1 and 2.

One should note that the observed accident rate over each 2-year period is calculated as $N \times 10^6 / 2V \times 365$, and thus is normalized to represent the accident rate per million vehicles entering the intersection. The last two columns represent the probability that the site is hazardous on the basis of the two criteria developed in the section on results; these elements are discussed in further detail in that section.

EMPIRICAL RESULTS

In order to compare the results of an analysis based on classical statistical methods (e.g., those based on statistical confidence intervals) with the results of an analysis based on the Bayesian methodology, the two analyses must use analogous criteria in identifying a hazardous location. In practice, two commonly used criteria can be stated as follows.

C1. Site i is hazardous if the observed rate, $\hat{\lambda}_i$, exceeds the observed average rate across the region, x, with a level of confidence equal to δ.

C2. Site i is hazardous if the observed accident rate, $\hat{\lambda}_i$, exceeds the site's critical rate, which is a function of the observed regional accident rate, the traffic volume at site i, and the level of confidence desired, δ.

Typically, δ is a reasonably high number, such as 0.99, 0.95, or 0.90. C1 is the standard confidence-based crite-

TABLE 1 DATA OBTAINED FROM JULY 1981 THROUGH JUNE 1983

site number	observed accident rate (#/MVE)	number of accidents (N)	daily volume (V)	prob. (B1)	prob. (B2)
1	0.957	20	28644	0.4308	0.3861
2	1.192	46	52891	0.8684	0.8331
3	0.947	20	28950	0.4145	0.3701
4	1.437	43	40994	0.9813	0.9738
5	0.588	9	20965	0.0742	0.0614
6	1.043	14	18393	0.5402	0.5005
7	1.418	17	16422	0.8609	0.8377
8	0.779	9	15825	0.2573	0.2285
9	1.375	14	13953	0.8046	0.7776
10	1.007	18	24496	0.5058	0.4621
11	1.074	14	17863	0.5740	0.5349
12	1.174	21	24515	0.7280	0.6897
13	0.660	14	29054	0.0754	0.0608
14	1.040	22	28998	0.5632	0.5167
15	1.133	18	21773	0.6634	0.6237
16	0.675	11	22343	0.1198	0.1007
17	0.846	5	8100	0.3700	0.3411
18	0.742	11	20323	0.1897	0.1640
19	0.617	9	19989	0.0965	0.0810
20	0.282	4	19450	0.0061	0.0047
21	0.709	11	21253	0.1545	0.1318
22	1.003	8	10924	0.4805	0.4468
23	1.010	15	20360	0.5034	0.4626
24	0.088	1	15550	0.0025	0.0019
25	1.848	17	12605	0.9627	0.9543
26	0.567	6	14500	0.1138	0.0978
27	1.337	15	15376	0.7967	0.7683
28	1.471	46	42850	0.9891	0.9842
29	1.604	14	11957	0.8908	0.8727
30	1.032	15	19915	0.5311	0.4904
31	0.963	13	18502	0.4441	0.4054
32	1.184	20	23147	0.7308	0.6935
33	0.589	9	20953	0.0745	0.0616

rion, whereas C2 corresponds to the rate-quality criterion, developed by Norden et al. (*2*).

To identify hazardous locations by using Criterion C1, one must calculate both the sample mean $\bar{x}$ and the sample standard deviation s. Associated with each value of δ is a constant k_δ (e.g., $k_{0.95} = 1.645$), and if

$$\hat{\lambda}_i > \bar{x} + k_\delta s$$

then site i is said to be hazardous at the δ confidence level (*11*). To identify hazardous locations using the Bayesian methodology, a criterion that is analogous to C1 can be stated as follows:

B1. Site i is hazardous if the probability is greater than δ that its true accident rate, $\hat{\lambda}_i$, exceeds the observed average rate across the region.

Recall that the Bayesian methodology treats the accident rate at a particular location as a random variable and obtains a refined estimate of its probability distribution. As such, if

$$P\{\tilde{\lambda}_i > \bar{x} \mid N_i, V_i\} > \delta$$

then site i is said to be hazardous. Thus, the identification of hazardous locations using Criterion B1 involves the computation of

$$P\{\tilde{\lambda}_i > \bar{x} \mid N_i, V_i\} = 1 - P\{\tilde{\lambda}_i \leq \bar{x}\} = 1 - \int_0^{\bar{x}} \frac{\beta_i^{\alpha_i}}{\Gamma(\alpha_i)} \lambda^{\alpha_i - 1} e^{-\beta_i \lambda} d\lambda \quad (1)$$

If the computed value exceeds δ, site i is identified as hazardous.

Similarly, to identify hazardous sites using Criterion C2, one must calculate the regional accident rate,

$$x_R = \frac{\sum_i N_i}{\sum_i V_i}$$

For a given level of confidence, δ, the critical rate associated with location i is computed as follows:

$$\lambda c_i = x_R + k_\delta \sqrt{\frac{x_R}{V_i} + \frac{1}{2V_i}}$$

TABLE 2 DATA OBTAINED FROM JULY 1984 THROUGH JUNE 1986

site number	observed daily rate (#/MVE)	number of accidents (N)	daily volume (V)	prob. (B1)	prob. (B2)
1	0.710	15	28950	0.0727	0.0616
2	1.006	51	69450	0.3890	0.3405
3	1.047	27	35350	0.4845	0.4458
4	0.897	22	33600	0.2398	0.2118
5	1.051	23	30000	0.4885	0.4522
6	0.625	12	26300	0.0414	0.0347
7	0.875	16	25050	0.2443	0.2190
8	0.726	11	20750	0.1266	0.1114
9	0.865	13	20600	0.2544	0.2303
10	0.648	13	27500	0.0463	0.0388
11	1.028	18	24000	0.4520	0.4191
12	1.357	36	36350	0.9083	0.8903
13	1.237	31	34350	0.7828	0.7528
14	1.162	24	28300	0.6560	0.6222
15	0.704	15	29200	0.0681	0.0575
16	0.689	14	27850	0.0648	0.0548
17	0.857	6	9592	0.3238	0.3033
18	1.109	27	33350	0.5909	0.5533
19	1.318	28	29100	0.8436	0.8202
20	0.386	8	28400	0.0017	0.0013
21	1.342	29	29600	0.8669	0.8456
22	1.269	15	16200	0.7032	0.6770
23	0.947	18	26050	0.3340	0.3036
24	0.684	12	24050	0.0790	0.0679
25	2.289	34	20350	0.9998	0.9997
26	1.006	13	17700	0.4275	0.3984
27	1.300	24	25300	0.8034	0.7782
28	2.177	66	41550	1.0000	1.0000
29	1.912	30	21500	0.9951	0.9937
30	1.288	22	23400	0.7795	0.7534
31	0.988	16	22200	0.3983	0.3675
32	1.266	20	21650	0.7445	0.7173
33	1.160	26	30700	0.6630	0.6284
34	0.601	5	11400	0.1361	0.1230
35	0.813	11	18550	0.2145	0.1935
36	0.525	12	31300	0.0081	0.0064
37	0.602	9	20500	0.0569	0.0489

This critical rate is based on the assumption that the number of accidents at location i is Poisson distributed with a mean of $x_R V_i$ (*2*), which is similar to Assumption A1 in this paper. The critical rate is defined so that with this assumption, the observed accident rate will be less than or equal to the critical rate with probability δ. An investigation into the early development of this rate-quality method (*2*) suggests that an analogous criterion within the Bayesian methodology can be stated as follows:

B2. Site i is hazardous if the probability is greater than λ that its accident rate, $\tilde{\lambda}_i$, exceeds the observed regional accident rate, x_R.

That is, under Criterion B2, site i is identified as hazardous if

$$P\{\tilde{\lambda}_i > x_R \mid N_i, V_i\} > \delta$$

or equivalently, if

$$1 - \int_0^{x_R} \frac{\beta_i^{\alpha_i}}{\Gamma(\alpha_i)} \lambda^{\alpha_i - 1} e^{-\beta_i \lambda} d\lambda > \delta \qquad (2)$$

Note that the assumptions leading to Criteria C2 and B2, namely, those regarding the Poisson nature of accidents at a particular site, are very similar. The fundamental difference lies in the fact that in using C2, one implicitly assumes that the true accident rate is x_R (*2*). The authors who pioneered this method concede that the true rate (*2*) "is never known and we shall always have to be satisfied with an estimate of the expectation" (i.e., x_R). In using Criterion B2, one accounts for the inherent randomness associated with each accident rate, as reflected in Assumption A2.

In identifying hazardous locations on the basis of the Bayesian methodology (i.e., Criteria B1 and B2), one must perform the integrations identified in Equations 1 and 2. A computer program was written to numerically evaluate each of these integrals. The results of our empirical study are summarized in Figures 1 through 4. For each data set and for each value of δ (i.e., $\delta = 0.99$, $\delta = 0.95$, and $\delta = 0.90$), hazardous sites were identified on the basis of Criteria C1 and C2, corresponding to the classical statistical methods, and of the analogous Bayesian Criteria B1 and B2. These results are presented in Figures 1–4. The ele-

a. $\delta = 0.99$

	HC1	NC1
HB1	0	0
NB1	0	33

b. $\delta = 0.95$

	HC1	NC1
HB1	1 (25)	2 (4,28)
NB1	1 (29)	29

c. $\delta = 0.90$

	HC1	NC1
HB1	2 (25,28)	1 (4)
NB1	1 (29)	29

FIGURE 1 Distribution of sites based on B1 and C1 (July 1981–June 1983).

a. $\delta = 0.99$

	HC2	NC2
HB2	0	0
NB2	2 (25,28)	31

b. $\delta = 0.95$

	HC2	NC2
HB2	3 (4,25,28)	0
NB2	0	30

c. $\delta = 0.90$

	HC2	NC2
HB2	3 (4,25,28)	0
NB2	2 (7,29)	28

FIGURE 2 Distribution of sites based on B2 and C2 (July 1981–June 1983).

a. $\delta = 0.99$

	HC1	NC1
HB1	2 (25,28)	1 (29)
NB1	0	34

b. $\delta = 0.95$

	HC1	NC1
HB1	3 (25,28,29)	0
NB1	0	34

c. $\delta = 0.90$

	HC1	NC1
HB1	3 (25,28,29)	1 (12)
NB1	0	33

FIGURE 3 Distribution of sites based on B1 and C1 (July 1984–June 1986).

a. $\delta = 0.99$

	HC2	NC2
HB2	3 (25,28,29)	0
NB2	0	34

b. $\delta = 0.95$

	HC2	NC2
HB2	3 (25,28,29)	0
NB2	0	34

c. $\delta = 0.90$

	HC2	NC2
HB2	3 (25,28,29)	0
NB2	1 (12)	33

FIGURE 4 Distribution of sites based on B2 and C2 (July 1984–June 1986).

ments in the 4×4 matrices found in Figure 1 are organized as follows:

1. Columns

a. HC1 corresponds to the number of sites that were identified as hazardous on the basis of Criterion C1 (i.e., the number of sites whose observed accident rate exceeded $\bar{x} + k_\delta s$).

b. NC1 corresponds to the number of sites that were not identified as hazardous on the basis of Criterion C1.

2. Rows

a. HB1 corresponds to the number of sites that were identified as hazardous on the basis of Criterion B1 (i.e., the number of sites with $P\{\tilde{\lambda}_i > \bar{x}\} \geq \delta$).

b. NB1 corresponds to the number of sites that were not identified as hazardous on the basis of Criterion B1.

The numbers in parentheses correspond to the sites that are identified as hazardous, as represented in Tables 1 and 2.

Thus, for example, in Figure 1 we see that for July 1981–June 1983, for $\delta = 0.95$, one site (25) is identified as hazardous under both C1 and B1, two sites (4, 28) are identified as hazardous under B1 but not C1, and one site (29) is identified as hazardous under C1 but not B1. The remaining 29 sites are not identified as hazardous under either criterion.

The results of the comparison between C2 and B2 for July 1981–June 1983 are similarly arranged in Figure 2 and those for the analyses of the data collected between July 1984 and June 1983 are presented in Figures 3 and 4 for B1 and C1 and for B2 and C2, respectively.

Because there seems to be consistent disagreement between the various criteria, a discussion of the information conveyed in Figures 1–4 is in order. First, note that as expected, as δ decreases, the number of sites identified increases under all four criteria. That is, the more relaxed the identification requirement, the easier it is to be identified.

Second, there is very little difference between the sites identified by B1 and B2, the Bayesian criteria. The only difference is in Site 12, using the data collected from July 1984–June 1986. With $\delta = 0.90$, it is identified using B1 but not B2 (see Figures 3 and 4). However, the data presented in Table 2 indicate that the probability computed under Criterion B1 is 0.9083, whereas the probability computed under Criterion B2 is 0.8901. Thus, although the methods differ, the difference is not substantial. This significant agreement between the two methods is easily explained by the data. B1 uses the threshold value $\overline{\lambda} = \bar{x}$, whereas B2 uses $\overline{\lambda} = x_R$. For both data sets, x and x_R are not substantially different, as indicated by the summary statistics presented in Table 3.

One should note carefully that Criteria B1, B2, and C2 all tend to be more conservative than C1, in that they tend to identify more sites as hazardous. This suggests that C1 may be more susceptible to the identification of false negatives (i.e., those sites that are actually hazardous but are not identified as such).

In a review of Figures 3 and 4, it is clear that for the data collected between July 1984 and June 1986, the classical criteria are in relatively high agreement with their Bayesian counterparts. This is most likely due to the extremely high accident rates (2.289, 2.177, and 1.912, respectively) of three intersections (25, 28, and 29), compared with a mean accident rate of 1.0396 and a regional rate of 1.0578. Because of the extreme nature of these three accident rates, it is reasonable to expect that any justifiable procedure would identify these sites as hazardous and that all others may seem safe by comparison. Of course, when the Bayesian procedure is used, a change in the threshold value would affect the sites that are identified as hazardous. For the purposes of this study, the values of $\overline{\lambda} - \bar{x}$(B1) and $\overline{\lambda} - x_R$(B2) were chosen so that the Bayesian and classical methods could be easily compared.

TABLE 3 SUMMARY STATISTICS

Data Set	$\bar{x}$	x_R	s
June 1981 - July 1983	0.9815	1.0042	0.3756
June 1984 - July 1986	1.0396	1.0578	0.4196

The disagreement between Criteria B1 and C1 (Figures 1 and 3) is probably best explained in terms of the underlying assumptions. B1 is based on the widely accepted assumptions that accidents occur according to a Poisson distribution and that the accident rate has a gamma distribution (i.e., Assumptions A1 and A2 as stated in the section on Bayesian methodology in this paper). Similarly, Criterion C1 is based on the implicit assumption that the observed accident rates are normally distributed. The relatively large standard deviations when compared with the low sample means (e.g., $\bar{x} = 0.9815$ and $s = 0.3756$ for the July 1981–June 1983 data) combined with the fact that the accident rates must be nonnegative suggest that the normal distribution may yield an inappropriate model for these data sets.

Because Criteria B2 and C2 are based on a similar assumption (i.e., that accidents occur according to a Poisson distribution), the differences between the corresponding results, as summarized in Figures 2 and 4, are due solely to the treatments of the actual accident rate. The Bayesian method explicitly assumes that the accident rate at any given site is a random variable and accounts for this randomness in the identification process. The rate-quality method (C2) implicitly assumes that the accident rate at each site is equal to the regional rate x_R. Thus, the Bayesian method allows site-specific accident information to guide the identification process, and the rate-quality method does not.

Finally, it should be noted that Intersection 4, which was identified as hazardous by using the data collected from July 1981–June 1983 with $\delta = 0.95$ under Criteria B1, B2, and C2, underwent significant change during July 1983–June 1984. Subsequently, it was no longer identified as a hazardous intersection, and the probability that it is hazardous dropped from 0.9813 to 0.2398 on the basis of B1 or from 0.9738 to 0.2190 on the basis of B2. Clearly, these substantial drops indicate that the improvement program was successful.

CONCLUSIONS

Use of a Bayesian analysis in the identification of hazardous accident locations using accident rate data appears to be a fundamentally sound procedure, which is shown to have identification criteria analogous to those used in the classical identification scheme, although it is certainly not limited to these criteria. The Bayesian technique has the added advantage of allowing the assessment of the impact of varying the degree of confidence, δ, without requiring that the decision statistics be recomputed. Moreover, knowing the probability that the actual rate exceeds the regional rate, for example, provides added information that can be used to evaluate the trade-offs involved in deciding which sites are candidates for improvement funds.

The results presented in Figures 1 and 3 suggest that, in general, the confidence-based procedure, C1, may be inappropriate for identifying hazardous locations. Criterion C1 fails to identify as hazardous many sites that are flagged by Criteria B1, B2, and C2. The underlying assumption of normality in the distribution of the accident rate appears to cause C1 to err in the direction of false negatives. This is the least desirable characteristic for an identification procedure.

The results presented in Figures 2 and 4, combined with an analysis of the underlying assumptions, suggest that in many cases, the use of B2 may be preferable to the use of C2. This may be especially true when data are sparse or when numerous years of comparable data are not available. It is expected that the differences between B2 and C2 will be substantially reduced whenever a great deal of data are available.

REFERENCES

1. D. A. Morin. Application of Statistical Concepts to Accident Data. In *Highway Research Record 187,* HRB, National Research Council, Washington, D.C., 1967, pp. 72–79.
2. M. Norden, J. Orlansky, and H. Jacobs. Application of Statistical Quality-Control Techniques to Analysis of Highway-Accident Data. In *Bulletin 117,* HRB, National Research Council, Washington, D.C., 1956, pp. 17–31.
3. C. P. Brinkman. Toward Improved Highway Safety Analysis. *Public Roads,* Vol. 50, No. 2, 1986, pp. 49–54.
4. E. Hauer. On the Estimation of the Expected Number of Accidents. *Accident Analysis and Prevention,* Vol. 18, No. 1, 1986, pp. 1–12.
5. E. Hauer and B. Persaud. Common Bias in Before-and-After Accident Comparisons and Its Elimination. In *Transportation Research Record 905,* TRB, National Research Council, Washington, D.C., 1983, pp. 164–174.
6. E. Hauer and B. Persaud. The Problem of Identifying Hazardous Locations Using Accident Data. In *Transportation Research Record 905,* TRB, National Research Council, Washington, D.C., 1984, pp. 36–43.
7. B. Persaud and E. Hauer. A Comparison of Two Methods for De-Biasing Before-and-After Accidents Studies. In *Transportation Research Record 975,* TRB, National Research Council, Washington, D.C., 1984, pp. 43–49.
8. W. D. Glauz, K. M. Bauer, and D. J. Migletz. Expected Traffic Conflict Rates and Their Use in Predicting Accidents. In *Transportation Research Record 1026,* TRB, National Research Council, Washington, D.C., 1985, pp. 1–12.
9. J. O. Berger. *Statistical Decision Theory and Bayesian Analysis,* 2nd ed. Springer Verlag, New York, 1985.
10. M. D. DeGroot. *Optimal Statistical Decisions.* McGraw-Hill, New York, 1970.
11. J. C. Laughlin, L. E. Haefner, J. W. Hall, and D. R. Clough. *NCHRP Report 162: Methods for Evaluating Highway Safety Improvements.* TRB, National Research Council, Washington, D.C., 1975.

DISCUSSION

CARL N. MORRIS
Department of Mathematics and Statistical Sciences Center, University of Texas, Austin, Tex. 78712

The proposal by Higle and Witkowski to use hierarchical Bayesian and empirical Bayesian methods to identify hazardous highway locations is a sound one that can be expected to fare well in comparison with more traditional methods. Accident rate estimation is extremely uncertain because the number of accidents at any one intersection tends to be quite random and subject to the regression-to-the-mean phenomenon. Bayesian methods help because they effectively permit pooling of data from other relevant sites. It also makes good sense that the authors base hazard determinations on probabilistic assessments of the value $\tilde{\lambda}_i$ of the intrinsic hazard rate, as, for example, Criteria B1 and B2 require.

In Table 1 the exposure rates (daily volumes) V_i vary by a factor of 6.5 between the two extreme sites (Sites 2 and 17). Although the setting here is for Poisson data, this is otherwise analogous to the empirical Bayes estimation of means in the normal distribution case, which is now well understood (*1, 2*). In such cases the Bayesian ranking of extreme intersections differs from the observed daily rates because those extreme intersections with low volumes generally would regress to the mean more than high-volume intersections. Thus, Site 4 in Table 1 probably is more hazardous than Site 25, even though the observed rate of 1.437 MVE is less than 1.848 MVE for Site 25. This occurs because the daily volume for Site 4 is more than three times higher, 40,994 to 12,605.

It should be obvious that a methodology that properly weighs all evidence in ranking dangerous intersections is very valuable. Such features can only be revealed by using Bayesian and empirical Bayesian models, in which distributions are specified and estimated for both observed data and unobserved parameters.

Despite the virtues of the authors' general idea, there are features in their proposed methods that need further adjustment to correct for bias and to improve statistical efficiency. I will explain the difficulties partly on the basis of my own research with Olga Pendleton on accident analysis using these same Poisson-gamma models.

The authors' development fits within the General Model for Statistics (*2, 3*), which in their Poisson-gamma setting and notation is summarized in Table 4.

Note that Equations 3 and 4 correspond to the authors' Assumptions A1 and A2. The descriptive model is entirely equivalent to the inferential model (Table 5), which reverses the probabilistic conditioning and is more convenient for statistical analysis.

Expressions 5 and 7 refer to the negative binomial and gamma distributions, with their usual parameterization, whereas Expressions 6 and 8 and the square brackets signify that the means and variances are displayed.

TABLE 4 DESCRIPTIVE MODEL

Observed Data:	$N_i \mid \tilde{\lambda}_i \sim \text{Poisson}(V_i \tilde{\lambda}_i)$ $i = 1, \ldots, m$ independently.	(3)
Unobserved Parameters:	$\tilde{\lambda}_i \mid (\alpha, \beta) \sim \text{Gamma}\left(\alpha, \frac{1}{\beta}\right)$ $i = 1, \ldots, m$ independently $\phi \equiv (\alpha, \beta)$ unknown, $\alpha, \beta > 0$.	(4)

TABLE 5 INFERENTIAL MODEL

Observed Data:	$N_i \mid \alpha, \beta \sim \text{NegBin}\left(\alpha, p_i \equiv \frac{V_i}{V_i + \beta}\right)$	(5)
	$= \text{NegBin}\left[\frac{\alpha V_i}{\beta}, \frac{\alpha V_i}{\beta} + \frac{\alpha V_i^2}{\beta^2}\right]$	(6)
	$i = 1, \ldots, m$ independently.	
Unobserved Parameters:	$\tilde{\lambda}_i \mid N_i, \alpha, \beta \sim \text{Gamma}\left(\alpha + N_i, \frac{1}{\beta + V_i}\right)$	(7)
	$= \text{Gamma}\left[\frac{\alpha + N_i}{\beta + V_i}, \frac{\alpha + N_i}{(\beta + V_i)^2}\right]$	(8)
	$i = 1, \ldots, m$ independently.	

The analysis proceeds using (N_i, . . . , N_m) to estimate $\phi = (\alpha, \beta)$ from Expression 5, and then carries this information to 7 or 8 to assess the posterior distribution. The simplest method for doing this, often called "empirical Bayes," simply develops a point estimate $\hat{\phi} = (\hat{\alpha}, \hat{\beta})$ and substitutes these values into 7 or 8. [This can be risky if ($\hat{\alpha}$, $\hat{\beta}$) are not accurately estimated, an issue that could be assessed in the data example.] The authors follow this empirical Bayes approach, although not quite correctly.

Note that the marginal distribution (Expression 5) for the data N_i is negative binomial, not gamma, as the authors indicate when discussing the MLE. Thus, the maximum likelihood conditions in the second section of their paper are incorrect.

Similarly, the MME technique is improperly applied by the authors, and the estimates of the "hyperparameter" $\phi = (\alpha, \beta)$ are biased and inefficient. Define $X_i = N_i/V_i$, $X = 1/m\Sigma X_i$ and $s^2 = [1/(m - 1)]\Sigma(X_i - \bar{X})^2$, as in the second section. Then from Expression 6, $EX_i = \alpha/\beta$ and so $E\bar{X} = \alpha/\beta$, as claimed. However, from Expression 6,

$$\text{Var}(X_i) = \frac{\alpha}{\beta V_i} + \frac{\alpha}{\beta^2} \equiv \sigma^2 \tag{9}$$

and so

$$Es^2 = \frac{1}{m} \sum_1^m \sigma_i^2 \tag{10}$$

$$= \frac{\alpha}{\beta^2} + \frac{\alpha}{\beta V^*} \tag{11}$$

This exceeds α/β^2, the value claimed by the authors, by the amount $\alpha/\beta V^*$, V^* the harmonic mean of (V_i, . . ., V_m). It follows that, instead of the authors' formula, the MME is

$$\hat{\beta} = \frac{V^* \bar{X}}{V^* s^2 - \bar{X}} \qquad \hat{\alpha} = \bar{X}\hat{\beta} \tag{12}$$

Of course additional modifications are required if the denominator of β is not positive or is close to zero.

For the data set of Table 1 (1981–1983), the authors' formulas give values for α and β) of about 7 and 7, respectively, but the correct MME equations give estimates closer to 14 and 14. Thus, the model in Expression 4 should provide about twice as much information as the authors have estimated.

We could further improve the estimation. The unweighted mean $\bar{X}$ is not the best estimate of α/β because those X_i based on larger exposures V_i deserve more weight. A similar statement would apply to the use of s^2, but determining the correct weights is quite complicated in that case. The benefits of these improvements with m = 33 or 37, as in the examples, should be of second order, however, and so perhaps this use of a simple method is not costly for the data considered.

A nice feature of using the optimally weighted $\bar{X}$ is that it then could be used instead of $\bar{X}$ and X_R in Criteria B1 and B2 and would be preferable to either.

The simple device of substituting the estimates (α, β) from Expression 5 into 7 fails to acknowledge the uncertainty in knowledge of (α, β). More accurate methods, which are hard to derive [see discussion by Morris (3)] would spread out the posterior distribution. The effect of properly accounting for this in the Higle-Witkowski application would be to lower somewhat the probabilities for Criteria B1 and B2 for those locations with high accident rates.

To summarize, the Higle-Witkowski Bayesian model promises to have many advantages over standard methods for identifying hazardous locations. It will take more time before the most appropriate analytical methods are available, however. The development of such methods promises to be a rewarding and interesting task.

ACKNOWLEDGMENT

This discussion was based on research supported by FHWA at the Texas Transportation Institute, Texas A&M University, College Station, Texas. Olga Pendleton is the principal investigator.

REFERENCES

1. B. Efron and C. N. Morris. Data Analysis Using Stein's Estimator and Its Generalizations. *Journal of the American Statistical Association,* Vol. 70, 1975, pp. 311–319.
2. C. N. Morris. Parametric Empirical Bayes Inference: Theory and Applications (with discussion). *Journal of the American Statistical Association,* Vol. 78, 1983, pp. 47–65.
3. C. N. Morris. Determining the Accuracy of Bayesian Empirical Bayes Estimators in the Familiar Exponential Families. In *Statistical Decision Theory and Related Topics, IV,* Vol. 1 (S. S. Gupta and J. O. Berger, eds.), Springer-Verlag, New York, 1988, pp. 251–263.

The opinions expressed in this discussion are those of the author alone, and have not been reviewed by FHWA.

DISCUSSION

OLGA J. PENDLETON
Texas Transportation Institute, College Station, Tex. 77843.

Morris, in his discussion of this paper, has eloquently and didactically covered its weaknesses. The following discussion is offered to paraphrase and reemphasize some of his comments and to add a few comments regarding the numerical example.

Basically, this study uses a method that is inefficient. Furthermore, there are computational errors in the analysis. The authors' estimates are biased and inefficient and the method of moments technique is improperly applied, as Morris has noted. With regard to the computational errors, Morris has shown that the discrepancy in the parameter estimates is twofold when the computations are done correctly.

My second comment pertains to the numerical example. Although the results are computationally incorrect, I will base my comments on the data as they were originally presented, to show that many of the authors' claims are not supported by their numerical example.

The authors attempt to show by the numerical example that the empirical Bayesian (EB) methods (Criteria B1 and B2) are superior to two classical methods—one that assumes a normal distribution (C1) and one that assumes the more correct Poisson distribution (C2). Careful inspection of the results (Figures 1–4) does not support this claim.

In Figure 1 the classical estimate, C1, identifies one site as hazardous that the EB method fails to recognize (Site 29), and fails to identify one site that the EB method does identify (Site 4). Note that the accident rate for the site identified by the classical method is higher (1.604) than that identified by the EB method (1.437), and hence might appear to be "more logical."

In Figure 2 the classical estimate using the more correct Poisson distribution assumption, C2, identifies the same hazardous sites as does the comparable EB method (B2) and recognizes them at an even higher δ than does the EB method. The classical method identifies two sites (7 and 29) that EB does not at $\delta = 9$. Thus, one could rephrase the authors' first sentence in the second paragraph of page 5 in support of the classical estimator as follows: "Criterion B2 (EB) fails to identify as hazardous many sites that are flagged by C1 and C2." The "many" here would be only three sites; however, this is the same number of sites the authors refer to as "many" in their original statement denouncing C1 (Figure 1).

In Figures 3 and 4 we see one site (Figure 3) that the EB method identified and the classical one did not, namely, Site 12, but then in Figure 4, the classical method identifies Site 12 as hazardous when the EB method does not.

In summation, this numerical example does not show, as stated in the abstract, "that some classically based statistical techniques may be prone to err in the direction of false negatives" any more than EB methods. If, after correcting the computational error discovered by Morris, this example still fails to support such a claim, a better example should be found. Otherwise, the would-be user of this methodology is left to conclude that here we have a much more complicated procedure that does no better (maybe even worse) than the more simplistic classical methods. In its present form, this paper appears to do a disservice to a methodology that may, in fact, be superior by (a) containing mathematical errors and (b) presenting a weak example.

DISCUSSION

EMMANUEL QUAYE
Department of Civil Engineering, University of Toronto, Box 10, 321 Bloor Street West, Toronto, Ontario M5S 1S5, Canada.

In this paper the authors suggest ways of estimating the mean and the variance of true accident rates for m sites. The expressions given in the paper for calculating the variance by the method of moments and the formulation of the likelihood function are incorrect.

It can be shown (*1*) that, using the method of moments,

$$\hat{\text{Var}}\{\lambda\} = [1/(m-1)]\left[\sum_{1}^{m}(N_i^2 - N_i)/V_i^2 - (1/m)\left(\sum_{1}^{m} N_i/V_i\right)^2\right] \quad (13)$$

and not

$$[1/(m-1)]\sum_{1}^{m}\left(N_i/V_i - 1/m\sum_{1}^{m} N_i/V_i\right)^2 \quad (14)$$

as given in the paper.

As can be seen from the expressions above, Equation 2 leads to an overestimation of the variance. The difference between the two expressions is

$$[1/(m-1)]\sum_{1}^{m}(N_i/V_i^2) \quad (15)$$

Using data from Tables 1 and 2, the following comparisons, as shown below, can be made. It can be observed that there is a substantial overestimation of the variance by Equation 2.

	Variance	
Data Set	*Equation 1*	*Equation 2*
Table 1	0.0673	0.1410
Table 2	0.1164	0.1759

Similarly, it can be shown (*1*) that the correct likelihood function is

$$\prod_{1}^{m}[\alpha/(\alpha + V_iE\{\lambda\})]^{\alpha}[\Gamma(\alpha + N_i)/(\Gamma(\alpha)N_i!)] \times [V_iE\{\lambda\}/(\alpha + V_iE\{\lambda\})]^{N_i} \quad (16)$$

and not the corresponding expression given in the paper. The performance of this likelihood function was confirmed by simulation (*1*) and shown to give good results.

REFERENCE

1. E. Hauer and P. Garder. Research Into the Validity of the Traffic Conflict Technique. *Accident Analysis and Prevention*, Vol. 18, 1986, pp. 471–481.

AUTHORS' CLOSURE

First, let us state unequivocably that it is truly an honor to have these authors share their thoughts on this paper. It is with pleasure that we respond to their comments.

In our paper, we present a two-stage method for obtaining information about accident rates in a Bayesian fashion. In the first stage, we combine accident histories from various sites within a region to estimate a regional distribution, whereas in the second stage, we use site-specific data to update the regional distribution, thereby obtaining refined estimates of the distribution associated with each site. The comments in the discussions offered by Quaye and Morris concern some of the details associated with the first stage. We shall address their comments first, and save our discussion of Pendleton's criticisms for last.

FIRST-STAGE CONCERNS

Both Quaye and Morris question the manner in which we compute our estimate of the variance of the regional distribution, $f_R(\lambda)$, which has an obvious impact on our initial choice for the parameters α and β. Quaye's objection arises from the fact that we have treated the observed accident rates $\hat{\lambda}_i = N_i/V_i$, $i = 1, \ldots, m$, as our sample of observations instead of the collection of paired values (N_i, V_i). Morris points out that even when the observed rates are used, the manner in which we calculate the sample variance yields a biased estimate of the distributional variance.

Admittedly, because each observed accident rate, $\hat{\lambda}_i$, is derived from two pieces of data, N_i and V_i, Quaye's interpretation of the "sample" may be preferred to ours. However, the method of estimating the variance of the accident rates, presented by Hauer and Garder (*1*), should be considered with caution. These authors verify that the estimate presented by Quaye can provide negative estimates of the variance, a quantity that is necessarily nonnegative. This can cause difficulties, and thus the procedure should be carefully examined before it is used. Although the estimate we used cannot yield a negative sample variance, it does yield a biased estimate of the true variance, as discussed by Morris.

Regardless, herein lies a major difficulty associated with some classical identification techniques. Each of these three methods represents a reasonable or common method used to estimate the variance of the regional distribution, yet each provides a different estimate (although the Quaye and Morris estimates are in very close agreement). Because of the direct dependence of the classical techniques on the computed sample mean and variance, it is clear that the resulting set of sites identified as hazardous depends on the manner in which the data are presented and the statistics are computed. Different estimates of the variance of the accident rates across the region will necessarily lead to different sets of sites that are identified as hazardous. As a result, the tremendous differences in the tabulated values of the computed estimates of the variance presented in the Quaye discussion, corresponding to factors of 2.09 and 1.51 for the first and second data sets, respectively, might cause some concern for the integrity of some of the classical procedures.

To see that the Bayesian technique does not suffer from these same shortcomings, one need only compare the distributions obtained with the two estimation techniques. Clearly, different estimates of the variance of the accident rates will result in different estimates of both the regional and the refined distributions. Figures 5 and 6 show the regional distribution and a representative refined distribution that result from the two estimation techniques. These figures are based on the moments computed from the second data set, because it best shows the differences between the various distributions. The more peaked curve in Figure 5 (i.e., the one with the smaller variance) corresponds to the regional distribution based on the estimate of the variance obtained with Quaye's Equation 13, which is very nearly equal to Morris's estimate. In Figure 6, one can see that the refined distributions are virtually indistinguishable. As we discussed in the paper, this is because in

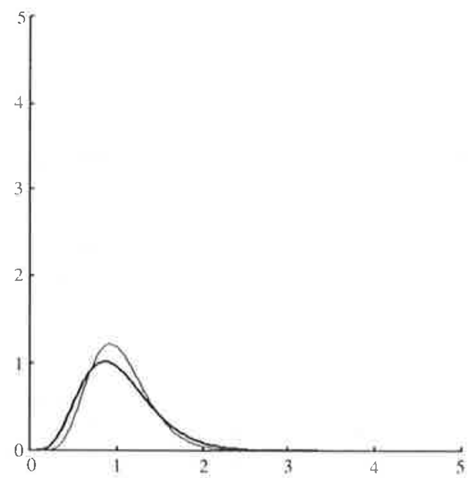

FIGURE 5 Regional distribution (July 1984–June 1986).

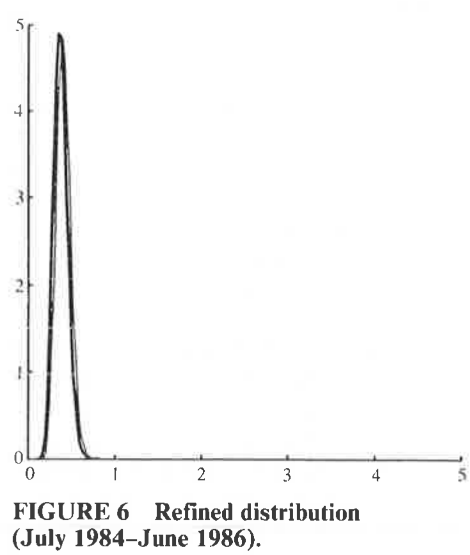

FIGURE 6 Refined distribution (July 1984–June 1986).

updating the parameters associated with the regional distribution to obtain the site-specific distribution, one necessarily overpowers the original parameters (which in our case are based on the computed sample mean and variance) with the site-specific data. Thus, the refined distributions, which provide the basis for the identifications and are therefore of paramount importance, are largely insensitive to the original parameter selection. It follows that one can be reasonably assured that although the classical methods are highly dependent on the sample variance, the Bayesian methods are not.

In addition, both Quaye and Morris question the likelihood function presented in the second section of the paper. In performing our analysis, we worked exclusively with the model that Morris has described as the descriptive model. As such, we have explicitly dealt only with the distribution of the accident rates, which are initially assumed to follow a gamma distribution. This is our Assumption A2. As a result, the likelihood function for the gamma distribution is correct as stated. Naturally, if we had performed our investigation on the basis of N_i, using the inferential model, the negative binomial model suggested by Morris would have been correct, and the likelihood function would change accordingly.

Unfortunately, Quaye's claim regarding the likelihood function (i.e., Quaye's Equation 16) is simply incorrect. A likelihood function is a mathematical entity representing the relative likeliness of the observed data (e.g., $\{(N_i, V_i)\}_{i=1}^{m}$ for a given set of distribution parameters (e.g., α and β for a gamma distribution). Thus, the maximum likelihood parameter estimates for the gamma distribution are those values of α and β that are most likely to have generated the observed data. In a very real sense, they provide the values of α and β that best fit the observed data, although the resulting theoretical mean and variance need not agree with the sample mean and variance. The "likelihood function" offered by Quaye, which also appears in the paper by Hauer and Garder (*1*), is not a true likelihood function. Instead, it is a form of a likelihood function that has been artificially constrained so that the resulting theoretical mean agrees with the observed sample mean. Thus, in general, the parameter estimates that are obtained by using it are not the maximum likelihood estimates, whereas those obtained from the expression in our paper are. In addition, the statement that somehow Quaye's Equation 16 has been "confirmed" by simulation is most disturbing indeed. As a mathematically known function, a likelihood function need not be subjected to empirical validation. Such a validation procedure suggests that there may be a "gray area" associated with the functional definition. Because it is a well-known mathematical entity, there is no such gray area requiring empirical validation.

INTERPRETATIONAL CONCERNS

Pendleton's discussion begins with a simple reiteration of Morris's comments. Because we have already discussed these comments at length, there is no need to further expand on them here. Instead, we shall focus on Pendleton's remaining criticism, which pertains to our interpretation of the results of our empirical study.

First, note that we are not in a position to claim that one technique is superior to another. To do so, one would have to know which sites are actually hazardous so that one can correctly determine which technique tends to provide correct identifications most often. Naturally, in the absence of perfect information, one can only interpret pieces of evidence or results as they become available.

The interpretation of empirical results such as those presented in our paper is necessarily subjective, and solid conclusions are often difficult to reach. Our conjecture that the classical technique, C1, "may be prone to err in the direction of false negatives" is based on an in-depth analysis of the magnitude of the differences in the levels at which sites are identified as hazardous.

To illustrate these differences, consider the three sites from Figure 1 for which techniques B1 and C1 provide differing results, namely, Sites 2, 28, and 29. From Table 1 one can obtain the probabilities computed using the Bayesian technique. A simple algebraic expression identifies the maximum confidence level at which Criterion C1 will identify these sites as hazardous, $\delta_{\max}$. Similar quantities can be obtained from the analogous sites associated with Figure 3 (i.e., Sites 12 and 29). These values are summarized in Table 6.

On the basis of this information, it seems clear that B1 provides very strong evidence that Sites 4 and 28 are hazardous (0.9813 and 0.9891, respectively), whereas C1 provides substantially weaker evidence (0.8874 and 0.9037, respectively). Similarly, although Site 12 receives a lower degree of support from B1 than do Sites 4 and 28, it still receives a substantially higher level of support from B1 (0.9083) than from C1 (0.7753). Of course, in the first data set, Site 29 receives a lower level of support from B1 (0.8908) than from C1 (0.9512), but the difference is smaller in this case. Because we believe that this type of analysis provides a better understanding of the difference between the methods than does Pendleton's method, we stand by our earlier claim. Of course, further investigation of the differences between the Bayesian and classical meth-

TABLE 6 SUMMARY OF DIFFERENCES BETWEEN B1 AND C1

Figure 1:	site	prob. (B1)	$\delta_{\max}$ (C1)
	4	0.9813	0.8874
	28	0.9891	0.9037
	29	0.8908	0.9512
Figure 3:	**site**	**prob. (B1)**	**$\delta_{\max}$ (C1)**
	12	0.9083	0.7753
	29	0.9951	0.9812

ods is called for. In addition, because Pendleton's observation regarding the apparent reordering of Sites 4 and 29 in Figure 1 is eloquently explained in Morris's discussion, we shall not endeavor to expand on his explanation here.

In conducting the research reported in this paper, it was our intention to offer a Bayesian technique for identifying hazardous intersections and to begin to understand how our technique differs from some of the classical techniques. It was not, as Pendleton states, to show that the Bayesian methods "are superior to two classical methods." Pendleton suggests that because our data set fails to provide evidence of superiority of the Bayesian method, another example should be "found" that will support such a claim. The purpose of this research was to explore the truths of a situation, not to discard or manufacture data in an effort to support a desired result.

CONCLUSION

In conclusion, we would like to agree with the closing remarks made by Morris. The application of Bayesian analyses to accident data does appear to provide a fruitful avenue of exploration. There are numerous modeling techniques to be explored. In addition, a further understanding of the differences in the results provided by Bayesian and classical identification techniques is of obvious importance. The pursuit of knowledge in this exciting field promises to offer its own rewards.