PREDICTING IBNYR EVENTS AND DELAYS

II. Discrete Time

BY WILLIAM S. JEWELL

Engineering Systems Research Center
University of California at Berkeley

ABSTRACT

An IBNYR event is one that occurs randomly during some fixed exposure interval and incurs a random delay before it is reported. A previous paper developed a continuous-time model of the IBNYR process in which both the Poisson rate at which events occur and the parameters of the delay distribution are unknown random quantities; a full-distributional Bayesian method was then developed to predict the number of unreported events. Using a numerical example, the success of this approach was shown to depend upon whether or not the occurrence dates were available in addition to the reporting dates. This paper considers the more usual practical situation in which only discretized epoch information is available; this leads to a loss of predictive accuracy, which is investigated by considering various levels of quantization for the same numerical example.

KEYWORDS

Incurred But Not Reported (IBNR) models; reporting delays; Bayesian estimation and prediction; Bayesian approximations; discrete-time models.

1. INTRODUCTION

An Incurred But Not Yet Reported claim in insurance is an event whose occurrence during some fixed exposure interval is not known until some later date because of random reporting delays. These claims, plus the Incurred But Not Fully Reported claims, which have been reported but whose cost development is incomplete, form the Incurred But Not Reported (IBNR) portfolio for a given policy exposure interval. The accurate prediction of the total number and the ultimate costs of such claims is a critical and recurring problem in many insurance lines.

In JEWELL (1989), hereinafter referred to as IBNYR-I, the author developed a continuous-time model for predicting the number of unreported IBNYR events, under the assumptions that the random (Poisson) rate of event occurrence as well as the parameters of the delay distribution are unknown.

Examination of the likelihood revealed not only a coupling between the unknown parameters for the number of occurrences and their associated random delays, but a strong dependence upon the type of epoch data available; for example, having only reporting dates but not occurrence dates led to predictions with wider variances than when both dates were available. A Bayesian development was then used to obtain a full predictive distribution and, from it, the interesting point predictors; natural conjugate priors were used for simplicity, although extensions to empirical priors are immediate. Either way, the key computational issue is the evaluation of the ratio of two integrals, for which various good approximation techniques are available. So predictive means, variances, and tail probabilities for IBNYR events are now easily obtained under continuous-time assumptions.

However, in most firms, exact epoch data is difficult to obtain, is unreliable, or, possibly, is dismissed as being unimportant. For instance, most models in the IBNR literature use quantized reporting intervals that are one year long, the same length as the usual exposure period. While this may give satisfactory results for the long-duration *cost evolution* of many casualty claims, *reporting delays* may be shorter than or comparable to the exposure interval, so that gross discretization can, as we shall see, lead to a significant loss in predictive power. Exceptions might be claims for industrial diseases (such as asbestosis) or for product liability, both of which may take a long time to develop.

The model we develop below is parallel to that of IBNYR-I, except that the reporting of dates is discretized into intervals equal to, or a submultiple of, the basic exposure interval. We model the equivalents of the first two cases of epoch data described in IBNYR-I (reporting dates always observed, occurrence dates may or may not be reported), since we know that both classical and Bayesian predictions are already bad in the other continuous cases where only occurrence dates, or only counts-to-date are available. To compare the effects of changing from continuous to quantized data, we consider the same numerical example as in the first paper.

Important references on the IBNR problem were given in IBNYR-I; supplemented by those below, they together give an overview of research in this area, most of which emphasizes point estimates for discrete-time cost-evolution models. Our results will not parallel these other efforts until a planned third paper on the "IBNR triangle" appears, in which the effect of *collateral discretized* data from several exposure periods is analyzed. As discussed in IBNYR-I, we believe that it is important to understand thoroughly the effect of various modelling assumptions upon event prediction before adding on the dynamics of random cost evolution.

2. THE MODEL

As in IBNYR-I, we assume that, during an *exposure interval* $(0, T]$, a random number of events, $\tilde{n}$, occurs according to a Poisson process with parameter λT. This implies that, given $\tilde{n} = n$, the *occurrence epochs* $(\tilde{x}_1, \tilde{x}_2, \ldots \tilde{x}_n)$ of the events are, *a priori*, independent and uniformly distributed over $(0, T]$.

Associated with each event indexed k is a random *reporting delay*, $\tilde{w}_k$, so that the actual *observation* or *reporting epochs* are $\tilde{y}_k = \tilde{x}_k + \tilde{w}_k$ $(k = 1, 2, \dots n)$. Each delay is assumed to be i.i.d. with a common probability density, $f(w \mid \theta)$, that depends upon one or more parameters, θ. It follows that, is given θ, each event pair $(\tilde{x}_k, \tilde{y}_k)$ is i.i.d. with joint density:

$$(2.1) \qquad p(x, y \mid \theta) = \frac{1}{T} f(y - x \mid \theta) \qquad (0 \leq x \leq T;\ x \leq y \leq \infty)$$

over the semi-infinite wedge-shaped region shown in Figure 1, and zero elsewhere. If we observe the reporting dates of the IBNYR events over some *observation interval* $(0, t]$, it is clear that only those pairs with $y_k \leq t$ will actually be reported, so that the total number of *reported events* will be some number R less than n.

As before, we assume that λ and θ are outcomes of the unknown random quantities $\tilde{\lambda}$ and $\tilde{\theta}$, respectively, for convenience *a priori* independent with known *prior densities*, $p(\lambda)$ and $p(\theta)$. Suppose that *epoch data* $\mathscr{D}_k$ is observed for each of the R reported events. Given these priors and the *total data*, $\mathscr{D} = \{R; \bigcup \mathscr{D}_k\}$, the *parameter estimation problem* is to determine $p(\lambda, \theta \mid \mathscr{D})$ and the *event prediction problem* is to determine $p(u \mid \mathscr{D})$, where $\tilde{u} = \tilde{n} - R$ is the unknown number of *unobserved IBNYR events* still outstanding.

To introduce the effects of discrete-time reporting, we imagine that the time axis is partitioned into equal *reporting intervals*, $\mathscr{T}_l = ((l-1)\Delta, l\Delta]$ $(l = 1, 2, \dots)$; thus $\Delta \leq T$ is the common length of the reporting intervals, and the precise values of any dates within that interval are lost. We assume that Δ is a submultiple of T, so that $I = T/\Delta$, the *quantization level*, is a positive integer. In practice, T is usually one year, and $I = 1, 2, 4$, or 12. The observation interval $(0, t]$ can now only be, say, $t = J\Delta$, with $J = 1, 2, \dots$.

We now consider two cases of quantized epoch reporting that correspond to the continuous data types I and II analyzed in IBNYR-I.

2.1. Type Iq Data. Quantized Occurrence and Reporting Dates

In this case, the continuous-time epoch data (x_k, y_k) for an observed event indexed k is mapped into $\mathscr{D}_k = (i_k, j_k)$, two positive integers indicating the reporting intervals, viz. $(i, j) \equiv (x \in \mathscr{T}_i) \cap (y \in \mathscr{T}_j)$. Obviously, $(1 \leq i \leq I)$ and $(j \geq i)$ always. Figure 1, which shows the joint partitioning of the allowed region for $I = 4$ and $t = 4.0$, gives a "tiling" that helps us to visualize the quantization. Most of the tiles are squares with sides Δ, but, if x and y are in the same interval, then (j, j) is reported in a triangular region, since $x \leq y$ always.

The probabilities associated with each tile can be expressed most easily with the aid of the function:

$$(2.2) \qquad \boldsymbol{\Phi}_h(\theta) = \frac{1}{T} \int_{(h-1)\Delta}^{h\Delta} F(w \mid \theta)\, dw \qquad (h = 1, 2, \dots)$$

$(\boldsymbol{\Phi}_0(\theta) = 0)$, which is monotonic over the integers and approaches I^{-1} for

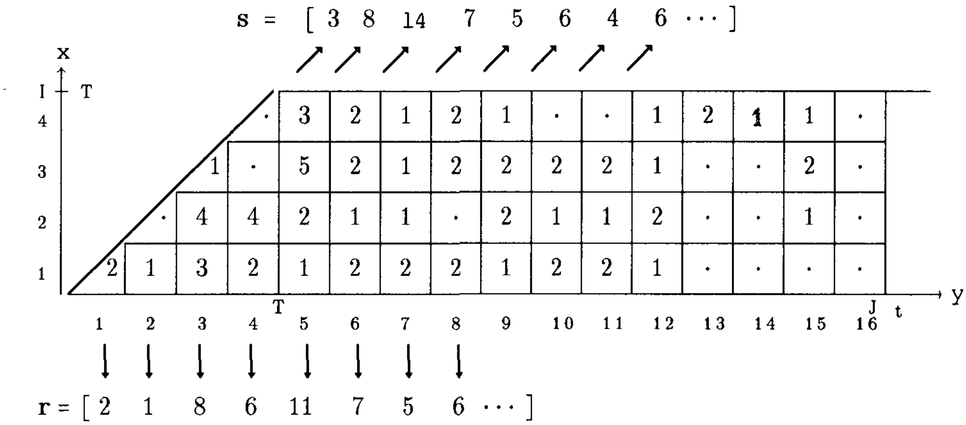

FIGURE 1. Regions of definition of quantized occurence and reporting dates, showing the distribution of 74 of the 100 events generated with $\theta = 0.5$ year^{-1}, for $I = 4$ and $J = 16$.

every θ as $h \to \infty$. Letting $\pi_{ij}(\theta) = \mathscr{P}r\,\{\mathscr{D}_k = (i,j) \mid \theta\}$ (any k) be the mass associated with tile (i,j), we find from (2.1) that:

$$\pi_{ij}(\theta) = \Phi_{j-i+1}(\theta) - \Phi_{j-i}(\theta) \qquad (1 \le i \le I)\,(j \ge i). \tag{2.3}$$

In other words, the mass of each cell along the "diagonals" with constant $h = j-i+1$ ($h = 1, 2, \ldots$) is the same, which might be expected from first principles. This is the discrete equivalent of a likelihood that depends only on $w = y-x$ ($w > 0$), as in the Type I continuous-time data models; in fact, if $(j-i) = w$ and $\Delta \to 0$, (2.3) approaches $f(w \mid \theta)\Delta^2/T$, so that events with about the same w carry the same information in the limit.

Suppose a total of R events were reported during the observation interval $(0, J\Delta]$; this includes only events for which $j \le t/\Delta$. Rather than reporting the discrete dates (i,j) for each event k, we can imagine that the epoch data represents a *distribution* of the R events into r_{ij} events for each tile (i,j), following a multinomial law with probabilities equal to $\pi_{ij}(\theta)$, normalized by dividing by the sum of probabilities over all cells in the observation interval. However, because of the structure of (2.3), the $\{r_{ij}\}$ can be accumulated over cells of equal mass on each diagonal, reducing them to the *sufficient statistics for Type Iq data*:

$$s_h = \sum_{i=1}^{\min(I,\, J+1-h)} r_{i,\, i+h-1} \qquad (h = 1, 2, \ldots J). \tag{2.4}$$

The complicated upper limit restricts the length of the observable "diagonal" elements as h approaches J and if $J < I$.

Figure 1 shows how the 74 counts for $J = 16$ and $I = 4$ in the numerical example are distributed over the cells. We find easily that $s = [3, 8, 14, 7, 5, 6, 4, 6, 5, 6, 5, 2, 2, \underline{1}, \underline{0}, \underline{0}]^T$, but note that, because of (2.4), if we decide to increase J, then the last (underlined) $I-1$ numbers would have to be increased by any new counts on their diagonals!

If we express the probabilities (2.3) in terms of:

$$\varphi_h(\theta) = \boldsymbol{\Phi}_h(\theta) - \boldsymbol{\Phi}_{h-1}(\theta) \qquad (h = 1, 2, \ldots), \tag{2.5}$$

the multinomial *conditional data likelihood,* given R and θ, is:

$$p(\bigcup \mathscr{D}_k \mid R, \theta) = \binom{R}{\boldsymbol{s}} \prod_{h=1}^{J} [\varphi_h(\theta)]^{s_h} \Big/ \left[\sum_{l=1}^{J} \min(I, J+1-l)\, \varphi_l(\theta) \right]^R, \tag{2.6}$$

where $\boldsymbol{s} = [s_1, s_2, \ldots s_J]^T$ is defined over the discrete simplex, $0 \leq s_h \leq R$, $\Sigma s_h = R$. Note how the total normalizing mass requires a weighted sum of all the $\{\varphi_h(\theta)\}$ to account for the fewer tiles near $h = J$.

2.2. Type IIq Data. Quantized Reporting Dates

The situation is somewhat simpler with only reporting epochs, $\mathscr{D}_k = (j_k)$, given for each event, which means that all event counts and probabilities are merged in each "column" of cells in Figure 1. Thus, the *sufficient statistics for Type IIq data* are $\boldsymbol{r} = [r_1, r_2, \ldots r_J]^T$, where:

$$r_j = \sum_{i=1}^{\min(I,j)} r_{ij} \qquad (j = 1, 2, \ldots J). \tag{2.7}$$

This gives $\boldsymbol{r} = [2, 1, 8, 6, 11, 7, 5, 6, 6, 5, 5, 5, 2, 1, 4, 0]^T$ from Figure 1. (2.7) can also be thought of as the result of a multinomial sorting of R events, this time with probabilities:

$$\pi_j(\theta) = \sum_{i=1}^{\min(j,I)} \pi_{ij}(\theta) = \boldsymbol{\Phi}_j(\theta) - \boldsymbol{\Phi}_{j-I}(\theta) \qquad (j = 1, 2, \ldots), \tag{2.8}$$

where the second term vanishes if $j \leq I$.

Thus, for Type IIq epoch data, (2.6) is replaced by:

$$p(\bigcup \mathscr{D}_k \mid R, \theta) = \binom{R}{\boldsymbol{r}} \prod_{j=1}^{J} [\pi_j(\theta)]^{r_j} \Big/ \left[\sum_{l=1}^{J} \pi_l(\theta) \right]^R, \tag{2.9}$$

with $\boldsymbol{r}$ defined over the discrete R-simplex. Here the normalizing mass is simpler because each $\pi_j(\theta)$ is already the sum of individual tile probabilities in each column.

As $\Delta \to 0$, (2.8) reduces to Δ times the usual probability for continuous Type II data, that is, $[F(t \mid \theta) - F((t-T)^+ \mid \theta)]\, \Delta/T$. Of course, when $I = 1$ and $\Delta = T$, the distinction between discrete Cases Iq and IIq vanishes, since $s_j = r_j = r_{1j}$, and $\varphi_j(\theta) = \pi_j(\theta) = \pi_{1j}(\theta)$.

3. DATA LIKELIHOODS AND MLE ESTIMATES

In the next two sections, we assume that Type IIq data is available; however, all formulae in which $\{r_j\}$ and $\{\pi_j(\theta)\}$ are used can be changed to Type Iq simply by replacing them with $\{s_h\}$ and $\{\varphi_h(\theta)\}$, respectively. Our first step is to uncondition (2.6) and (2.9) on R by noting that, given n and θ, n can be considered as being partitioned binomially into $\tilde{R}$ and $\tilde{u}$. At this point, it is useful to introduce the continuous cumulative probability function defined in IBNYR-I:

$$(3.1)\quad \Pi(t\,|\,\theta) = \frac{1}{T}\int_{(t-T)^+}^{t} F(w\,|\,\theta)\,dw = \sum_j \pi_j(\theta) = \sum_h \min(I, J+1-h)\,\varphi_h(\theta),$$

with $t = J\Delta$ and $T = I\Delta$, as before. Thus, $\Pi(J\Delta\,|\,\theta)$ is the mass associated with $\tilde{R}$, and each event is unreported with probability $1-\Pi(J\Delta\,|\,\theta)$. The *total data conditional likelihood* becomes the multinomial:

$$(3.2)\qquad p(\mathscr{D}\,|\,\theta, n) = \begin{pmatrix} & n \\ r & n-R \end{pmatrix} \prod_{j=1}^{J} [\pi_j(\theta)]^{r_j}\,[1-\Pi(J\Delta\,|\,\theta)]^{n-R}.$$

Let $\tau = \min(T, t) = \Delta \min(I, J)$. Then, given λ, the *total* number of events generated (but not necessarily observed) in $(0, \tau]$ follows the Poisson law with parameter $\lambda\tau$. Setting $u = n-R$ in (3.2) and marginalizing over all values of u, we obtain the final *data likelihood* in terms of the underlying parameters:

$$(3.3)\qquad p(\mathscr{D}\,|\,\lambda, \theta) = \frac{1}{\Pi(r_j!)} \prod_{j=1}^{J} [\pi_j(\theta)]^{r_j}\,(\lambda\tau)^R\, e^{-\lambda\tau\Pi(J\Delta\,|\,\theta)}.$$

(The first term is uninformative, and may be dropped). (3.2) should be compared with (4.2) in IBNYR-I (where R was written r); it might, in fact, be argued directly from it. The last term in (3.3) reflects the coupling between $\tilde{\lambda}$ and $\tilde{\theta}$ induced by the data, so that, even if they are *a priori* independent, they will become *a posteriori* dependent.

Assuming θ represents a single delay parameter, the traditional point estimates of the parameters, the MLEs $(\hat{\lambda}, \hat{\theta})$, are found from:

$$(3.4)\qquad (\hat{\lambda}\tau)\,\Sigma\pi_j(\hat{\theta}) = R; \qquad \sum \frac{d\pi_j(\hat{\theta})}{d\theta}\left[\frac{r_j}{\pi_j(\hat{\theta})} - \frac{R}{\Sigma\pi_k(\hat{\theta})}\right] = 0.$$

(All sums are over observed intervals only). The second equation can be used to find $\hat{\theta}$ numerically, which is then used in the first equation to give $\hat{\lambda}$. The ML predictor would then be $\hat{u} = \hat{\lambda}\tau - R$.

4. BAYESIAN FORMULATION

As argued in IBNYR-I, we believe that a Bayesian formulation is the natural one for IBNYR problems, since in most applications there will always be rather good prior opinion and relevant experience data about the likely values of $\tilde{\lambda}$ (which will be linked to the number of risk contracts in the portfolio), and about the parameter(s) of the delay distribution (which reflects claim filing delays, administrative flow, adjustment procedures, etc., that are common to all claims in similar lines in each company). No actuary makes estimates in a complete vacuum. The Bayesian approach also has the great advantage of giving a complete *predictive distribution*, which is essential for setting aside portfolio fluctuation reserves.

For consistency with IBNYR-I, we again assume that $\tilde{\lambda}$ and $\tilde{\theta}$ are, *a priori*, independent, with $p(\lambda)$ a $\mathcal{Gamma}$ (a, b) density. For the rest of this section, we shall leave $f(\cdot \mid \theta)$ and $p(\theta)$ in general form, later specializing to exponential delays and another Gamma prior for $\tilde{\theta}$. As in IBNYR-I, these assumptions do not simplify the *joint posterior-to-data density*, $p(\lambda, \theta \mid \mathscr{D})$, because of the coupling term, $\exp[-\lambda\tau\boldsymbol{\Pi}(J\Delta \mid \theta)]$. However, when predicting the number of unreported events, $\tilde{u} = \tilde{n} - R$, we can follow the development in IBNYR-I and show that $\tilde{u}$, given (λ, θ), is Poisson with parameter $\lambda[T - \tau\boldsymbol{\Pi}(t \mid \theta)]$, because of a fortuitous cancellation of the coupling term. Thus, the *predictive density* factors into a product of two *shaping factors*:

$$(4.1) \qquad p(u \mid \mathscr{D}) \propto h_\lambda(u \mid \mathscr{D})\, h_\theta(u \mid \mathscr{D}),$$

with:

$$(4.2) \qquad h_\lambda(u \mid \mathscr{D}) = \frac{T^u}{u!} \int \lambda^{R+u} e^{-\lambda T} p(\lambda)\, d\lambda \propto \frac{\Gamma(a+R+u)}{u!} \left[\frac{T}{b+T}\right]^u$$

with a $\mathcal{Gamma}$ (a, b) prior, and:

$$(4.3) \qquad h_\theta(u \mid \mathscr{D}) = \int \prod_{j=1}^{J} [\pi_j(\theta)]^{r_j} \left[1 - \left(\frac{\tau}{T}\right) \boldsymbol{\Pi}(t \mid \theta)\right]^u p(\theta)\, d\theta$$

for Type IIq data, with a similar form for Type Iq. Note that the first shaping factor depends only on R and $p(\lambda)$, while (4.3) depends only on $\boldsymbol{r}$ or $\boldsymbol{s}$ and $p(\theta)$. As in IBNYR-I, we refer to the term involving u in (4.3) as the *kernel*, $K(\theta)$.

Computation of the predictive distribution is most easily accomplished using the recursive form:

$$(4.4) \qquad \frac{p(u+1 \mid \mathscr{D})}{p(u \mid \mathscr{D})} = \left(\frac{a+R+u}{u+1}\right)\left(\frac{T}{b+T}\right)\left(\frac{h_\theta(u+1 \mid \mathscr{D})}{h_\theta(u \mid \mathscr{D})}\right),$$

calculated by starting with $p(0 \mid \mathscr{D}) = 1$, then normalizing when finished. With no data, the marginal (preposterior) predictive density is simply a $\mathcal{Pascal}$ $(a, T/(b+T))$ density. As in IBNYR-I, (4.4) also provides a Bayesian point estimator, the *predictive mode*, $\hat{u}(\mathscr{D})$, as the smallest integer not less than the value u^* that satisfies:

$$(4.5) \qquad u^*+1 = \left(\frac{a+R+u^*}{b+T}\right) T \left(\frac{h_0(u^*+1 \mid \mathscr{D})}{h_0(u^* \mid \mathscr{D})}\right).$$

Note that only the *ratios* of h_θ are needed in (4.4), which means that simple approximations to the integrals will give quite accurate predictive densities (TIERNEY & KADANE, 1986); (KASS, TIERNEY & KADANE, 1988). We now consider how these integrals might be approximated if the delay distribution were exponential.

5. EXPONENTIAL DELAY DISTRIBUTION

Following the example in IBNYR-I, we set $f(w \mid \theta) = \theta \exp(-\theta w)$ $(w \geq 0)$, and recall that:

$$(5.1) \qquad \boldsymbol{\Pi}(t \mid \theta) = \left(\frac{\tau}{T}\right)(1-\psi(\theta\tau)\, e^{-\theta(t-T)^+}),$$

where the properties of the useful function $\psi(x) = [1-e^{-x}]/x$ were given in that paper.

Then, from (2.2), we find:

$$(5.2) \qquad \boldsymbol{\Phi}_h = I^{-1}(1-\psi(\theta\Delta)\, e^{-(h-1)\theta\Delta}) \qquad (h = 1, 2, \ldots).$$

and the Type Iq probabilities from (2.5) are:

$$(5.3) \qquad \varphi_h(\theta) = \left\{\begin{array}{ll} I^{-1}[1-\psi(\theta\Delta)] & (h=1) \\ I^{-1}[\theta\Delta\, \psi^2(\theta\Delta)\, e^{-(h-2)\theta\Delta}] & (h = 2, 3, \ldots) \end{array}\right\}.$$

The slightly more complicated Type IIq data probabilities are found from (2.8) as:

$$(5.4) \qquad \pi_j(\theta) = \left\{\begin{array}{ll} I^{-1}[1-\psi(\theta\Delta)\, e^{-(j-1)\theta\Delta}] & (j = 1, 2, \ldots I) \\ I^{-1}[\theta T\, \psi(\theta\Delta)\, \psi(\theta T)\, e^{-(j-I-1)\theta\Delta}] & (j = I+1, I+2, \ldots) \end{array}\right\}.$$

Rewriting h_θ as in IBNYR-I:

$$(5.5) \qquad h_\theta(u \mid \mathscr{D}) = \int L(\theta \mid \mathscr{D})\, [K(\theta)]^u\, p(\theta)\, d\theta,$$

the *epoch data likelihood*, $L(\theta)$, is then expressed for Type Iq data as:

$$(5.6) \qquad L(\theta \mid \mathscr{D}) = \prod_{h=1}^{J} [\varphi_h(\theta)]^{s_h} \propto [1-\psi(\theta\Delta)]^{s_1}\, [\theta\Delta\, \psi^2(\theta\Delta)]^{R-s_1}\, e^{-M_s\theta\Delta},$$

https://doi.org/10.2143/AST.20.1.2005486 Published online by Cambridge University Press

where uninformative constants have been dropped, and M_s is the moment:

$$(5.7) \qquad M_s = \sum_{h=2}^{J} (h-2)\, s_h .$$

In other words, with exponential delays, (s_1, R, M_s) becomes the reduced set of sufficient statistics for Type Iq data. Remember that, with each new value of J, the $I-1$ most recent values of s have to be recomputed from (2.4); otherwise, there is nothing special about the choice of J relative to I.

For Type IIq data, assuming $J > I$:

$$(5.8) \qquad L(\theta \mid \mathscr{D}) = \prod_{j=1}^{J} [\pi_j(\theta)]^{r_j} \propto \prod_{j=1}^{I} [1-\psi(\theta\Delta)\, e^{-(j-1)\theta\Delta}]^{r_j} \times [\theta T\, \psi(\theta\Delta)\, \psi(\theta T)]^{R_J} e^{-M_r\theta\Delta},$$

where uninformative constants have been dropped, and:

$$(5.9) \qquad R_J = \sum_{j=I+1}^{J} r_j; \qquad M_r = \sum_{j=I+1}^{J} (j-I-1)\, r_j .$$

In this case, $(r_1, r_2, \ldots r_I; R_J; M_r)$ become the sufficient statistics. If $J \leq I$, the product term in (5.8) has an upper limit of J, the terms on the second line are dropped since $R_J = M_r = 0$, and the sufficient statistics revert to $(r_1, r_2, \ldots r_J)$. In contrast to Iq data, once all of the values in $\boldsymbol{r}$ are computed for a given I, they can be used for any J.

6. NUMERICAL EXAMPLE AND DATA ANALYSIS

To facilitate comparison with prediction using continuous data, we will use the same basic data and assumptions as in IBNYR-I, namely, that $\tilde{\lambda}$ has a $\mathscr{Gamma}$ (2, 0.02) prior density and $T = 1$, so that the no-data (marginal) prediction density is $\mathscr{Pascal}$ $(2, 1.02^{-1})$, with mean $\mathscr{E}\{\tilde{n}\} = 100$ events, mode $\hat{n} = 49$, and fractiles $n_{.05} = 16.5$, $n_{.25} = 47.0$, $n_{.75} = 134.5$, and $n_{.95} = 238.1$. The delay is assumed to be exponentially distributed, with a $\mathscr{Gamma}$ (4,6) prior density on $\tilde{\theta}$, so that the prior mean delay is $\mathscr{E}\{\tilde{\theta}^{-1}\} = 2.0$ years, with $\mathscr{V}\{\tilde{\theta}^{-1}\} = 8.0$ years^2.

For the purposes of simulation, we "stacked the deck" by using the same 100 samples (x_k, y_k) as IBNYR-I, where the x_k were drawn from a uniform distribution over (0,1), and the delays, $w_k = y_k - x_k$, were drawn from an exponential density with true parameter $\theta = 0.5$ years^{-1}. As shown in Table 1 of IBNYR-I, this gave continuous delay samples from 0.163 to 12.402 years, with a sample average delay of 2.35 years, somewhat larger than the true mean. Thus, our experiment assumes accurate but not too precise prior knowledge, so that the behavior below shows primarily the effects of quantization and the two different data types. Clearly, with vaguer prior information, we would see a

https://doi.org/10.2143/AST.20.1.2005486 Published online by Cambridge University Press

PREDICTION - TYPE Ic DATA

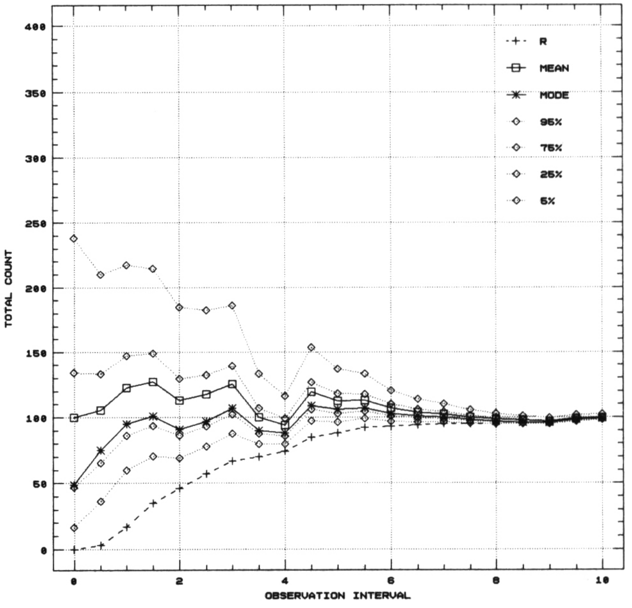

PREDICTION - TYPE IIc DATA

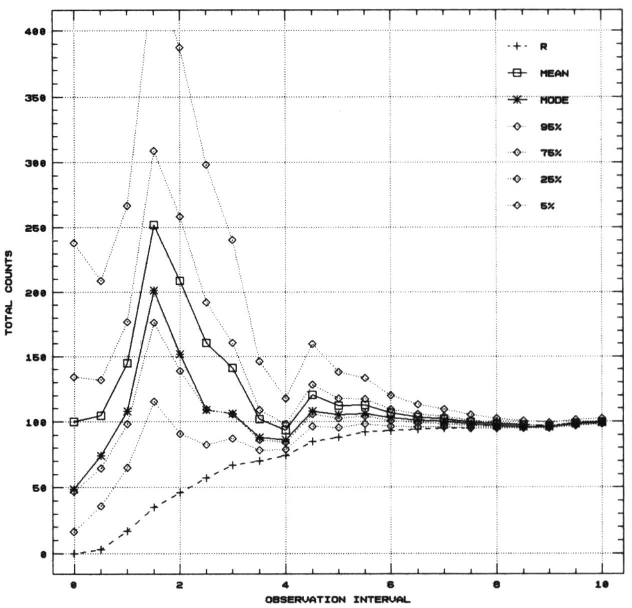

FIGURES 2a & 2b. Predictive mean, mode, and fractiles versus t for Types Ic and IIc continuous data.

https://doi.org/10.2143/AST.20.1.2005486 Published online by Cambridge University Press

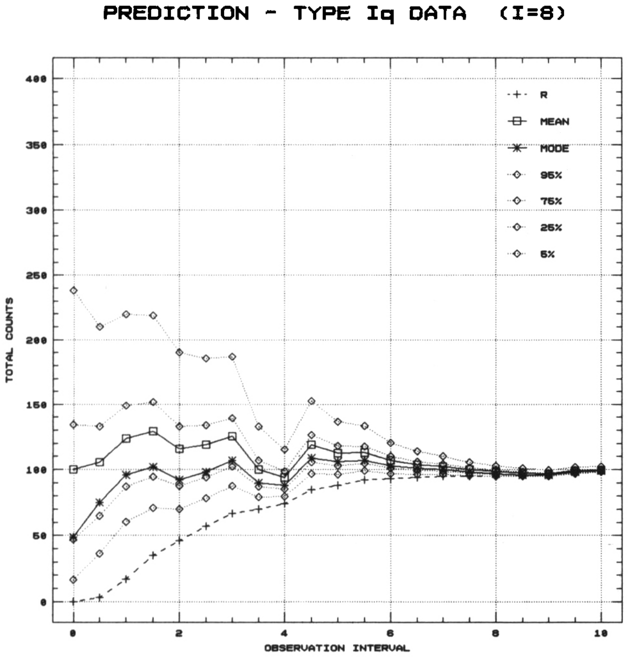

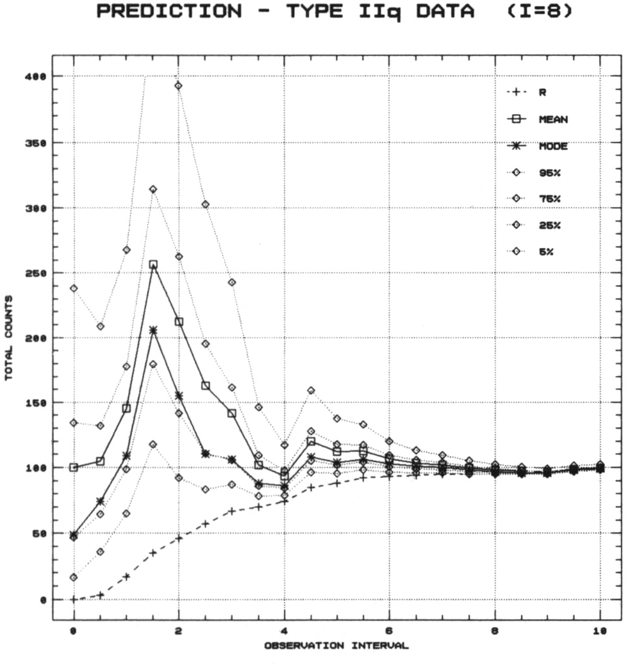

FIGURES 3a & 3b. Predictive mean, mode, and fractiles versus t for Types Iq and IIq quantized data ($I = 8$).

further degradation of the predictive power for the smaller values of $t(J)$. Figure 1 shows the individual cell counts for this sample when $\Delta = 0.25$ years ($I = 4$), and $t = 4.0$ years ($J = 16$). The values for the statistics s and r were given above in Section 2.

As the effects of quantization are the main interest of this paper, computations were carried out for many different values of I, with $I = 1, 2, 4$, and 8 finally chosen as representative, with complete predictive densities computed for observation intervals $t = 0(0.5)10.0$, except when $I = 1$, when only $t = 0(1.0)10.0$ is possible. Approximations for the shaping factor integral h_0 were computed using the Gammoid method outlined in IBNYR-I, in which a numerical search for the mode, $\hat{\theta}$, of the combination $L(\theta \mid \mathscr{D})\, p(\theta)$ is made, and the unimodal curve then approximated at the mode by a curve of the form $g(\theta) = (A\theta)^G e^{-D\theta}$. Since, to a good approximation, the kernel $K(\theta) \approx e^{-\delta\theta}$ in the neighborhood of this mode, the integral (4.3) can be computed exactly, giving a final recursive relationship like that in (10.1) of IBNYR-I. Initially, the mode was chosen from the prior density as $\hat{\theta} = 0.5$; from two to five iterations were then necessary to find the true value of the mode, which ranged from 0.46 to 1.98 in the cases examined. For smaller values of t and I, $p(u \mid \mathscr{D})$ is heavy in the tails, so, to obtain stable means, the recursion (10.1) was carried out over the range [0,1000], and, in a few cases, [0,2000]. As the no-data ($t = 0$) case is known analytically, a total of $2 \times (10 + 3 \times 20) = 140$ complete densities, $p(u \mid \mathscr{D})$, were computed for Figures 3-6 below. This task took 5-10 seconds per density on a PC-AT. The densities themselves look much like Figures 5 and 6 in IBNYR-I, and are not shown. But from these, the means, modes, and fractiles shown in the figures below were computed for the total count $n = R + u$.

Our standard of comparison will be the continuous data predictions, the results for which are reproduced from IBNYR-I in Figures 2a & 2b; for short, we shall refer to these as the Ic and IIc results, respectively. For ease in comparison, we keep the same vertical scale in all plots against the observation interval, $t(J)$.

Figures 3a & 3b show the Types Iq and IIq results for a fine quantization level, $I = 8$. At this level, it is practically impossible to see the effects of discrete reporting, as the only differences are a few percent in the upper fractiles in the interval $1.5 \leq t \leq 2.5$.

When we coarsen the quantization level to $I = 4$, as shown in Figures 4a & 4b, there begins to be a noticeable increase in the Case Iq upper fractiles and the predictive mean in the interval [1.0, 3.0], but still less than 4 % in the worst case. However, the degradation of Type IIq predictions is noticeably worse, with increase in the fractiles, the mean, and the mode in the region [0.5, 3.5], up to 11 % in the worst cases. It should be remembered that $I = 4$ means that the reporting interval is *one-eighth* the mean delay, which is already more frequent than many implementations encountered in practice.

Then, with $I = 2$, Figures 5a & 5b both show the instability in the interval [1.0, 3.5] that before was characteristic of only Type II data. In fact, the Type IIq predictions in the unstable region are now so bad as to be unreliable

PREDICTION - TYPE Iq DATA (I=4)

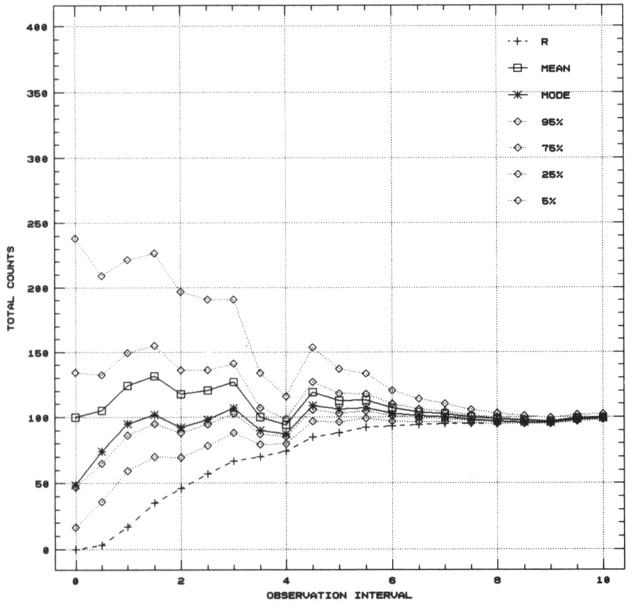

PREDICTION - TYPE IIq DATA (I=4)

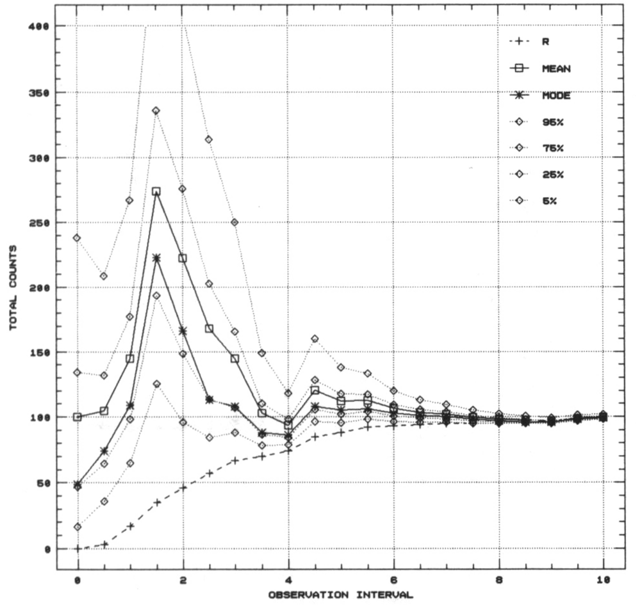

FIGURES 4a & 4b. Predictive mean, mode, and fractiles versus t for Types Iq and IIq quantized data ($I = 4$).

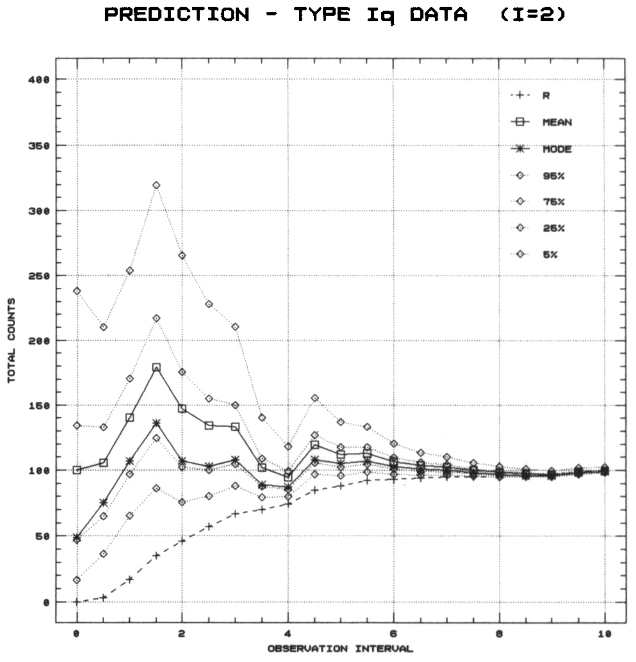

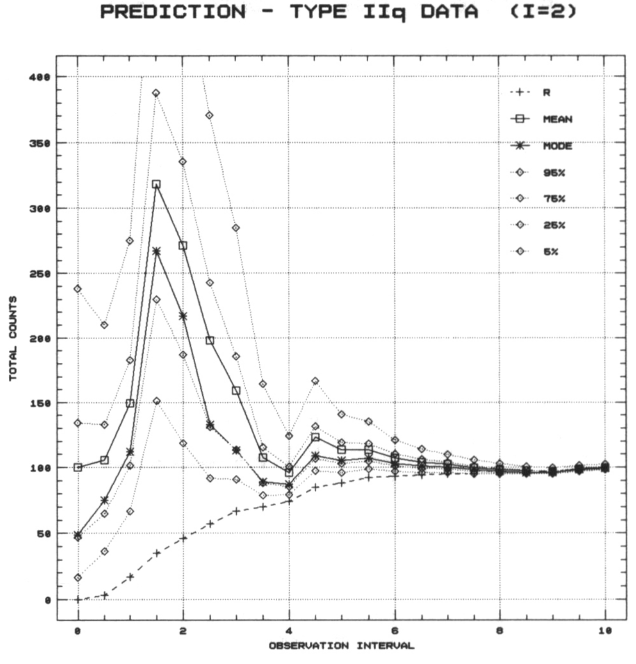

FIGURES 5a & 5b. Predictive mean, mode, and fractiles versus t for Types Iq and IIq quantized data (I = 2).

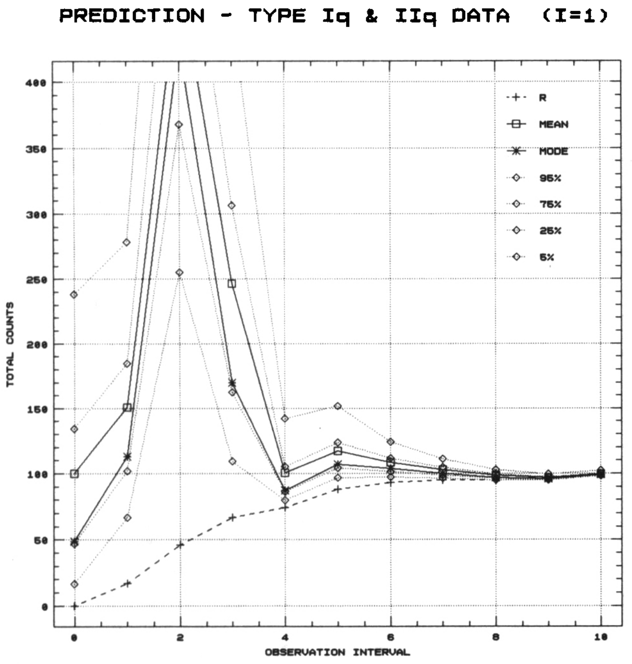

FIGURE 6. Predictive mean, mode, and fractiles versus t for Type Iq = Type IIq quantized data ($I = 1$).

unless no other estimates are available. Even the region $t \geq 4.0$, which heretofore had given similar results for both types of data because over 74% of the counts were reported, now shows some "bobbling around" due to the changing aggregation of data.

Finally, we have the case $I = 1$ in which Cases Iq and IIq coalesce. To illustrate the extreme degradation in this case, we have chosen to plot the results in Figure 6 on the *same* vertical scale as previous graphs, rather than changing the scale to show all the results. For $t = 2.0$ ($t = 1.5$ cannot be computed), the missing predictive mean count is 481.1, the mode is 430, and the upper fractiles are 575 and 763, respectively! Clearly, the use of a quantization interval that is *one-half* the mean delay is much too coarse when $1.0 \leq t \leq 6.0$. Admittedly, the region above that is reasonable, but that is prediction with at least 93% of the events already reported!

Figures 7a & 7b give a "cross-sectional" impression of the changing level of quantization, in the case for $t = 2.0$, which is in the region of instability with 46% of the events reported. The vertical scale has now been doubled, so that one may now clearly see how bad the cases $I = 1$ and $I = 2$ truly are. In my opinion, one should pick *at least* $I = 4$ in Case Iq and $I = 8$ in Case IIq to get "good" predictions, which means that, given a mean delay of 2.0 years, one must have semi-annual or quarterly data, respectively!

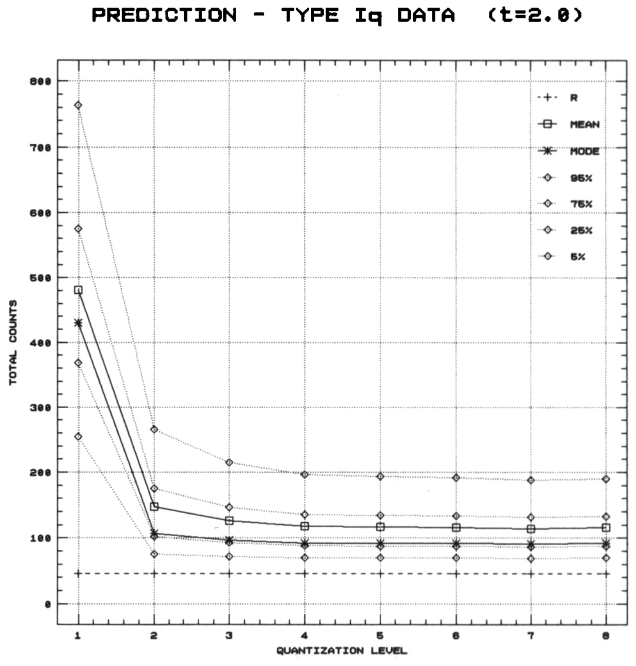

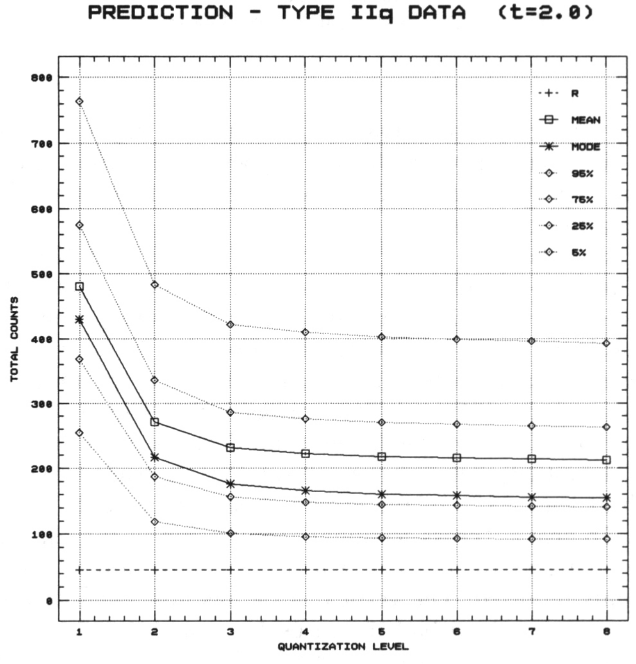

FIGURES 7a & 7b. Predictive mean, mode, and fractiles versus I for Types Iq and IIq quantized data (t = 2.0).

7. DISCUSSION AND SUMMARY

We should perhaps emphasize once more that the results obtained with changing levels of quantization (for a fixed observation interval) are due solely to changes in Δ and data type upon the epoch data likelihood. $L(\theta \mid \mathscr{D})_2$ in (5.5). This is because the part of the prediction that depends upon $\tilde{\lambda}$ is unaffected by changing Δ; R reflects *all* of the relevant information we can obtain about the event rate for the purposes of *prediction*. On the other hand, (3.3) shows that the computation of the *joint estimates* of $\tilde{\lambda}$ and $\tilde{\theta}$ will be much more difficult.

The effect of quantization upon the epoch data likelihood can be visualized in Figures 8a & 8b, which show this function when $t = 2.0$ for $I = \infty$ (continuous data), 4, 2 and 1, for the two different data types. Although the mean and mode shift somewhat as I decreases from ∞ towards 2, the predominant effect is an increased spread in the likelihood. These likelihoods are multiplied by the prior density (dotted line), the results approximated by a Gammoid, and then used with the kernel to find the shaping factors $h_0(u \mid \mathscr{D})$, and, from the recursion (4.4), the final predictive density. Note that Type IIq data likelihoods, although converging faster with finer quantization, do not shift the mode as much as Type Iq; since the true value of θ is 0.5 (mode of prior density), this means that Type IIq data will give less accurate predictions. The case $I = 1$ is, well, hopeless.

Keeping in mind the summary observations that were already made in IBNYR-I about the continuous-data prediction problem, the main lessons to be drawn from this paper are:

(1) The introduction of quantized reporting of epochs into the IBNYR model requires no new concepts and only a modest increase in algebra and computational effort.

(2) Case IIq data (no occurrence dates reported) continue to give poorer predictions than Case Iq (both occurrence and reporting epochs known) and the predictions degrade more quickly with coarser quantization.

(3) The predictive accuracy of these discrete-time models, in comparison with the continuous case, declines dramatically as Δ increases from, say, one-sixteenth the mean delay to one-quarter the mean delay. A tentative rule-of-thumb seems to be to choose Δ to be at least one-eighth the mean delay with Iq data and one-sixteenth the mean delay with IIq data, if at all possible.

(4) The case $I = 1$ (Δ is one-half the mean delay), while coalescing the two data types and simplifying the sufficient statistics, is so poor as to be unusable in the region of interest.

Admittedly, it is dangerous to extrapolate from one numerical example to practice. For instance, one may be able to be much more precise *a priori* about the parameters of the delay distribution; this narrower prior will, to some extent, counteract the imprecise data likelihoods obtained with coarse quantization. And, as always, the final predictive spreads can be greatly reduced if we

EPOCH LIKELIHOODS & PRIOR - TYPE I DATA

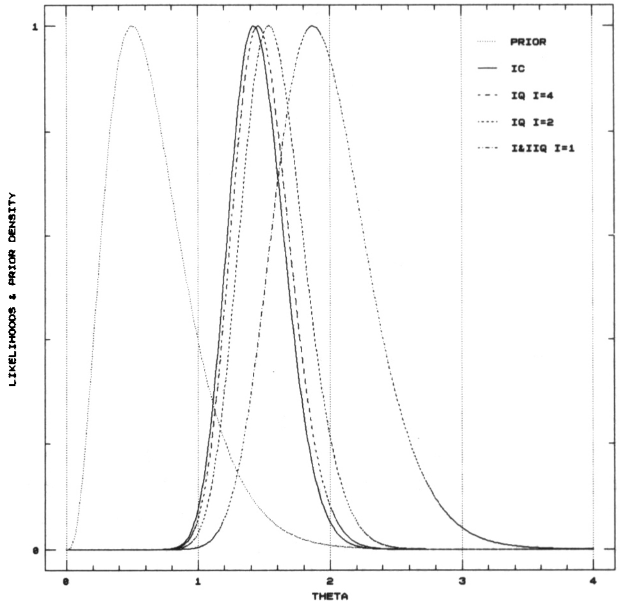

EPOCH LIKELIHOODS & PRIOR - TYPE II DATA

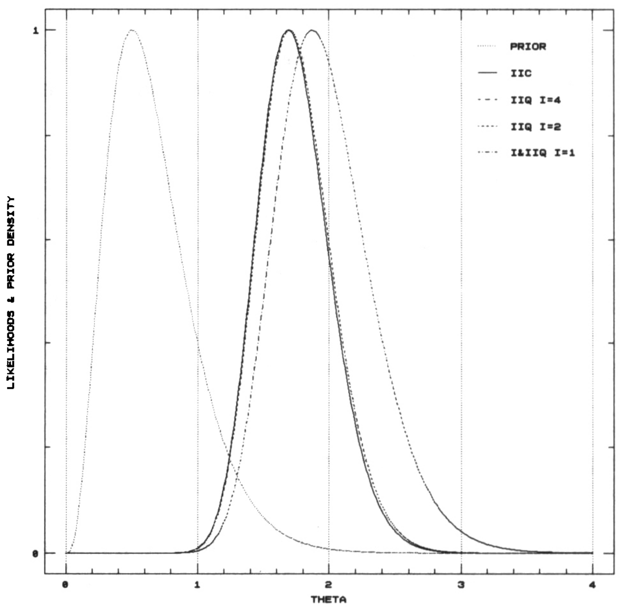

FIGURES 8a & 8b. Epoch Data Likelihood, $L(\theta|\mathscr{D})$, and Prior Density, $p(\theta)$, versus θ for Types Iq and IIq quantized data (t = 2.0).

can provide better prior information about the occurrence rate, perhaps by incorporating the underlying business volume into the model.

With this understanding of the potential hazards of quantized reporting, our next paper will consider the question of whether or not *cohort data* from an IBNR traingle can sharpen our estimation of the unknown delay distribution and improve our predictions of the unreported events.

I would like to thank M. LIN for her substantial computational and proofing assistance in developing these results. Any comments or criticisms on this paper are welcome, as are suggestions for making the basic model more realistic and useful.

ACKNOWLEDGMENT

This research has been partially supported by the Air Force Office of Scientific Research (AFOSR), USAF, under Grant AFOSR-81-0122 with the University of California. Reproduction in whole or in part is permitted for any purpose of the United States Government.

PARTIAL REFERENCES

JEWELL, W. S. (1989) Predicting IBNYR events and delays. *ASTIN Bulletin* **19**, 25–55.

KASS, R. E., TIERNEY, L. and KADANE, J. B. (1988) Asymptotics in Bayesian computation. In *Bayesian Statistics 3*, J. M. Bernardo, M. H. DeGroot and A. F. M. Smith (Eds.), Oxford University Press, 261–278.

KAHANE, Y. (1989) A modern approach to loss reserving in long-tail lines — the case of automobile insurance. *Transactions, XXIth ASTIN Colloquium, New York*, 175–197.

Reinsurance Association of America (1987) Loss development study. *Transactions, XXIth ASTIN Colloquium, New York*, 525–583.

SCHNIEPER, R. (1989) A pragmatic IBNR method. *Transactions, XXIth ASTIN Colloquium, New York*, 405–425.

STIERS, D. (1989) Applications de méthodes d'évaluation des réserves. *Transactions, XXIth ASTIN Colloquium, New York*, 441–463.

TIERNEY, L. and KADANE, J. B. (1986) Accurate approximations for posterior moments and marginal densities. *J. Amer. Statist. Assoc.* **81**, 82–86.

WELLER, A. O. (1989) Generalized Bondy development. *Transactions, XXIth ASTIN Colloquium, New York*, 585–609.

WILLIAM S. JEWELL

Department of Industrial Engineering & Operations Research, 4173 Etcheverry Hall, University of California, Berkeley, California 94720 USA.