Bayesian Reserving Models Inspired by Chain Ladder Methods and Implemented Using WinBUGS *

David P.M. Scollnik †

Abstract

This paper examines some new Bayesian models for loss reserving inspired by a consideration of some of the methods and techniques appearing in the traditional chain ladder literature. This includes a possibly order restricted hierarchical Bayesian model for the year over year development factors. The issues of reserve variability and ranges of reasonable reserve estimates are considered. A new Bornhuetter-Ferguson styled Bayesian method of estimate revision is also discussed. These Bayesian models can be implemented using Markov chain Monte Carlo (MCMC) methods in a variety of programming languages, and can serve as starting points for more advanced models. This paper describes their implementation in WinBUGS (a specialized software program for MCMC simulations). Illustrative WinBUGS code is included.

*This work was supported by a grant from the Actuarial Education and Research Fund (AERF).

†David P.M. Scollnik, A.S.A., Ph.D., is a Professor in the Department of Mathematics and Statistics, University of Calgary, 2500 University Drive NW, Calgary, Alberta, Canada, T2N 1N4, phone: 403-220-7677, e-mail: scollnik@ucalgary.ca.

1 Introduction

Chain ladder reserving methods have been discussed in the actuarial literature for many years. Taylor (1986) dates the original methodology back to Harnek (1966). Taylor (2000, page 26) describes chain ladder models for reserving as those that chain a sequence of ratios together (i.e., age to age factors) 'into a ladder of factors ... which enable one to climb (i.e. project) from experience recorded to date to its predicted ultimate value.' Taylor (2000, Chapter 3) provides a detailed overview of traditional chain ladder reserving methods, as does Booth, Chadburn, Cooper, Haberman, and James (1999, Chapter 16). An example using a traditional chain ladder method is given below (see Section 2).

In recent years, a number of authors have considered Bayesian methods for claims reserving. While not an exhaustive list, notable contributions on the topic are made by Verrall (1989, 1990), Ntzoufras and Dellaportas (2002), de Alba (2002), England and Verrall (2002), and Lamps (2002a, 2002b, 2002c). Verrall (1989, 1990) describes some ways in which the traditional chain ladder method can be analyzed using the theory of Bayesian linear models. These are inspired by work of Kremer (1982), who shows how to transform the multiplicative chain ladder model into the form of the linear model associated with the two-way analysis of variance by taking logarithms. Assuming positive and lognormally distributed incremental claim amounts, Verrall (1990) demonstrates how this model can be analyzed using either an empirical Bayesian or a parametric fully Bayesian approach in the context of hierarchical linear modeling. Verrall (1989) emphasizes a Kalman filter (state space) approach in the Bayesian analysis.

Ntzoufras and Dellaportas (2002) consider the Bayesian analysis of four models for claim amounts using Markov chain Monte Carlo (MCMC) methods. Their models begin with an assumption of lognormally distributed incremental claim amounts in an analysis of variance setting. They go on to explore the use of state space modeling and the inclusion of additional information in the form of inflation factors, measures of exposure (such as portfolio size), and associated claim counts. In his discussion of Ntzoufras and Dellaportas (2002), Scollnik (2002a) shows how the models of Ntzoufras and Dellaportas can be implemented using

WinBUGS (specialized software for MCMC discussed later in this paper).

The paper by de Alba (2002) considers several model specifications that are essentially the same as two of those appearing in Ntzoufras and Dellaportas (2002). Note that de Alba (2002, page 8) states that Ntzoufras and Dellaportas (2002) require the number of claims per year of origin be known and fixed in order for their MCMC procedure to be implemented. In fact, Ntzoufras and Dellaportas (2002, page 128) permit this value to be a stochastic variable and describe how to modify the MCMC procedure in this case. It turns out that the principle difference between de Alba (2002) and Ntzoufras and Dellaportas (2002) is in the method of implementation. Whereas Ntzoufras and Dellaportas (2002) emphasize the use of a MCMC simulation method, de Alba (2002) describes how direct Monte Carlo can be used instead. The former approach may be more accessible to the average practitioner, especially if he or she wishes to introduce changes to the basic models. The paper by de Alba (2002) is valuable for its summary of Bayesian reserving methods that have appeared to date.

England and Verrall (2002a, 2002b) gift us with a detailed treatise on the state of current claim reserving methodologies. The methodologies they describe are primarily non-Bayesian, but they also discuss the Bayesian analysis of an over-dispersed Poisson chain ladder model (England and Verrall, 2002a, pages 19-20 and 49-52). These authors also describe a variety of generalized linear models appropriate for loss reserving; in principle, any of these models could be implemented using a Bayesian method. Readers interested in pursuing this approach may find the text "Generalized Linear Models: A Bayesian Perspective" of interest (Dey, Ghosh, and Mallick, 2000).

Numerous authors, including Cairns (2000), Rosenberg and Young (1999), Scollnik (2001, 2002b), Ntzoufras and Dellaportas (2002), and de Alba (2002), along with many others referenced in these cited works, have described the advantages of Bayesian over traditional non-Bayesian methods of analysis and given reasons for the popularity of Bayesian models in recent years. Included among the commonly stated advantages of the Bayesian method of analysis are these: it requires that the statistical model be completely specified; it allows one to incorporate relevant and significant prior information in the analysis when it is available; it allows parameter and model uncertainty to be explicitly modeled and accounted for in the analysis and in the resulting predictive forecasts; it yields complete posterior distributions for quantities of interest, and posterior predictive distributions for unobserved future observations, not just point estimates or confidence intervals. In addition, Rosenberg and Young (1999, page 132) note that the output of their particular 'Bayesian method is easier for actuaries to understand than output from non-Bayesian methods because actuaries are used to thinking in terms of probability.' This is probably true for most Bayesian methods applied in actuarial contexts, including loss reserving. Readers unfamiliar with Bayesian methods, or wishing an accessible review of modern Bayesian statistical modeling, will find the books by Gelman, Carlin, Stern, and Rubin (2004) and Congdon (2001, 2003) of interest.

No doubt the development of MCMC simulation algorithms well-suited for Bayesian analyses, together with the ubiquity of fast personal computers, have played a significant role in popularizing Bayesian methods. Specialized software programs facilitating the application of these MCMC algorithms to statistical models, notably WinBUGS, have for the first time brought modern Bayesian data analysis into the hands of the average practitioner. The WinBUGS program, documentation, and related resources are freely available from The BUGS Project website found at www.mrc-bsu.cam.ac.uk/bugs. Congdon (2001, 2003) emphasizes the use of WinBUGS in his Bayesian texts. Scollnik (2001, 2002a, 2002b, 2002c) discusses the operation of WinBUGS and illustrates its use in a variety of actuarial contexts.

The purpose of this paper is to examine some new Bayesian models for loss reserving, inspired by a consideration of some of the traditional methods and techniques appearing in the chain ladder literature. This paper will also direct the reader as to how these Bayesian models can be implemented using WinBUGS. However, note that these Bayesian models can be implemented using MCMC methods in a variety of programming languages, and can serve as starting points for more advanced models. The reader should understand that a feature of modern Bayesian statistical modeling with MCMC is its flexibility. Consequently, it would be incorrect and misleading for this paper to suggest that a particular model must be used exactly as it is presented herein. The room exists for many subtle variations, as the case warrants or as the practitioner desires. For this reason, although the paper will mention a number of possible model variations that a practitioner may wish to explore, it will not always go into explicit detail when it does so. We will assume that readers have at least a basic understanding of the ideas presented in Scollnik (2001), to save us from repeating summary reviews of Bayesian statistical modeling, MCMC, and WinBUGS that have already appeared many times before in the actuarial (and statistical) literature.

In the next section, we will introduce a data set and develop its analysis using traditional chain ladder methods. For instructional reasons, we will discuss how this basic analysis can be implemented in WinBUGS. Subsequently, a variety of Bayesian models will be introduced and applied to the same data set. The implementation of these Bayesian models in WinBUGS will be an integral part of the discussion. Section 6 concludes and describes how code for these illustrative models may be obtained.

2 A Data Set and Traditional Chain Ladder Analysis

Let $Y[\,i,j\,]$ denote the claim amounts paid by the insurance company with a delay of $j-1$ years for accidents reported in the year i, with $i,\, j = 1,\ldots,r$. The value of j is commonly known as the development period. Let $Z[\,i,j\,]$ denote the cumulative claim amount for accidents reported in the year i with a delay of $j-1$ years or less. For convenience, we will assume that the observed data is in the traditional upper triangular form such that $Y[\,i,j\,]$ and $Z[\,i,j\,]$ are observed for $i=1,\ldots,r$ and $j=1,\ldots,r-i+1$, and unobserved elsewhere. We assume that the observed data has been adjusted for inflation. The incremental and cumulative claims for just such a data set are given in Tables 1 and 2, respectively. This data set appears in Chapter 16 (see Table 16.26) of Booth *et al.* (1999).

Define the single cell development factor devfac$[\,i,j\,]$ as

$$\text{devfac}[\,i,j\,] = \frac{Z[\,i,\,j+1\,]}{Z[\,i,j\,]}\,, \tag{1}$$

for $i = 1,\ldots,r$ and $j = 1,\ldots,r-1$. At the end of reporting year r, these factors are only observed for $i = 1, \ldots, r-1$, with $j = 1, \ldots, r-i$. For our example, these factors are reported in Table 3 (with $r = 6$). A traditional chain ladder reserving method uses the observed data in order to estimate the missing single cell development factors in the lower triangle. Then these estimated development factors are used, in conjunction with (1), to develop estimates of the cumulative claim amounts in the lower triangle and, hence, of the missing incremental claim amounts and the loss reserve.

There are many possible ways in which to construct estimates of the missing single cell development factors in each column. Just to name a few possibilities, a practitioner might use the arithmetic (or a weighted) mean of the observed factors in each column, the most recent factor appearing in a column, or the average of some number (e.g., two or three) of the most recent factors appearing in a column in order to complete each column's missing entries (see Brown and Gottlieb, 2001, pages 123-124). The popular set of estimates known as the volume weighted development factors are given by

$$\text{volwtdevfac}[\,j\,] = \frac{\sum_{i=1}^{r-j} Z[\,i,\, j+1\,]}{\sum_{i=1}^{r-j} Z[\,i, j\,]}\,, \tag{2}$$

for $j = 1, \ldots, r-1$. Observe that the volume weighted development factors are weighted averages of the single cell development factors, with the cumulative claim amounts appearing in the denominator of the latter used as the weights involved in the calculation of the former.

Appearing in Table 3 are the volume weighted development factors for the data appearing in Table 2. These factors were used to develop the estimates of the future cumulative and incremental claim amounts reported in Tables 4 and 5, respectively. The example in Booth *et al.* (1999) assumes future inflation of 12% per annum, and a 6% effective rate of interest for discounting the inflated claim amounts back to the present. Applying this inflation rate to the current money incremental claims in Table 5 produces the nominal incremental claim amounts in Table 6. Discounting these back to the present using the assumed 6% effective rate of interest produces the net present value incremental claims in Table 7. The current money, nominal, and net present value reserves in Table 8 are obtained by summing the rows in Tables 5, 6, and 7, respectively.

Nothing presented in the discussion above is new. However, it has provided us with a review of the basic chain ladder method along with notation and an example that we will refer to throughout the remainder of the paper.

3 Implementing the Chain Ladder Method in WinBUGS

WinBUGS can be used to develop point estimates of the outstanding claims via the traditional chain ladder route making use of the volume weighted development factors (or some other function of the observed single cell development factors) in the manner described above. This is not a particularly exciting or useful application of WinBUGS as it merely reproduces results that are as easy to develop using any spreadsheet package. However, the WinBUGS code used in the process is instructive and will be used again in subsequent sections.

Assume that the observed incremental claim amounts $Y[\,i,j\,]$ are read in as data. Then the following WinBUGS code calculates the cumulative claim amounts defined up to and including the present date, along with the volume weighted development factors. Minuscule normal errors are added to the constant volume weighted development factors in order so that they can be monitored as stochastic nodes in WinBUGS. Note that the second parameter in the `dnorm` definition of a normal density is the precision, or inverse variance, parameter.

```
# Define the first column of the cumulative claims triangle.
for( i in 1 : r ) {
   z[ i, 1 ] <- y[ i, 1 ]
}

# Compute the cumulative claim amounts in successive development periods
#  up to and including the present date.
for( i in 1 : ( r - 1 ) ) {
   for( j in 1 : ( r - i ) ) {
      z[ i, j + 1 ] <- z[ i, j ] + y[ i, j + 1 ]
   }
}
```

```
# Calculate the traditional volume weighted development factors.
for( j in 1 : ( r - 1 ) ) {
   volwtdevfac[ j ] <- sum( z[ 1 : ( r - j ), j + 1 ] ) /
                       sum( z[ 1: ( r - j ), j ] ) + normerror[ j ]
   normerror[ j ]  ~ dnorm( 0, 1000000000 )
}
```

Now, let $Y.ns[i, j]$ and $Z.ns[i, j]$ denote the non-stochastic chain ladder estimates for the future incremental and cumulative claim amounts, developed on the basis of the observed volume weighted development factors. The WinBUGS code below uses the volume weighted development factors to project the claim amounts forward into the future development periods, thus yielding $Y.ns[i, j]$ and $Z.ns[i, j]$.

```
# Traditional model for incremental and cumulative claims.
for( i in 2 : r ) {
   z.ns[ i, r - i + 1 ] <- z[ i, r - i + 1 ]
   for( j in ( r - i + 1 ) : ( r - 1 ) ) {
      z.ns[ i, j + 1 ] <- z.ns[ i, j ] * volwtdevfac[ j ]
      y.ns[ i, j + 1 ] <- z.ns[ i, j + 1 ] - z.ns[ i, j ]
   }
}
```

As mentioned earlier, the future incremental claim amounts can be adjusted for inflation and present value discounting. Let the vectors lambda.inflate and lambda.discount contain the anticipated inflation and appropriate discounting rates for future years. (Recall that in the continuing example, the inflation and discounting rates are constant at 12% and 6% effective per annum, respectively.) The following WinBUGS code uses these rate vectors to develop the inflation and net present value adjusted incremental claims. It assumes that claim payments are at the end of each year but the code is easily modified to fit other timing assumptions.

```
# Adjust claim amounts for inflation and present value discounting.
for( i in 2 : r ) {
   for( j in ( r - i + 2 ) : r ) {
      y.ns.inflate[ i, j ] <- y.ns[ i, j ] *
         prod( lambda.inflate[ 1 : ( i + j - r - 1 ) ] )
      y.ns.discount[ i, j ] <-  y.ns.inflate[ i, j ] /
         prod( lambda.discount[ 1 : ( i + j - r - 1 ) ] )

   }
}
```

At this stage, the current money, nominal, and net present value reserves for each year of origination can be obtained by simply summing the constituent future incremental claim amounts over the appropriate cells. These calculations are coded below. The outstand.ns.row variable label reminds us that the outstanding claim amounts, or reserves, are being calculated by row. (The reserves can also be obtained by calendar year, in which case the sums would be defined over cells with fixed values of $i+j$. See Scollnik (2002a, pages 134-135) for details.)

```
for( i in 2 : r ) {
   outstand.ns.row[ i ] <- sum( y.ns[ i, ( r + 2 - i ) : r ] )
   outstand.ns.row.nom[ i ] <- sum( y.ns.inflate[ i, ( r + 2 - i ) : r ] )
   outstand.ns.row.npv[ i ] <- sum( y.ns.discount[ i, ( r + 2 - i ) : r ] )
}

outstand.ns.row[ r + 1 ] <- sum( outstand.ns.row[ 2 : r ] )
outstand.ns.row.nom[ r + 1 ] <- sum( outstand.ns.row.nom[ 2 : r ] )
outstand.ns.row.npv[ r + 1 ] <- sum( outstand.ns.row.npv[ 2 : r ] )
```

4 A Bayesian Chain Ladder Model

Our first Bayesian model builds on the presentation above and is based on the observation that the single cell development factors tend to be similarly valued, given the development year j. This is clearly illustrated by the continuing example in Table 3. Thus, our first Bayesian analysis begins by specifying stochastic models with equal means for single cell development factors sharing a common development year. We will assume normal models, although others could be entertained (e.g., gamma or lognormal). Thus,

$$\text{devfac}[\,i, j\,] \sim \text{normal}\,(\,\theta_j\,,\,\tau_{i,j}\,)\,, \tag{3}$$

for $i = 1, \ldots, r$ and $j = 1, \ldots, r-1$. Observe that negative incremental claims are permitted by this model, since devfac$[\,i, j\,]$ may be less than 1. The second parameter appearing in this normal distribution is a precision, or inverse variance, parameter (in accordance with the definition of the normal distribution in WinBUGS). The precision parameters $\tau_{i,j}$ may be modeled in a variety of ways. For example, they may be set equal to a common value, say τ, for all values of i and j, or they may be set equal to a common value, say τ_j, for all values of i given a particular j. Possibly, the parameters $\tau_{i,j}$ may be scaled by relevant weights. In this case, we could define $\tau_{i,j} = \tau * \mathsf{weight}[\,i, j\,]$. The weights could be related to the written premium associated with the different years of origination, or perhaps to the number of claims associated with the different years of origination that settled in the latest development period (for example). Ideally, the actuarial practitioner should make use of his or her experience in order to construct a reasonable set of assumptions for model parameters like $\tau_{i,j}$, given the line of business under consideration. If a practitioner must make a choice between several possible model specifications, then the method of posterior predictive checks (discussed later in this section) can be used to select the one that is most appropriate in light of the data.

If desired, bounds on the unobserved development factors could be imposed by applying interval restrictions to the relevant normal distributions. For example, imagine imposing restrictions, based on extensive past experience, such that future cumulative claims cannot decline by more than 5% year over year. This may be accomplished in WinBUGS by including a I(0.95,) term with each normal distribution definition. This is illustrated in the WinBUGS code provided below. Note that interval restrictions coded with the I(,) construct in WinBUGS only apply to model parameters with unobserved values; that is, they are automatically ignored for any model parameter that is observed. See the WinBUGS User Manual and related documentation for more details on the use of I(,). Observed model parameters (including observed development factors) can still be modeled with interval restricted distributions using the 'zeros trick' or 'ones trick' described in the WinBUGS User Manual. See Scollnik (2001, 2002b, 2000c) for examples of this.

Next, we suppose that the underlying θ_j parameters are drawn from a common normal distribution,

$$\theta_j \sim \text{normal}\,(\,\mu_\theta\,,\,\tau_\theta\,)\,, \tag{4}$$

for $j = 1, \ldots, r-1$. Bounds or order restrictions may be placed on the θ_j parameters (e.g., so that they form a decreasing sequence) by using the I(,) construct in WinBUGS as discussed above. Order restrictions are probably superfluous when the development factors are as clearly differentiated between different development years as they are in Table 3.

The model specification described above can be coded in WinBUGS as below (i.e., by extending the code previously given so as to include the new definitions).

```
# Define the development factors for the observed data.
for( i in 1 : ( r - 1 ) ) {
    for( j in 1 : ( r - i ) ) {
      devfac[ i, j ] <- z[ i, j + 1 ] / z[ i, j ]
   }
}

# Define the stochastic model for future cumulative claims.
for( i in 2 : r ) {
   for( j in ( r - i + 1 ) : ( r - 1 ) ) {
      z[ i, j + 1 ] <- z[ i, j ] * devfac[ i, j ]
```

```
      y[ i, j + 1 ] <- z[ i, j + 1 ] - z[ i, j ]
   }
}

# Define a distribution on the development factors devfac.
for( i in 1 : r  ) {
   for( j in 1 : ( r - 1 ) ) {
      # Uncomment / modify the I( 0.95, ) terms as desired.
      devfac[ i, j ] ~ dnorm( mu.df[ i, j ], tau.df[ i, j ] ) # I( 0.95, )
      mu.df[ i, j ] <- theta[ j ]
      tau.df[ i, j ] <- tau * weight[ i, j ]
      weight[ i, j ] <- 1 # Or whatever is appropriate.
   }
}

# Code for unrestricted theta[ j ] nodes.
for( j in 1 : ( r - 1 ) ) {
   theta[ j ] ~ dnorm( mu.theta, tau.theta )
}

# Illustrative code for order restricted (decreasing) theta[ j ] nodes.
# theta[ 1 ] ~ dnorm( mu.theta, tau.theta ) I( theta[ 2 ], )
# for( j in 2 : ( r - 2 ) ) {
#    theta[ j ] ~
#       dnorm( mu.theta, tau.theta ) I( theta[ j + 1 ], theta[ j - 1 ] )
# }
# theta[ r - 1 ] ~ dnorm( mu.theta, tau.theta ) I( 0, theta[ r - 2 ] )
```

The remaining parameters τ, μ_θ, and τ_θ must be assigned prior density specifications in order to complete the definition of a full (i.e., fully specified) probability model. These distributions may be diffuse proper priors or very informative ones, as the problem at hand may warrant. In the first analysis of the continuing example to follow, we will assume that

$$\tau \sim \text{gamma}\,(\,0.001\,,\,0.001\,)\,,$$

$$\mu_\theta \sim \text{normal}\,(\,1\,,\,0.001\,)\,,$$

$$\tau_\theta \sim \text{gamma}\,(\,0.001\,,\,0.001\,)\,.$$

These distributions all have means of 1 and variances of 1,000. They are meant to be illustrative, and are by no means necessarily appropriate for default application. Other prior density specifications will be presented in short order. Illustrative WinBUGS code is presented below.

```
tau ~ dgamma( 0.001, 0.001 )
mu.theta ~ dnorm( 1, 0.001 )
tau.theta ~ dgamma( 0.001, 0.001 )
```

As in the previous section, we can adjust the incremental claim amounts for inflation and present value discounting. The current money, nominal, and net present value reserves for each year of origination can be obtained by simply summing the constituent future incremental claim amounts over the appropriate cells. These calculations are coded below.

```
# Adjust claim amounts for inflation and present value discounting.
for( i in 2 : r ) {
   for( j in ( r - i + 2 ) : r ) {
      y.inflate[ i, j ] <- y[ i, j ] *
         prod( lambda.inflate[ 1 : ( i + j - r - 1 ) ] )
      y.discount[ i, j ] <-  y.inflate[ i, j ] /
         prod( lambda.discount[ 1 : ( i + j - r - 1 ) ] )

   }
}

for( i in 2 : r ) {
   outstand.row[ i ] <- sum( y[ i, ( r + 2 - i ) : r ] )
   outstand.row.nom[ i ] <- sum( y.inflate[ i, ( r + 2 - i ) : r ] )
```

```
   outstand.row.npv[ i ] <- sum( y.discount[ i, ( r + 2 - i ) : r ] )
}

outstand.row[ r + 1 ] <- sum( outstand.row[ 2 : r ] )
outstand.row.nom[ r + 1 ] <- sum( outstand.row.nom[ 2 : r ] )
outstand.row.npv[ r + 1 ] <- sum( outstand.row.npv[ 2 : r ] )
```

The Bayesian analysis of this model yields the posterior distribution (i.e., not just point estimates) for all unknown model parameters. This includes the posterior predictive distribution of the unobserved claim and cumulative claim amounts (i.e., the reserves). Using WinBUGS, it is a simple matter to estimate these posterior distributions, and their corresponding means, standard deviations, desired percentiles, or any other summary statistics of interest. Graphical summaries, such as histograms and estimated posterior density plots, may also be easily generated in WinBUGS. See the WinBUGS User Manual or Scollnik (2001) for details.

As previously noted, Scollnik (2001) provides a complete discussion addressing the specifics of model checking and data loading, compiling and initializing, running, monitoring for convergence and summarizing the results from a simulation coded in WinBUGS. These steps were implemented for the present model. Four independent simulations were performed and each was allowed to burn-in for 50,000 iterations. Each simulation was then run for an additional 50,000 iterations, over which the values of the unknown variables of interest (i.e., the missing data and some unknown model parameters) were monitored and then summarized. Thus, our posterior summaries are based on $4 * 50{,}000 = 200{,}000$ simulated values for each unknown variable of interest.

Summary statistics for the predictive distribution of unobserved claim and cumulative claim amounts are given in Tables 9, 10, and 11. Table 9 lists the predictive means, standard deviations, and 95% probability intervals for the future incremental claim amounts in successive periods. These summary statistics are in current money; summary statistics accounting for inflation and present value discounting are also available from the simulation run in WinBUGS, but are not reported due to space considerations. Table 10 lists the predictive means, standard deviations, and 95% probability intervals for the reserves associated with the different years of origin. This table does include results for when the future claim amounts are adjusted for inflation and present value discounting. As before, the inflation and discounting rates are held constant at 12% and 6% effective per annum, respectively. The predictive means reported in Tables 9 and 10 are generally close to the corresponding traditional chain ladder estimates given in Tables 5 and 8. Generally, the predictive means slightly exceed the traditional chain ladder estimates. However, each chain ladder estimate is contained within the corresponding 95% posterior probability interval produced by the Bayesian analysis.

From Tables 9 and 10, it is evident that the predictive distributions for the claim amounts in future years and for the reserves are rather dispersed, with large standard deviations and wide 95% posterior probability intervals. It is fair to ask how the variability in these results compares with the variability inherent in standard non-Bayesian reserving methods. Booth *et al.* (1999, Chapter 16) determined estimates of the total (current money) reserve and variability thereof for the data set under consideration using the bootstrap method (Efron and Tibshirani, 1993) and also a distribution free approach set forth in Mack (1993). See Booth *et al.* (1999) for details of these calculations. The bootstrap method gave an estimate of the expected total reserve equal to 153,077 with a sample standard deviation of 9,922. Mack's distribution free approach gave an estimate of the expected total reserve equal to 151,600 with the standard error of this prediction equal to 7004.5. The important point to note is that these estimates of variability are with regard to the estimate of the mean (i.e., expected) total reserve, and not to the 'actual' value of the total reserve as given in Tables 9 and 10. Thus, it is incorrect to compare the measures of variability in Tables 9 and 10 for the actual total reserve to the standard errors for the bootstrap and distribution free estimates of the mean reserve. It is more sensible to compare the standard errors associated with the traditional estimates to the Monte Carlo standard error estimates of the predictive mean values. These Monte Carlo standard errors are also produced in WinBUGS and are reported in Table 11 for the reserves (under the 'Base Model' column heading). On this basis, the Bayesian predictive estimates of the mean reserves in this example are actually somewhat less variable than the estimates of the mean reserves developed using the two non-Bayesian methods. Nonetheless, the results in Tables 9 and 10 serve to remind us just how variable the predictive results for the actual (rather than the mean) values of the reserves really are.

A reviewer of this paper suggested that the two measures of variability presented in the tables above, i.e., the standard deviation for the predicted actual reserves in Tables 9 and 10 and the Monte Carlo standard error for the predictive mean values in Table 11, might be interpreted in a context of "acceptable ranges" or "reasonableness" for casualty loss reserves. That is, if another actuary using a different reserving procedure arrives at a reserve estimate that is within a narrow interval centered about one of the predictive mean values given in Table 11, say within two Monte Carlo standard errors of the given predictive mean value, then that reserve estimate should also be considered a reasonable value. If the actual results subsequently fall in a wider range but still within, say, two standard deviations of the predicted actual reserve or alternatively within the central 95 % predictive probability interval for the actual reserve as given in Table 10, then the model and the original point estimate cannot be faulted. Of course, other interval widths could be used to define acceptable ranges. This is easy to implement, as the output of the Bayesian method we've described can easily be used to estimate percentiles of the predictive distribution for the actual reserve. So, for example, we could use the 30 th and 70 th percentiles of the predictive distribution to define an acceptable range. The issues of reserve variability and ranges of reasonable reserve estimates were recently discussed in papers by Hayne (2003), McClenahan (2003), and Shapland (2003).

Table 11 also contains results for two variations on the base model described above. Both correspond to simple variations on the prior density specifications and are provided for illustration. The first model variation assumes that

$$\tau \sim \text{gamma}\,(\,1.5625\,,\,0.0025\,)\,,$$

$$\mu_\theta \sim \text{normal}\,(\,1\,,\,1\,)\,,$$

$$\tau_\theta \sim \text{gamma}\,(\,1.5625\,,\,0.0025\,)\,.$$

This implies that $1/\sqrt{\tau}$ and $1/\sqrt{\tau_\theta}$ both have a prior mean and standard deviation of (approximately) 0.054 and 0.038, respectively. Recall that τ and τ_θ are the precision parameters (inverse variances) associated with the conditional prior distributions (3) and (4) for devfac$[\,i,j\,]$ and θ_j, for $i = 1,\ldots,r$ and $j = 1,\ldots,r-1$. Hence, $1/\sqrt{\tau}$ and $1/\sqrt{\tau_\theta}$ are the corresponding conditional prior standard deviations for the parameters devfac$[\,i,j\,]$ and θ_j, respectively. Also, recall that the parameter μ_θ is the prior conditional mean for θ_j, for $j = 1,\ldots,r-1$. The prior density specification given above assigns μ_θ a conditional prior mean and standard deviation that are both equal to 1. The prior density specification inherent in this model variation is quite different from the one we started with, in which the parameters τ, μ_θ, and τ_θ all shared common means and variances equal to 1 and 1,000, respectively. The second model variation is the same as the first variation described above, except that the unobserved devfac$[\,i,j\,]$ parameters are now restricted to be no less than 0.95. Both model variations were easily implemented in WinBUGS, and summary output from the corresponding simulations is presented in Table 11.

Which of the models presented above is correct? None of them, of course. It is rare that any statistical model fit to real world data is correct (i.e., completely and accurately describes the underlying process) since the true state of affairs is unknowable and, no doubt, often far more complicated than we can imagine. The best a practitioner can hope for is to select a sensible and interpretable model that is consistent with both the past experience and the observed data, and that also yields good predictions. See Klugman, Panjer, and Willmot (1998) for a general discussion of the modeling process and the stages it encompasses, including model checking and selection. In the Bayesian setting, one approach towards model checking and selection depends on the use of posterior predictive checks. This approach is described in detail in Chapter 6 of Gelman, Carlin, Stern, and Rubin (2004, pages 159-176). Scollnik (2002c) illustrates it in the context of size-of-loss models for exact data. The discussion in the next few paragraphs draws from these sources.

The main idea behind posterior predictive checking as described in Gelman *et al.* (2004) is to simulate a new sample of replicated data for each set of simulated model parameters, and then compare the characteristics of the replicated data samples with those of the original data set. If the model is a good fit to the orginal data set, then the characteristics of the original and replicated data should be similar. To explain this further in a general context, let x denote a vector of observed data, let θ denote the vector of unknown model parameters for the model associated with the observed data, and let x^{rep} denote a vector of replicated data that could be observed if a new sample of observations is sampled from the same model and with the same model parameter values used to generate x. Of course, the observed data x is fixed whereas the values of x^{rep} and θ, written together as (x^{rep}, θ), are unknown and have a joint posterior distribution. We suppose that values of (x^{rep}, θ) are jointly simulated in the course of a simulation based Bayesian analysis (e.g., using WinBUGS; see the WinBUGS User Manual and Scollnik (2002c) for some examples and tips on coding replicated data).

Let $T(x, \theta)$ denote a test quantity or discrepancy measure, defined as a scalar summary of the data and possibly also of the model parameters, that will be used as a standard to compare the observed data to the replicated data. The particular choice(s) of $T(x, \theta)$ will depend on the problem context and the particular characteristic(s) of the model and data that the practitioner is interested in examining. Suppose interest lies in the spread or dispersion of the data. In this case, some obvious test quantity candidates are $T_1(x, \theta) = \max(x_i) - \min(x_i)$ and $T_2(x, \theta) = \sum_{i=1}^{n}(x_i - \bar{x})^2$. Since the test quantity is allowed to depend on the model parameters as well as the data, the sum of the squared deviations of the observed data values around their theoretical means given θ, that is, $T_3(x, \theta) = \sum_{i=1}^{n}(x_i - \mathrm{E}(X_i \mid \theta))^2$, could also be employed. Note that $T_1(x, \theta)$ and $T_2(x, \theta)$ are both fixed, given the observed data, for all values of θ, whereas $T_3(x, \theta)$ will vary with θ. This is inconsequential so far as the method of posterior predictive checking is concerned, except that it affords greater flexibility in the construction of test quantities. On the other hand, recall that traditional non-Bayesian statistics used in classical testing are functions of the data only, and never of the model parameters.

Now, observe that any given discrepancy measure can also be calculated using the posterior simulations of (x^{rep}, θ) in order to obtain values we denote $T(x^{\text{rep}}, \theta)$. These values of $T(x^{\text{rep}}, \theta)$ may be monitored and compared to those of $T(x, \theta)$. In this manner, we can judge whether the model under consideration is generating replicated data with characteristics similar to those exhibited by the original data set. If not, another model may be entertained. In order to measure the lack of fit of the observed data to the model and its associated posterior predictive distribution, Gelman *et al.* (2004) suggest using the tail-area posterior probability, or 'Bayes p-value', of the test quantity. This p-value is defined as the posterior probability that the replicated data is more extreme than the observed data, as measured by the test quantity and given the assumed model. Mathematically, this can be written as

$$\text{Bayes } p\text{-value} \quad = \quad \Pr(\, T(\, x^{\text{rep}}, \theta\,) \geq T(\, x, \theta\,) \,|\, x\,) \;, \tag{5}$$

with the probability understood to be taken over the joint posterior distribution of (x^{rep}, θ). In WinBUGS, this probability can be estimated by monitoring the value of an indicator function defined to be 1 whenever the inequality in the expression above holds, and 0 otherwise (see the WinBUGS User Manual and Scollnik (2002c) for examples). If the tail-area probability (5) is close to 0 or 1 for some meaningful test quantity, then the assumed model is suspect and serious consideration should be given towards entertaining another.

For each of the Bayesian reserving models discussed in this section, we monitored the p-values associated with two test quantities. These test quantities were taken equal to the sample mean and the sample standard deviation. The p-values for these two test quantities are reported in Table 12 in the rows labeled (1) and (2), respectively, for each model. On the basis of these p-values, none of the models is judged to be particularly suspect. We also monitored the posterior distribution for another quantity, this being the sum of squared deviations between the replicated data and the observed data values. As this quantity is a function of both the observed and replicated data it cannot serve as a formal discrepancy measure in the sense of Gelman *et al.* (2004) as discussed above, but it still provides a useful overall posterior summary of a model's fit to the observed data. Its posterior mean values are reported in the row labeled (3) in Table 12, for each model. It appears that the base model performs slightly better, overall, than the two variants. However, this should be monitored over future years and changes introduced to the model if and when necessary.

Other Bayesian approaches to model checking and selection do exist. Ntzoufras and Dellaportas (2002) considered several models for outstanding liabilities, and used a predictive measure for model selection based on a criterion given in Laud and Ibrahim (1995). Scollnik (2000, 2002b) uses the negative log likelihood to assist in Bayesian model selection in a variety of actuarial contexts. Spiegelhalter, Best, Carlin, and van der Linde (2002) present various other Bayesian measures of model complexity and fit. Interested readers are directed to these references for more information.

5 A Bayesian Bornhuetter-Ferguson Model

In the previous section, a full probability model was presented for the single cell development factors. This Bayesian model was analyzed using WinBUGS. The output of this analysis included knowledge concerning the posterior distribution associated with the single cell development factors. The idea pursued in this section is to combine knowledge concerning the claim development pattern given by the posterior distribution of the development factors with additional knowledge concerning the level of ultimate claims associated with each year of origination (i.e., we are adding another level to the model discussed previously). This approach is inspired by the traditional Bornhuetter-Ferguson reserving method (Bornhuetter and Ferguson, 1972), although our implementation and interpretation are plainly different (i.e., being stochastic and Bayesian in nature). As the reader will observe, this approach could be combined with many other Bayesian reserving models (e.g., such as those described in Ntzoufras and Dellaportas, 2002, or de Alba, 2002), not just the one described in the previous section.

Let $R[\,i,j\,]$ denote the proportion of ultimate claims associated with year of origination i, observed by the end of development period j. In terms of the single cell development factors defined previously we have

$$R[\,i,j\,] \;=\; \frac{1}{\prod_{k=j}^{r-1}\,\mathrm{devfac}[\,i,k\,]}\,, \tag{6}$$

for $i = 1,\ldots,r$ and $j = 1,\ldots,r-1$. In WinBUGS, this may be coded as below.

```
for( i in 1 : r ) {
   for( j in 1 : ( r - 1 ) ) {
      R[ i, j ] <- 1 / prod( devfac[ i, j : ( r -1 ) ] )
   }
}
```

The value $1-R[\,i,j\,]$ is equal to the proportion of the ultimate claims associated with year of origination i observed (or to be observed) after development period j, while the value $R[\,i,j-$

$1\,] - R\,[\,i, j\,]$ is equal to the proportion observed (or to be observed) in development period j (i.e., between the ends of period $j-1$ and period j). In the Bayesian context, the $R\,[\,i, j\,]$ values are stochastic variables or nodes, as they are functions of the stochastic development factors. The traditional Bornhuetter-Ferguson method uses fixed point estimates for the $R\,[\,i, j\,]$ values, possibly constructed using observed volume-weighted development factors in place of the development factors appearing in the denominator on the right hand side of (6).

The traditional Bornhuetter-Ferguson reserving method next takes an independent estimate of the ultimate claims associated with each year of origination, and parcels it into the different development periods in a manner consistent with the proportions suggested by the observed development process. (By an independent estimate, we mean one coming from a new source of information.) The total amount of the ultimate claims parceled into development periods that are as yet unobserved for a given year of origination constitute the reserves for that year. Let ultimate[i] denote the independent estimate of the ultimate claims associated with year of origination i, for $i = 1, \ldots, r$. Then, for year of origination i, the amount of the ultimate claims arising in development period 1 is given by

$$\text{ultimate}[\,i\,] * R\,[\,i, 1\,]\,; \tag{7}$$

the amount arising in development period j is given by

$$\text{ultimate}[\,i\,] * (\,R\,[\,i, j\,] \;-\; R\,[\,i, j-1\,]\,) \tag{8}$$

for $j = 2, \ldots, r-1$; and the amount arising in development period r is given by

$$\text{ultimate}[\,i\,] * (\,1 \;-\; R\,[\,i, r-1\,]\,)\,. \tag{9}$$

It is important to note that these amounts are consistent with (i.e., defined on the basis of) the development patterns defined by the development factors in the paragraphs above, and with the estimates of the ultimate claims.

The values (7), (8), and (9) are easily defined in WinBUGS, and then monitored as the simulation proceeds. Illustrative WinBUGS code corresponding to (7), (8), and (9) is given below.

```
for( i in 1 : r ) {
   bfcell.curr[ i, 1 ] <- ultimate[ i ] * R[ i, 1 ]
   for( j in 2 : ( r - 1 ) ) {
      bfcell.curr[ i, j ] <- ultimate[ i ] * ( R[ i, j ] - R[ i, j - 1 ] )
      }
   bfcell.curr[ i, r ] <- ultimate[ i ] *  ( 1 - R[ i, r - 1 ] )
}
```

The revised reserve associated with a given year of origination is equal to the sum of the bfcell.curr[i, j] values over the unobserved development periods for that year. Note that the bfcell.curr[i, j] current money amounts are easily converted to nominal or net present value amounts, if so desired (as were other current money amounts in the previous sections).

The simulated values generated for (7), (8), and (9) are random draws from the posterior distribution for these random quantities (i.e., assuming, of course, that the simulation has been allowed to converge). The simulated values associated with unobserved development periods in a given year of origination are random draws from the posterior distribution for a set of reserves for these periods that have been modified in accordance with a procedure in the spirit of the traditional Bornhuetter-Ferguson method. They allow us to estimate the posterior distribution (i.e., its mean, variance, shape, percentiles, etc.) for the revised reserve, in light of our knowledge concerning the ultimate claim level, whereas the traditional Bornhuetter-Ferguson method yields only a point estimate.

Often, the independent estimate of the ultimate claim amount associated with a given year of origination is constructed from the written premium and the initial expected loss ratio (i.e., defined as their product) for that year. This is our assumption, but others are easily accommodated. Let premium[i] and lossratio[i] denote the written premium and a loss ratio (e.g., some measure of total losses divided by premium) associated with year of origination i, for $i = 1, \ldots, r$. Then,

$$\text{ultimate}[\,i\,] = \text{premium}[\,i\,] * \text{lossratio}[\,i\,]\,,$$

for $i = 1, \ldots, r$. These values are all defined in terms of current money, in order to be consistent with the previous definition of the development factors. The written premium values are assumed to be known. The values of the loss ratios may be fixed (i.e., non-stochastic) or random (i.e., stochastic). In either case, they may also vary, or not, with the year of origination.

The model described in this section will be applied to the same data as before. The necessary written premium values for each year are given in Table 13. (These values appear in Table 16.43 of Booth *et al.* (1999). We assume that they have already been adjusted into current money.) We consider four different analyses, corresponding to four extensions of the base model discussed in the previous section (as summarized in Tables 11 and 12) along the lines described above. These analyses differ only in the assumed manner in which the loss ratio values lossratio[i] are modeled, for $i = 1, \ldots, r$. As usual, these approaches are illustrative and a practitioner may find that another set of assumptions is more appropriate for the particular context under consideration.

For the first loss ratio specification, the loss ratio associated with each year of origin is simply set equal to a fixed value of 0.8. The ultimate claim for a particular year of origin is defined as the product of the written premium and the loss ratio for the year in question. The WinBUGS code for this simple case is given below.

```
for( i in 1 : r ) {
   lossratio[ i ] <- 0.8
   ultimate[ i ] <- premium[ i ] * lossratio[ i ]
}
```

Table 14 lists the posterior means, standard deviations, and 95% probability intervals for the reserves associated with the different years of origin under the base model, revised in light of the additional knowledge concerning the level of ultimate claims provided by the loss ratios as above. These summaries are reported in current money, and also for when adjustments are made for inflation and present value discounting. As before, the inflation and discounting rates are held constant at 12% and 6% effective per annum, respectively. The first column in Table 15 reports the posterior means along with their associated Monte Carlo standard errors. All summaries are based on four chains run for 50,000 iterations, following 50,000 burn-in iterations as before.

For the second specification, an underlying loss ratio parameter mu.lossratio is taken to be normally distributed with a mean of 0.8 and a standard deviation of 0.1 (i.e., a precision of 100). Conditional on this parameter, the loss ratios associated with the different years of origination are independently and identically normally distributed with their common means equal to mu.lossratio and their common standard deviations equal to 0.025 (i.e., precisions equal to 1,600). The WinBUGS code for this case is as follows.

```
mu.lossratio ~ dnorm( 0.8, 100 )

for( i in 1 : r ) {
   lossratio[ i ] ~ dnorm( mu.lossratio, 1600 )
   ultimate[ i ] <- premium[ i ] * lossratio[ i ]
}
```

Abbreviated summary results (i.e., posterior means and Monte Carlo standard errors) for the revised reserves under this analysis are given in the second column of Table 15. The posterior means are similar to those in the first column, but the Monte Carlo standard errors are now slightly larger. This is as expected since the loss ratios are now allowed to vary symmetrically about their previous fixed level of 0.8, thus increasing the variability in the predictions.

The third and fourth loss ratio specifications are somewhat more interesting than the first two, and illustrate a novel modeling twist. The WinBUGS code for the third specification is given below, followed by an explanation.

```
mu.lossratio <- 0.8
```

```
for( i in 1 : r ) {
   lossratio[ i, 1 ] ~ dnorm( mu.lossratio, 1600 )
   lossratio[ i, 2 ] <- sum( y[ i, 1 : r ] ) / premium[ i ]
   lossratio[ i, 3 ] <- cut( lossratio[ i, 2 ] )
   # lossratio[ i, 3 ] ~ dnorm( lossratio[ i, 1 ], 1600 )
   # Use the zeros trick code below in place of line above in
   # order to avoid defining the lossratio[ i, 3 ] node twice.
   #
   zero[ i ] <- 0
   zero[ i ] ~ dpois( lambda[ i ] )
   lambda[ i ] <- 1600 * ( lossratio[ i, 3 ] - lossratio[ i, 1 ] ) *
                        ( lossratio[ i, 3 ] - lossratio[ i, 1 ] ) / 2
   ultimate[ i ] <-  premium[ i ] * lossratio[ i, 1 ]
}
```

In this specification, the nodes lossratio[i, j], for $j = 1, 2, 3$, correspond to loss ratio parameters associated with year of origination i, for $i = 1, \ldots, r$. For ease of discussion, consider just the year of origin k (i.e., $i = k$). The other years will be treated in an analogous fashion. The first of the nodes, lossratio[k, 1], represents the 'true' loss ratio for this year and is sampled from an underlying normal population with mean mu.lossratio. This is a statement of prior knowledge and can be modified as necessary. The second node, lossratio[k, 2], is set equal to the actual loss ratio given the (observed and current simulated) values of $Y[k, j]$ and the known value of premium[k]. Although our assumption is that the unobserved values of $Y[k, j]$ are being simulated in WinBUGS on the basis of the model defined in the previous section, we observe again that they could be simulated by another Bayesian model (e.g., as in Ntzoufras and Dellaportas, 2002, or de Alba, 2002). We would like to 'borrow' the value of the actual loss ratio, lossratio[k, 2], from the first part of the model, and treat it as a random draw from a normal population with mean lossratio[k, 1]. This knowledge would enable us to adjust our prior opinion for the 'true' loss ratio, lossratio[k, 1], by computing (via WinBUGS) a posterior distribution for this parameter, on the basis of the sample observation lossratio[k, 2]. The value of the ultimate claim ultimate[k] used in the calculation of the revised (i.e., Bornhuetter-Ferguson styled) reserves would be determined on the basis of the posterior distribution for lossratio[k, 1], as coded above.

Conceptually, the procedure described above is reasonable but it is complicated by the fact that sample information generally cannot be included more than once in the overall likelihood function. This means that lossratio[k, 2] cannot be assigned a normal distribution since it is already defined as a deterministic function of the stochastic nodes $Y[\,k, j\,]$, for $j = 1, \ldots, r$. In other words, the node lossratio[k, 2] already has a stochastic distribution implicitly defined for it, through the $Y[\,k, j\,]$, for $j = 1, \ldots, r$. However, it is not our intention to define a distribution on the node lossratio[k, 2] in order to adjust the posterior distribution of the unobserved $Y[\,k, j\,]$ nodes used in the calculation of the original reserves for year k. Doing so would effectively 'double-count' the effect of the $Y[\,k, j\,]$ nodes. Rather, we wish to borrow its value from one part of the analysis in order to adjust a prior distribution used in another part, in the calculation of an entirely new set of reserves.

Borrowing information from one part of a model for use in another in a situation like that described above can be accomplished in WinBUGS via the "cut" function. The WinBUGS User Manual describes the cut function as a kind of valve that allows information to flow in one direction but not the other. This allows evidence from one part of a model to play a role in a second part, without feedback from the second part of the model contaminating the first. Hence, the introduction of the node lossratio[k, 3] set equal to the cut value of lossratio[k, 2]. The node lossratio[k, 3] (with numerical value equal to lossratio[k, 2]) can be assigned the desired normal distribution with mean lossratio[k, 1], thus adjusting the prior distribution of lossratio[k, 1], without this assignment filtering back and impacting the posterior distribution of the unobserved $Y[\,k, j\,]$ nodes.

Unfortunately, WinBUGS will not allow lossratio[k, 3] to be directly assigned as a stochastic node in the usual fashion. The obvious standard declaration in WinBUGS,

$$\mathsf{lossratio[i,3]} \sim \mathsf{dnorm(lossratio[i,1], 1600)}\,, \tag{10}$$

results in a compilation error. This is due to the fact that lossratio[k, 3] is already defined and set equal to (the cut value of) lossratio[k, 2]. For this reason, the standard declaration mentioned above is commented out in the WinBUGS code given earlier. In its place are several lines of code implementing the "zeros trick" that is described in the WinBUGS User Manual. This trick constructs terms in the overall model's likelihood function that are proportional to those that would be generated by (10). The zeros trick has been used, along with a related "ones" trick, in actuarial modeling contexts by Scollnik (2002b, 2002c). Both tricks are very well described in the WinBUGS User Manual.

The normal populations employed in the third loss ratio specification are all assigned precisions equal to 1,600 (i.e., standard deviations equal to 0.025). These values may be adjusted as the practitioner sees. They could also be defined so as to vary with the year of origination. Equal precisions are used in the third specification to illustrate the situation in which equal weight is being given to the prior information and the sample information concerning the true loss ratio lossratio[k, 1]. The fourth loss ratio specification is identical to the third, except that the precision involved in (10) is changed to 400 (i.e., a standard deviation of 0.05). This illustrates the case in which more faith is being placed in the prior information than in the sample based information, as discussed below.

Readers familiar with the Kimeldorf-Jones method of graduation may recall how normal distributions blend together in accordance with the following basic normal distribution theory (see Section 5.4 of London, 1985). Suppose the random variable T has a marginal distribution that is normal,

$$T \sim \text{normal}\,(\,M\,,\,A\,)\,,$$

and suppose that the conditional distribution of U, given T, is normal as well, such that

$$T\,|\,U \sim \text{normal}\,(\,T\,,\,B\,)\,.$$

The first parameters in these normal distributions are means while the second are precisions (inverse variances). Then, the conditional distribution of T, given U, is also normal and is such that

$$U\,|\,T \sim \text{normal}\,(\,V\,,\,C\,)\,.$$

In this expression, the conditional mean and precision are equal to

$$V = \frac{AM + BU}{A + B} = \left(\frac{A}{A + B}\right) M + \left(\frac{B}{A + B}\right) U$$

and

$$C = A + B \,,$$

respectively. These relations can be used as guides when setting the precisions in the model defining the revised reserves. For example, the values $A = 1{,}600$ and $B = 400$ will assign weights of 0.8 and 0.2 to the prior mean and the single sample observation, respectively, in the calculation of the conditional mean for T, given U.

The situation with the revised reserves is a little more complicated since the node corresponding to U (i.e., lossratio[k, 3]) is dynamic and changes values as the simulation progresses. Still, an examination of the abbreviated summary results in Table 15 for the revised reserves under the third loss ratio specification reveals that they are approximately midway between the reserves given for the base model in Table 11 and the revised reserves under the first loss ratio specification in Table 15; under the fourth specification, they are approximately equal to a 20% / 80% weighted average of the two.

England and Verrall (2002a, Section 6.3) discuss an over-dispersed Poisson chain ladder model in which (using their notation) the incremental claims $C_{i,j}$ are independently distributed with means $x_i y_j$ and variances $\phi x_i y_j$. Here, x_i is the expected ultimate claim in year of origination i, y_j is the proportion of ultimate claims to emerge in development year j, and the parameter ϕ is positive (and greater than unity if the model is over-dispersed). See Renshaw and Verrall (1998) for more details of this model. England and Verrall (2002a, page 19) state that 'the Bornhuetter-Ferguson technique assumes that there is prior knowledge about' the row parameters $\{x_i\}$, 'making the Bornhuetter-Ferguson technique analogous to a Bayesian approach.' These authors then assign independent gamma prior distributions to the x_i parameters, with expected values set equal to the anticipated value of the ultimate claims. Clearly, their approach differs from the one introduced in this section.

Whether the approach described by England and Verrall (2002a, 2002b) or the one described in this section is closer in spirit to the traditional Bornhuetter-Ferguson technique is debatable and irrelevant. However, we will point out one last time that the Bornhuetter-Ferguson style of approach developed in this section can be combined with many other Bayesian reserving models when the additional written premium information and prior knowledge regarding the loss ratios are available. The impact of the Bornhuetter-Ferguson styled revision can be adjusted by the practitioner through the precisions of the normal distributions involved, as discussed and illustrated above.

6 Final Remarks

This paper explored some Bayesian models for loss reserving along with their implementation using WinBUGS. WinBUGS enabled a variety of prior density specifications to be considered, allowed modifications in the models to be introduced, and yielded reliable information about the predictive accuracy associated with the reserve estimates. Model selection via posterior predictive checking was also considered. A secondary modeling approach, inspired by the traditional Bornhuetter-Ferguson method of reserving, allowed reserve revision to take place in order to incorporate external knowledge relating to written premiums and loss ratios. WinBUGS code implementing the models presented in this paper (under WinBUGS 1.4) is available upon request from the author. Hopefully, the ideas presented in this paper will inspire and encourage others to consider how other traditional reserving methods may be embedded in a Bayesian framework.

References

Booth, P., R. Chadburn, D. Cooper, S. Haberman, and D. James. 1999. *Modern Actuarial Theory and Practice*. New York, NY: Chapman & Hall/CRC Press LLC.

Bornhuetter, R.L., and R.E. Ferguson. 1972. "The Actuary and IBNR," *Proceedings of the Casualty Actuarial Society* **LIX**: 181-195.

Brown, R.L., and L.R. Gottlieb. *Introduction to Ratemaking and Loss Reserving for Property and Casualty Insurance*, Second Edition. Winstead, CT: ACTEX Publications, Inc.

Cairns, A.J.G. 2000. "A Discussion of Parameter and Model Uncertainty in Insurance," *Insurance: Mathematics and Economics* **27** (3): 313-330.

Congdon, P. 2003. *Applied Bayesian Modelling*. New York, NY: John Wiley & Sons, Inc.

Congdon, P. 2001. *Bayesian Statistical Modelling*. New York, NY: John Wiley & Sons, Inc.

de Alba, E. 2002. "Bayesian Estimation of Outstanding Claim Reserves," *North American Actuarial Journal* **6** (4): 1-20.

Efron, B., and R.J. Tibshirani. 1993. *An Introduction to the Bootstrap*. New York, NY: Chapman and Hall.

England, P., and R.J. Verrall. 2002a. "Stochastic Claims Reserving in General Insurance," Sessional Meeting Paper Presented to the Institute of Actuaries on January 28, 2002. Online at www.actuaries.org.uk in the library and publications section.

England, P., and R.J. Verrall. 2002b. "Stochastic Claims Reserving in General Insurance," (with discussion), *British Actuarial Journal* **8**: 443-518.

Dey, D.K., S.K. Ghosh, and B.K. Mallick. 2000. *Generalized Linear Models: A Bayesian Perspective*. New York, NY: Marcel-Dekker, Inc.

Gelman, A., J.B. Carlin, H.S. Stern, and D.B. Rubin. 2004. *Bayesian Data Analysis*, Second Edition. New York, NY: Chapman and Hall/CRC.

Harnek, R.F. 1966. "Formula Loss Reserves," *Insurance Accounting and Statistical Proceedings*.

Hayne, Roger M. 2003. "Measurement of Reserve Variability," *Casualty Actuarial Society Forum* **Fall 2003**: 141-171.

Kimeldorf, G.S., and D.A. Jones. 1967. "Bayesian Graduation," *Transactions of the Society of Actuaries*, **XIX**.

Klugman, S.A., Panjer, H.H., and Willmot, G.E. 1998. *Loss Models: From Data to Decisions*. New York, NY: John Wiley & Sons, Inc.

Kremer, E. 1982. "IBNR-Claims and the Two-Way Model of ANOVA," *Scandinavian Actuarial Journal* **1**: 47-55.

Lamps, D. 2002a. "A Set of Bayesian Models for Analysis of Claim Lags," *Actuarial Research Clearing House* **2003.1**. (Proceedings of the 37th Actuarial Research Conference, The University of Waterloo, Waterloo, Ontario, August 7-10, 2002.) Available on-line at www.soa.org.

Lamps, D. 2002b. "An Application of Markov Chain Monte Carlo Theory to Claim Run Off Analysis." Presented at the Annual Meeting of the Society of Actuaries, Boston, Massachusetts, October 27-30, 2002.

Lamps, D. 2002c. "Bayesian Analysis of Claim Lag Data Using WinBUGS Software." Presented at the Annual Meeting of the Society of Actuaries, Boston, Massachusetts, October 27-30, 2002.

Laud, P.W., and J.G. Ibrahim. 1995. "Predictive Model Selection," *Journal of the Royal Statistical Society, Series B* **57**: 247-62.

London, D.L. 1985. *Graduation: The Revision of Estimates*. Winsted, Connecticut: ACTEX Publications.

Mack, T. 1993. "Distribution Free Calculation of the Standard Error of Chain Ladder Reserve Estimates," *ASTIN Bulletin* **23**: 213-225.

McClenahan, Charles L. 2003. "Estimation and Application of Ranges of Reasonable Estimates," *Casualty Actuarial Society Forum* **Fall 2003**: 213-230.

Ntzoufras, I., and P. Dellaportas. 2002. "Bayesian Modeling of Outstanding Liabilities Incorporating Claim Count Uncertainty," *North American Actuarial Journal* **6** (1): 113-128.

Renshaw, A.E., and R.J. Verrall. 1998. "A Stochastic Model Underlying the Chain-ladder Technique," *British Actuarial Journal* **4**: 903-923.

Rosenberg, M.A., and V.R. Young. 1999. "A Bayesian Approach to Understanding Time Series Data," *North American Actuarial Journal* **3** (2): 130-143.

Scollnik, D.P.M. 2001. "Actuarial Modeling with MCMC and BUGS," *North American Actuarial Journal* **5** (2): 96-124.

Scollnik, D.P.M. 2002a. "Implementation of Four Models for Outstanding Liabilities in WinBUGS: A Discussion of Ntzoufras and Dellaportas," *North American Actuarial Journal* **6** (1): 128-136.

Scollnik, D.P.M. 2002b. "Regression Models for Bivariate Loss Data," *North American Actuarial Journal* **6** (4): 67-80.

Scollnik, D.P.M. 2002c. "Modeling Size-of-Loss Distributions for Exact Data in WinBUGS," *Journal of Actuarial Practice* **10** (1-2): 193-218.

Shapland, Mark R. 2003. "Loss Reserve Estimates: A Statistical Approach for Determining "Reasonableness"," *Casualty Actuarial Society Forum* **Fall 2003**: 321-360.

Spiegelhalter, D.J., Best, N.G., Carlin, B.P., and A. van der Linde. 2002. "Bayesian Measures of Model Complexity and Fit," (with discussion and rejoinder), *Journal of the Royal Statistical Society, Series B* **64**: 583-639.

Taylor, G.C. 2000. *Loss Reserving: An Actuarial Perspective*. Norwell, MA: Kluwer Academic Publishers.

Verrall, R.J. 2001. "A Stochastic Model for the Bornhuetter-Ferguson Method of Claims Reserving." Paper presented at the Fifth International Congress on Insurance: Mathematics and Economics, Penn State University, July 23-25, 2001.

Verrall, R.J. 1990. "Bayes and Empirical Bayes Estimation for the Chain Ladder Model," *ASTIN Bulletin* **20**: 217-243.

Verrall, R.J. 1989. "A State Space Representation of the Chain Ladder Linear Model," *Journal of the Institute of Actuaries* **116**: 589-609.

Table 1

Incremental Claim Amounts, Adjusted for Inflation.

Year of	Development Period (j)					
Origin	1	2	3	4	5	6
1991	115,239	56,055	14,691	10,255	5,530	3,102
1992	121,528	57,911	15,221	9,339	5,336	
1993	115,427	53,957	12,498	8,543		
1994	113,008	46,666	11,050			
1995	109,881	53,408				
1996	128,982					

Table 2

Cumulative Claim Amounts, Adjusted for Inflation.

Year of	Development Period (j)					
Origin	1	2	3	4	5	6
1991	115,239	171,294	185,985	196,240	201,770	204,872
1992	121,528	179,439	194,660	203,999	209,335	
1993	115,427	169,384	181,882	190,425		
1994	113,008	159,674	170,724			
1995	109,881	163,289				
1996	128,982					

Table 3

Ratios of Cumulative Claim Amounts in Successive Periods.

Year of Origin	Development Period (j) 1 - 2	2 - 3	3 - 4	4 - 5	5 - 6
1991	1.486424	1.085765	1.055139	1.028180	1.015374
1992	1.476524	1.084825	1.047976	1.026157	
1993	1.467456	1.073785	1.046970		
1994	1.412944	1.069204			
1995	1.486053				
1996					
Volume Weighted Development Factor	1.466014	1.078642	1.050019	1.027149	1.015374

Table 4

Future Cumulative Claim Amounts in Successive Periods.

Year of Origin	Development Period (j) 1	2	3	4	5	6
1991						
1992						212,553
1993					195,595	198,602
1994				179,263	184,130	186,961
1995			176,130	184,940	189,961	192,882
1996		189,090	203,960	214,162	219,976	223,358

Table 5

Future Incremental Claim Amounts in Successive Periods.

Year of Origin	Development Period (j) 1	2	3	4	5	6
1991						
1992						3,218
1993					5,170	3,007
1994				8,539	4,867	2,831
1995			12,841	8,810	5,021	2,920
1996		60,107	14,870	10,202	5,814	3,382

Table 6

Future Incremental Claim Amounts in Successive Periods,
with Inflation of 12% Per Annum.

Year of Origin	Development Period (j) 1	2	3	4	5	6
1991						
1992						3,605
1993					5,790	3,772
1994				9,564	6,105	3,977
1995			14,382	11,051	7,054	4,595
1996		67,320	18,653	14,333	9,149	5,960

Table 7

Future Incremental Claim Amounts in Successive Periods, with Inflation of 12% Per Annum, Discounted at 6% Per Annum.

Year of Origin	Development Period (j) 1	2	3	4	5	6
1991						
1992						3,400
1993					5,462	3,357
1994				9,023	5,433	3,339
1995			13,568	9,835	5,923	3,640
1996		63,510	16,601	12,034	7,247	4,454

Table 8

Reserves, with and without Future Inflation and Present Value Discounting.

Year of	Reserve		
Origin	Current Money	Nominal	Net Present Value
1992	3,218	3,605	3,400
1993	8,177	9,562	8,820
1994	16,237	19,646	17,795
1995	29,592	37,082	32,966
1996	94,375	115,415	103,845
Total	151,599	185,310	166,826

Table 9

Predictive Means, SDs, and 95% Probability Intervals for Future Incremental Claim Amounts.

Year of	Development Period (j)				
Origin	2	3	4	5	6
1992					3,682 (8,050)
					(-12,260, 19,880)
1993				5,375 (6,362)	3,405 (7,566)
				(-7,273, 18,110)	(-11,490, 18,640)
1994			8,658 (5,379)	5,057 (6,010)	3,238 (7,095)
			(-2,052, 19,400)	(-6,832, 17,090)	(-10,760, 17,480)
1995		12,830 (4,990)	8,934 (5,556)	5,201 (6,203)	3,341 (7,325)
		(2,916, 22,770)	(-2,090, 20,010)	(-7,078, 17,630)	(-11,070, 18,150)
1996	59,910 (3,861)	14,840 (5,773)	10,310 (6,415)	6,048 (7,173)	3,861 (8,490)
	(52,160, 67,550)	(3,362, 26,380)	(-2,407, 23,160)	(-8,116, 20,500)	(-12,860, 20,960)

Table 10

Bayesian Estimation of Reserves, with and without Future Inflation and Present Value Discounting: Predictive Means, SDs, and 95% Probability Intervals.

Year of Origin	Reserve		
	Current Money	Nominal	Net Present Value
1992	3,682 (8,050)	4,123 (9,016)	3,890 (8,506)
	(-12,260, 19,880)	(-13,730, 22,270)	(-12,950, 21,010)
1993	8,780 (9,964)	10,290 (11,960)	9,481 (10,880)
	(-10,610, 29,100)	(-12,990, 34,760)	(-11,700, 31,720)
1994	16,950 (10,930)	20,590 (14,100)	18,610 (12,340)
	(-4,100, 39,440)	(-6,527, 49,670)	(-5,133, 44,070)
1995	30,300 (12,560)	38,140 (17,480)	33,830 (14,700)
	(6,239, 56,470)	(4,690, 74,610)	(5,719, 64,470)
1996	94,970 (15,240)	116,500 (23,290)	104,700 (18,640)
	(66,050, 126,700)	(72,400, 165,100)	(69,320, 143,400)
Total	154,700 (38,830)	189,700 (53,720)	170,050 (45,280)
	(79,510, 234,300)	(85,770, 299,800)	(82,920, 263,300)

Table 11

Predictive Mean Reserves and Monte Carlo Standard Errors.

Year of Origin	Reserve in Current Money		
	Base Model	Variant 1	Variant 2
1992	3,682 (18.38)	4,101 (19.81)	5,413 (38.12)
1993	8,780 (22.49)	9,351 (23.27)	10,970 (44.43)
1994	16,950 (25.49)	17,560 (25.38)	19,120 (47.34)
1995	30,300 (29.23)	31,000 (29.67)	32,600 (52.31)
1996	94,970 (34.61)	95,520 (36.19)	97,370 (58.8)
Total	154,700 (89.19)	157,500 (93.45)	165,500 (215.0)

Table 12

Illustrative Posterior Predictive Measures of Model Fit.

Measure of Fit	Model		
	Base Model	Variant 1	Variant 2
(1)	0.4897	0.4949	0.5189
(2)	0.5899	0.5725	0.5625
(3)	5387	5606	5447

Table 13

Written Premium for Each Year of Origin.

Year	Premium
1991	292,674
1992	303,647
1993	283,718
1994	267,086
1995	275,545
1996	319,082

Table 14

Bayesian Revised Reserves, with and without Future Inflation and Present Value Discounting: Predictive Means, SDs, and 95% Probability Intervals.

Year of Origin	Reserve		
	Current Money	Nominal	Net Present Value
1992	3,856 (9,075)	4,318 (10,160)	4,074 (9,588)
	(-15,030, 21,010)	(-16,830, 23,530)	(-15,880, 22,200)
1993	9,460 (10,910)	11,080 (13,110)	10,210 (11,920)
	(-13,390, 30,000)	(-16,320, 35,770)	(-14,740, 32,650)
1994	18,580 (11,390)	22,520 (14,790)	20,380 (12,900)
	(-5,428, 39,990)	(-8,553, 50,420)	(-6,785, 44,650)
1995	33,710 (12,080)	42,320 (17,140)	37,590 (14,260)
	(8,250, 56,470)	(6,397, 74,680)	(7,611, 64,530)
1996	107,600 (10,080)	131,600 (17,740)	118,400 (13,250)
	(86,300, 126,500)	(94,470, 165,100)	(90,580, 143,300)
Total	173,200 (36,450)	211,900 (51,500)	190,600 (42,930)
	(97,700, 243,100)	(105,400, 310,600)	(101,800, 273,000)

Table 15

Posterior Mean Revised Reserves and Monte Carlo Standard Errors.

Year of	Revised Reserve in Current Money			
Origin	Specification 1	Specification 2	Specification 3	Specification 4
1992	3,856 (20.13)	3,816 (21.17)	3,596 (19.59)	3,739 (20.31)
1993	9,460 (24.97)	9,471 (25.26)	8,897 (24.2)	9,245 (24.8)
1994	18,580 (26.77)	18,630 (27.26)	17,490 (25.6)	18,170 (25.96)
1995	33,710 (27.86)	33,690 (29.47)	31,660 (28.66)	32,890 (28.41)
1996	107,600 (22.6)	107,500 (37.51)	101,000 (35.57)	105,000 (30.25)
Total	173,200 (84.6)	173,100 (97.52)	162,600 (88.68)	169,000 (86.68)