Bayesian Solvency Analysis with Autocorrelated Observations

M. Mendoza and L.E. Nieto-Barajas
Departamento de Estadística, ITAM, México

Abstract

Most financial institutions are required to comply with a minimum capital rule in order to face their obligations during a given period of time. Due to the random nature of the financial flows involved, the problem of assessing the amount of capital required must be analised within a stochastic framework and the solution can be reduced to the estimation of a selected quantile. Given the financial impact of an specific capital requirement, a proper and careful choice of the underlying model is of great relevance. Here, we address the problem by proposing two autocorrelated models to describe the relative severity after a suitable transformation and compare the results with those of a model which assumes independence among observations. We undertake a full Bayesian analysis and posterior inference is obtained via a Metropolis-Hasting scheme. Results are illustrated with mexican insurance industry data.

Key Words: Solvency Analysis; Prediction; Reference priors; log-normal distribution; autocorrelated observations.

1 Introduction

As in many other countries, in Mexico financial institutions are supervised by the government through a number of specialised agencies. Given the fact that most of the assets that these institutions handle are not of their property but belong to their clients, one of the objectives of the supervision is to assure that institutions will have enough resources to face the requirements of their clients at any time. In other words, institutions are supervised in order to guarantee their solvency.

If we define Z_t and Y_t as the available resources and the requirements of the clients at time t respectively, a firm is said to be solvent at this point of time if $Y_t \leq Z_t$. Typically, $Z_t = X_t + C_t$ where X_t is the amount of money collected from the clients whereas C_t describes the money obtained from other sources (the shareholders of the institution, for example). In principle, X_t should be sufficient for solvency purposes. However, if that is not the case, the regulatory agency can require the additional resources C_t from the company. Alternative ways of assessing the solvency condition of an institution have been proposed, see for example, the document DAFFE (2001).

In any case and for the sake of simplicity, solvency is usually evaluated periodically (say, each year). In addition, apart from the periodical evaluation, the supervising agencies are interested in predicting the solvency condition one period ahead. Thus, given the current information $\{X_t, Y_t\}$ for $t = 1, 2, \ldots, T$ we are interested in describing the behaviour of Y_{T+1}, so that given X_{T+1}, the amount of additional resources C_{T+1} can be determined. Due to the random nature of the phenomena, the idea of a guaranteed solvency condition must be replaced by its stochastic equivalent. Then the relevant problem is that of finding a prespecified upper quantile for Y_{T+1}. The company is then said to meet the q-level of solvency if it is able to guarantee an amount of money K^q such that $P(Y_{T+1} \leq K^q) = q$. Usually, q is chosen to be some value close to one.

Given the usual strong positive relation between Y_t and X_t and from a purely statistical point of view, it might be reasonable to produce predictions for Y_{T+1} using a conditional model $P(y_{T+1}|x_{T+1})$ instead of the marginal distribution $P(y_{T+1})$.

If we consider the particular case of the insurance industry, the $\{X_t\}$ series describes the premiums collected by a company whereas the $\{Y_t\}$ series stands for the claims presented by the clients. In this case, the solvency condition means that the firm has a level of resources K^q that the future claims Y_{T+1} will exceed only with a low probability $1 - q$. In other words and recalling the usual regulatory language, the insurance company is protected against large amounts of claims with high probability.

Within the same context of insurance, the series of claims $\{Y_t\}$ is typically related to the series of premiums $\{X_t\}$ in a positive way. In fact, if for any year t the company sold more policies or got more clients, a larger amount of premiums would be collected and, at the same time, it would be natural to expect an increase in the amount of claims. This implies that that a large value of Y_t, is not necessarily related to a situation of potential insolvency. As a consequence, solvency analysis is usually conducted in terms of the relative severity of claims or loss rate defined as $V_t = Y_t/X_t$. As an additional advantage, this ratio cancels out some other irrelevant effects such as inflation.

In this paper we shall focus on finding the upper quantiles of V_t and thus of Y_t assuming a log-normal distribution under three scenarios: conditional independence, common autocorrelation for all t and different autocorrelation for each t. Whenever possible in all cases, a reference prior for all the parameters involved is used to obtain the final results.

In Section 2 we will explain the basic structure of the problem and will briefly comment on some elementary procedures, currently used by regulatory agencies (in Mexico). In Section 3 we will use a log-normal distribution and assume that the observations are conditionally independent. Section 4 deals with the case of assuming that observations are dependent with a common correlation structure. In Section 5 we propose an alternative model to deal with autocorrelated observations that allows the observations to have different correlations in pairs. Finally, we conclude with a discussion in Section 6.

2 The basic structure

In the sequel we shall assume that for any year t, the value of the random variable Y_t describing the total amount of claims, appears as the result of the following generating process:

$$Y_t = V_t\, X_t \tag{2.1}$$

where, as before, $X_t > 0$ is the random amount of premiums collected in year t and V_t is a positive random variable describing the relative severity for this specific kind of insurance at year t. This structure is particularly convenient if we notice that, given $X_{T+1} = x$, the random behaviour of Y_{T+1} is simply that of V_{T+1} times the constant x. So, the conditional distribution of interest $P(y_{T+1}|X_{T+1} = x)$ will be totally known as soon as we specify the value x and a conditional model for V_{T+1} given $X_{T+1} = x$.

In we assume that the relative severity series $\{V_t\}$ is independent of the series of premiums $\{X_t\}$, then under these conditions, a solution may be obtained by selecting a model $P(v_{T+1}|\phi)$ where ϕ is a vector of possibly unknown parameters. The problem of estimating ϕ can be solved if we assume that, for any t and for a given ϕ, $V_1, V_2, \ldots, V_t$ are conditionally independent random observations from the same model $P(v|\phi)$.

In the past, Mexican regulations have been operating on the basis of some naive solutions to this problem (and probably that is the case for some other countries as well). Initially, V_t was assumed to follow a normal distribution, a very unlikely assumption since this variable can only take positive values and historical data show evidence of skewness. Recently, the density for V_t has been approximated through a mixture of normal kernels. In both cases the (frequentist) adjusted models have been considered as the true models and the required quantiles have been calculated on that basis. In the following section we introduce a first, very simple, Bayesian solution to this problem assuming a log-normal distribution.

3 Independence case

3.1 The model

From equation (2.1), if we define $U_t = \log(V_t)$, we have that

$$\log(Y_t) = U_t + \log(X_t) \tag{3.1}$$

so that, in particular, if $X_{T+1} = x$, we get $\log(Y_{T+1}) = U_{T+1} + \log(x)$. In this case and given the nature of V_t, we might assume that $\mathbf{U}_{(T)} = U_1, U_2, \ldots, U_T$ is a random sample from a normal distribution $\mathrm{N}(\mu, \tau)$ with mean μ and precision τ.

3.2 Prior and posterior distributions

In order to produce the predictive distribution for the future observation U_{T+1}, we adopt the usual reference prior $P(\mu, \tau) \propto \tau^{-1}$ (Bernardo & Smith, 1994), and get the well known result:

$$P(\mu, \tau | \mathbf{U}_{(T)}) = \mathrm{N}(\mu | \bar{u}, T\,\tau)\mathrm{Ga}(\tau | (T-1)/2, S_u^2/2) \tag{3.2}$$

where $S_u^2 = \sum_{t=1}^{T} (u_t - \bar{u})^2$, $\mathrm{N}(\cdot|\mu, \tau)$ is a normal distribution with mean μ and precision τ and $\mathrm{Ga}(\cdot|a, b)$ stands for a Gamma distribution with mean a/b. As an immediate consequence we obtain

$$U_{T+1} | \mathbf{U}_{(T)} \sim \mathrm{St}\left(\bar{u}, \frac{T(T-1)}{(T+1)S_u^2}, T-1\right). \tag{3.3}$$

where $\mathrm{St}(c, d, r)$ refers to a Student distribution with location parameter c, precision parameter d and r degrees of freedom. Recalling that $\log(Y_{T+1}) = U_{T+1} + \log(x)$ it follows that the posterior predictive distribution for $\log(Y_{T+1})$ is given by a Student distribution with location parameter $\bar{u} + \log(x)$, where the precision parameter and the degrees of freedom are those in (3.3). Alternatively, we might say that the required predictive distribution for Y_{T+1} is log-Student with the same parameters. In this case, we have that K^q is given by

$$K^q = \left(\prod_{t=1}^{T} v_t^{1/T}\right) \exp\left[t_{(T-1)}^q \left\{(1 + 1/T)\tilde{\sigma}_u^2\right\}^{1/2}\right](x) \tag{3.4}$$

where $t_{(T-1)}^q$ is the corresponding q-quantile of a standard Student distribution with $T-1$ degrees of freedom and $\tilde{\sigma}_u^2 = S_u^2/(T-1)$.

3.3 Examples

As an illustration of the results that can be obtained with this approach, let us consider the data set in Table 1 which corresponds to the premiums (X_t) and claims

(Y_t) for the Car insurance and Accident and Disease insurance in Mexico during the years 1973 to 1998. Relative severity $\{V_t\}$ is shown as a dotted line in Figure 2 for the car (top graph) and accident and disease (bottom graph) industries. Figure 1 displays the normal probability plots for the logarithms of the relative severity $\{U_t\}$ for the car (top graph) and accident and disease (bottom graph) industries. From these graphs, no extreme evidence can be observed against the assumption of normality in either case.

Table 1: Premiums and claims for Car insurance (X_{1_t}, Y_{1_t}) and Accident and Disease insurance (X_{2_t}, Y_{2_t}). Source: Comisión Nacional de Seguros y Fianzas, Mexico.

Year	X_{1_t}	Y_{1_t}	X_{2_t}	Y_{2_t}
1973	1388	1017	232	95
1974	1741	1257	285	112
1975	2120	1455	353	159
1976	3283	1938	441	217
1977	4642	2899	550	257
1978	6045	3835	683	304
1979	8163	5862	847	376
1980	11633	9365	1129	537
1981	18259	14425	1641	858
1982	29020	22945	2327	1389
1983	47811	35914	3721	2180
1984	81996	43115	6559	3608
1985	122253	69698	10866	6705
1986	177045	152267	21656	15347
1987	369098	420049	57978	45952
1988	996911	1026628	197162	171477
1989	1650096	1013856	357664	278147
1990	2223373	1523317	598800	446188
1991	3310847	2136837	937137	658831
1992	5018762	2967816	1231516	849876
1993	5741010	3512496	1481307	954805
1994	6096368	3869226	1773392	1208944
1995	6367129	4884735	2132637	1590105
1996	7865581	5471037	3028912	2050374
1997	10655257	7224987	4062261	2819082
1998	14762652	9321247	5192831	3473580

Example 1 (*Car insurance*). In the case of the car insurance data, the posterior

predictive distribution for the relative severity in the forthcoming year V_{1999} is a log-Student with location -0.35, precision 32.70 and 25 degrees of freedom. This distribution is plotted as the solid line in the top graph of Figure 3 and the corresponding q-quantile for the claims Y_{1999} is given by

$$K^q = (0.7015) \exp\left(0.1779\, t^q_{(25)}\right) (x). \tag{3.5}$$

So that, for example, if we use $q = 0.95$ we find that $K^{0.95} = 0.9505(x)$. Therefore, we can say that, in 1999, from each peso collected as premium the firm must be prepared to pay up to 95.05 cents to their clients if it is desired to guarantee the solvency condition with probability 0.95. To have an idea how large this quantile is compared with the relative severity achieved by previous years, see the solid line of the top graph in Figure 2.

Example 2 (*Accident and disease insurance*). A similar calculation can be made using the accident and disease insurance data. In this case, the conditional predictive distribution for V_{1999}, is a log-Student with location -0.51, precision 19.50 and 25 degrees of freedom. The solid line of bottom graph in Figure 3 corresponds to this conditional predictive distribution. The corresponding 0.95-quantile of Y_{1999} is then $K^{0.95} = (0.8825)x$ and thus, for each peso collected by an insurance company from accidents and diseases, 88.25 cents should be kept as reserve if the (0.95)-level of solvency is to be attained. Figure 2 (bottom graph) shows the relative severity of previous years and the quantile $K^{0.95}$ for the future claim Y_{1999} given $X_{1999} = 1$ (see the solid line).

4 Autocorrelated case with common ρ

Results obtained from the former examples show a typical behaviour for this type of models. Since the observations are considered i.i.d., the most extreme data have a noticeable impact on the required quantile disregarding the time period when occurred. In the particular context discussed in this paper this fact might imply an over estimation of the reserves required for year $T + 1$ with the corresponding financial costs.

As one of the simplest strategies available to deal with this situation we might consider a trend component for the series as well as some dependence structure. Here, we shall focuss on some models where the observations are assumed to be correlated and hence, in this section we propose a first order autorregresive process.

4.1 The model

Let us assume that $U_1, U_2, \ldots, U_{T+1}$ are normal random variables such that

$$U_1 \sim \mathrm{N}(\mu, \tau)$$

$$U_t|U_{t-1} \sim \mathrm{N}(\mu + \rho(u_{t-1} - \mu), \tau^{-1}(1-\rho^2)); \quad t = 2, \ldots, T+1.$$

In other words, we are assuming that the vector $\mathbf{U}$ of dimension $T+1$ follows a multivariate normal distribution $\mathrm{N}(\mu\mathbf{1}, \mathbf{R})$ where the covariance matrix $\mathbf{R}$ has typical element $r_{ij} = \tau^{-1}\rho^{|i-j|}$.

It should be noticed that with this assumptions, $\{U_t\}$ follows a stationary process with marginal distributions $N(\mu, \tau)$, for each $t = 1, \ldots, T+1$. See Nieto-Barajas and Walker (2002) for similar constructions of stationary beta and gamma processes. In fact, the model in Section 3.1 can be obtained as a particular case of this structure just by taking $\rho = 0$.

Now, the problem is that of calculating the predictive distribution for U_{T+1}, given the past observations $\mathbf{U}_{(T)} = (U_1, U_2, \ldots, U_T)$. For this purpose we shall consider the expression

$$U_{T+1}|\mathbf{U}_{(T)} \sim \mathrm{N}\left(\mu + \rho(U_T - \mu),\, \tau^{-1}(1-\rho^2)\right) \tag{4.1}$$

together with the posterior distribution $P(\mu, \tau, \rho\,|\mathbf{U}_{(T)})$.

4.2 Prior and posterior distributions

In order to obtain a posterior distribution for the parameters (μ, τ, ρ) we decided to follow a reference approach. In particular, using Berger & Bernardo (1992) we obtained that the reference prior corresponding to the ordered parameterisation $\{\rho\}, \{\tau\}$ and $\{\mu\}$ is

$$P(\mu, \tau, \rho) \propto \tau^{-1}\frac{(1+\rho^2)^{1/2}}{(1-\rho^2)}; \qquad \mu \in \Re,\ \tau > 0,\ \rho \in (-1, 1). \tag{4.2}$$

It is interesting to notice that this prior is the same for any permutation of these three elements. This distribution implies independence *a priori* between ρ and (μ, τ), and for the latter, the marginal distribution is the same as that used for the independence case.

Inferences from the joint posterior distribution can be achieved by noting that, conditional on ρ, we have basically a conjugate normal-gamma model. That is, the conditional posterior distribution for μ is

$$\mu|\mathbf{U}_{(T)}, \tau, \rho \sim \mathrm{N}(\mu_1, \tau_1),$$

where

$$\mu_1 = \frac{\frac{1}{1+\rho}\sum_{t=1}^{T-1}(u_{t+1} - \rho u_t) + u_1}{\frac{(n-1)(1-\rho)}{1+\rho} + 1} \quad \text{and} \quad \tau_1 = \tau\left\{\frac{(n-1)(1-\rho)}{1+\rho} + 1\right\},$$

the conditional posterior distribution for τ is given by

$$\tau|\mathbf{U}_{(T)}, \mu, \rho \sim \mathrm{Ga}(\alpha_1, \beta_1),$$

where,

$$\alpha_1 = \frac{n}{2} \quad \text{and} \quad \beta_1 = \frac{1}{2}\left[\frac{1}{1-\rho^2}\sum_{t=1}^{T-1}\left\{u_{t+1}-\mu-\rho(u_t-\mu)\right\}^2 + (u_1-\mu)^2\right],$$

and finally, the conditional posterior distribution for ρ has the form

$$f(\rho|\mathbf{U}_{(T)},\mu,\tau) \propto (1+\rho^2)^{1/2}(1-\rho^2)^{-(n+1)/2} \times \exp\left[-\frac{\tau}{2(1-\rho^2)}\sum_{t=1}^{T-1}\left\{u_{t+1}-\mu-\rho(u_t-\mu)^2\right\}\right].$$

Posterior simulations can then be obtained via a Gibbs sampling (see for example, Smith and Roberts, 1993) with a Metropolis-Hastings step (see for example, Tierney, 1994) to simulate from the posterior conditional distribution of ρ. Posterior predictive quantiles of U_{T+1} can be estimated straightforward by simulating from expression (4.1).

4.3 Examples

Example 1 (*cont.*) We implemented the autocorrelated model to the car insurance industry data using the reference prior (4.2). The Gibbs sampling was run for 50,000 iterations with a burn-in of 5,000. A histogram of the posterior distribution of ρ is presented in Figure 4 (top graph). From this graph we can clearly see that there is a very low probability of ρ being zero. Moreover, Figure 4 shows that ρ should be positive. This suggests that the assumption of independence in the series $\{V_t\}$ is not supported by the data.

The conditional predictive distribution of V_{1999} is shown in Figure 3 (top graph, dashed line). This predictive distribution is similar to the corresponding conditional predictive distribution for the independence case, but with the right tail decaying faster.

The estimate value for the 0.95-quantile $K^{0.95}$ of the conditional predictive distribution of Y_{1999} is $(0.87)x$. The dashed line of Figure 2 (top graph) corresponds to this estimate value ($x = 1$), which is smaller than the corresponding 0.95-quantile obtained using the independence model. Therefore, when introducing the common autocorrelation in the model, the percentage of premiums from car insurance kept as reserve is smaller.

Example 2 (*cont.*) Now, implementing this autocorrelated model to the accident and disease insurance data, after running the Gibbs sampling for 50,000 iterations with a burn-in of 5,000, and using the reference prior (4.2) we obtained the following results.

Given $X_{1999} = 1$, the conditional predictive distribution of the claims Y_{1999} is plotted as the dashed line in Figure 3 (bottom graph). Differing from the car insurance case, this predictive distribution is noticeably different from the corresponding conditional predictive distribution for the independence case (solid line). The former one has larger precision than the latter one.

The estimate value for the 0.95-quantile $K^{0.95}$ of the conditional predictive distribution of Y_{T+1} is $(0.763)x$. In Figure 2 (bottom graph), taking $x = 1$, we can observe that this estimate value (dashed line) is smaller than the corresponding 0.95-quantile obtained using the independence model (solid line). This would imply in a reduction in the reserve for the next year in the accident and disease insurance.

The posterior distribution of ρ is presented in a histogram in the bottom graph of Figure 4. This figure shows that almost with probability one ρ is positive. This is a strong evidence that the assumption of independence is invalid.

5 Autocorrelated case with different ρ_t

Even though the model in Section 4.1 not only incorporates a certain structure of correlation in the data but allows the user to investigate how adequate the independence assumption might be, it can be generalised in many different ways to deal with more complex dependence patterns. In this section we shall introduce one of these generalisations which uses again the autoregressive structure but makes possible for the autocorrelation coefficient to change over time.

5.1 The model

Let us assume that $U_1, U_2, \ldots, U_{T+1}$ are again normal random variables. However, in this case we shall consider the following generating process:

$$U_1 \sim \mathrm{N}(\mu, \tau)$$

$$U_t | U_{t-1} \sim \mathrm{N}\left(\mu + \rho_{t-1}(u_{t-1} - \mu), \tau^{-1}(1 - \rho_{t-1}^2)\right); \quad t = 2, \ldots, T+1.$$

Now, the joint distribution of $U_1, U_2, \ldots, U_{T+1}$ depends on the (T+2)-vector of parameters $(\mu, \tau, \rho_1, \ldots, \rho_T)$ and the vector $\mathbf{U}$ follows a multivariate normal distribution $N(\mu\mathbf{1}, \mathbf{R})$ where the covariance matrix $\mathbf{R}$ has now typical element

$$r_{ij} = \begin{cases} \tau^{-1} & \text{if } i = j \\ \tau^{-1} \prod_{r=i}^{j-1} \rho_r & \text{if } i < j \end{cases}$$

For this model we have again that the marginal distribution for each U_t is $N(\mu, \tau)$ and the model in Section 4 can be obtained as a particular case if we take $\rho_t = \rho$ for $t = 1, \ldots, T$.

5.2 Prior and posterior distributions

It should be noticed that for this model we actually have more parameters than observations and inferences for ρ_t are expected to involve only the observations U_t and U_{t+1}. This structure implies that the initial distribution for ρ_T will be hardly modified by the data unless we introduce some additional structure. For this purpose, we will consider a hierarchical model for the correlation parameters. We shall assume that, conditional on a hyperparameter ϕ, the coefficients $\rho_1, \ldots, \rho_T$ are independently generated from the model

$$P(\rho|\phi) = \frac{\phi}{2^\phi}(1+\rho)^{\phi-1} I_{(-1,1)}(\rho)\,; \qquad \phi > 0 \tag{5.1}$$

which can be interpreted as a particular instance of a translated beta distribution to the interval $(-1,1)$, and the hyperparameter ϕ follows *a priori* an exponential distribution:

$$P(\phi) = \lambda_0 e^{-\lambda_0 \phi} I_{(0,\infty)}(\phi)\,; \qquad \lambda_0 > 0. \tag{5.2}$$

For the sake of simplicity, and in order to provide a marginal prior for each correlation coefficient which might be considered as an approximation to the uniform distribution over the interval $(-1,1)$, we have set $\lambda_0 = 1$.

Finally, using the usual reference prior for (μ,τ) and assuming prior independence between the rest of the parameters, we have that the joint prior distribution for the $T+2$-vector of parameters is given by

$$P(\mu, \tau, \rho_1, \ldots, \rho_T|\phi) \propto \tau^{-1} \prod_{t=1}^{T} (1+\rho_t)^{\phi-1} \tag{5.3}$$

and $P(\phi) = e^{-\phi} I_{(0,\infty)}(\phi)$.

Once more we have that, given $\rho_1, \ldots, \rho_T$ and ϕ the posterior distribution for μ and τ is a normal-gamma. That is, the conditional posterior distribution for μ is

$$\mu|\mathbf{U}_{(T)}, \tau, \rho_1, \ldots, \rho_T, \phi \sim \mathrm{N}(\mu_1, \tau_1),$$

where

$$\mu_1 = \frac{\sum_{t=1}^{T-1}\left(\frac{u_{t+1}-\rho_t u_t}{1+\rho_t}\right) + u_1}{\sum_{t=1}^{T-1}\left(\frac{1-\rho_t}{1+\rho_t}\right) + 1} \quad \text{and} \quad \tau_1 = \tau\left\{\sum_{t=1}^{T-1}\left(\frac{1-\rho_t}{1+\rho_t}\right) + 1\right\},$$

and the conditional posterior distribution for τ is given by

$$\tau|\mathbf{U}_{(T)}, \mu, \rho_1, \ldots, \rho_T, \phi \sim \mathrm{Ga}(\alpha_1, \beta_1),$$

where,

$$\alpha_1 = \frac{T}{2} \quad \text{and} \quad \beta_1 = \frac{1}{2}\left[\sum_{t=1}^{T-1}\left(\frac{1}{1-\rho_t^2}\right)\{u_{t+1} - \mu - \rho_t(u_t - \mu)\}^2 + (u_1 - \mu)^2\right].$$

Furthermore, the conditional posterior distribution for each ρ_t, has the form

$$f(\rho_t|\mathbf{U}_{(T)},\mu,\tau,\phi) \propto (1+\rho_t)^{\phi-1}(1-\rho_t^2)^{-1/2} \times \exp\left[-\frac{\tau\left\{u_{t+1}-\mu-\rho_t(u_t-\mu)^2\right\}}{2(1-\rho_t^2)}\right],$$

for $t = 1,\ldots,T-1$ and $f(\rho_T|\mathbf{U}_{(T)},\mu,\tau,\phi) = f(\rho_T)$. Finally, the conditional posterior distribution for ϕ is

$$\phi|\mathbf{U}_{(T)},\mu,\tau,\rho_1,\ldots,\rho_T \sim \text{Ga}\left(T+1, 1-\sum_{t=1}^{T}\log(1+\rho_t)+T\log 2\right).$$

Again, posterior simulations can be obtained via a Gibbs sampling (see for example, Smith and Roberts, 1993) with a Metropolis-Hastings step (see for example, Tierney, 1994) to simulate from the posterior conditional distribution of the ρ_t's, and posterior predictive quantiles of U_{T+1} can be estimated straightforward.

5.3 Examples

Example 1 (*cont.*) We implemented the autocorrelated model with different ρ_t to the car insurance industry data. The Gibbs sampling was run again for 50,000 iterations with a burn-in of 5,000. 95% posterior bands for the posterior distribution of all ρ_t $t = 1,\ldots,T$ is presented in Figure 5 (top graph). From this graph we can observe that the ρ_t's vary over time. Almost half of them are above the zero line. This would imply that the correlation is not the same for all times. Again, the assumption of independence is rejected.

The conditional predictive distribution of Y_{1999} given $X_{1999} = 1$ is presented in Figure 3 (dotted-dashed line, top graph). This predictive distribution has slightly smaller variance than the corresponding predictive distributions obtained with the other two models.

The estimate value for the 0.95-quantile $K^{0.95}$ of the conditional predictive distribution of Y_{1999} is $(0.875)x$. The dotted-dashed line of the top graph in Figure 2 corresponds to this estimate value (for $x = 1$), which is slightly larger than the corresponding quantile from the common correlation model, but smaller than the corresponding quantile from the independence model.

Example 2 (*cont.*) Finally, we implemented the autocorrelated model with different ρ_t to the accident and disease insurance data. We run the Gibbs sampling for 50,000 iterations with a burn-in of 5,000.

The conditional predictive distribution of V_{1999} is plotted as the dotted-dashed line of the bottom graph in Figure 3. The precision of this predictive distribution is smaller than the one obtained from the common correlation model but larger than the one from the independence model.

The estimate value for the 0.95-quantile $K^{0.95}$ of the conditional predictive distribution of Y_{T+1} is $(0.788)x$. In Figure 2 (bottom graph) we can observe (taking $x = 1$) that this estimate value (dotted-dashed line) is in the middle of the other two estimates.

95% posterior bands for each ρ_t are presented in Figure 5 (bottom graph). From this graph we can see that all of the intervals are above zero and that the value around 0.65 falls in all of them. This is again a strong evidence that the assumption of independence is invalid, but in agreement with the assumption of a common correlation in the data.

6 Discussion and concluding remarks

The problem of asessing the amount of resources required by a firm to attain a certain level of solvency involves at least two opposite forces. On one hand, the regulatory agencies play a conservative role and thus prefer an overestimation of the quantiles of interest. The companies, looking for business profitability, would be willing to eliminate any unnecessary reserves which might led to severe financial costs.

From a decision-theoretical point of view, this contradictory preferences should be reconciliated and properly reflected in a utility function and thus, in the level q. In any case, it is clear that an accurate model and a robust estimation procedure for these quantiles is crucial. As a first step towards this goal, in this paper we have assumed an underlying log-normal distribution and attention has been focussed on the correlation structure of the observed severity data.

We have proposed a way of estimating the solvency requirement by modelling the logarithm of the relative severity with three different processes (independence, common autocorrelation and different autocorrelation for each t). Inference from these three models can be easily obtained by implementing a Gibbs sampling.

As for the specific data considered, the autocorrelated models seem to be more appealing since in both examples there is strong evidence against the independence hypothesis even though the suggested dependence pattern is not the same. The most important practical consequence is that the estimated quantiles are smaller for the autocorrelated models as compared with the independence structure. Reductions range from 8.8% to 9.3% for the car data and from 10.7% to 13.5% for the accident data. From a financial perspective these differences might be overwhelming.

The analysis in this paper is intended to shed some light on the practical consequences of assuming different models in the context of solvency evaluation. The particular set of structures proposed here constitute a nested sequence of models. All of them are simple, can be fitted without any problem and seem to describe fairly well the data from the examples considered. In any case, some further research could led to some, more general, models. Particularly, a robust version of the log-normal distribution might be considered, In addition some other transformations, apart from the logarithm, could be explored.

References

Berger, J.O. and Bernardo, J.M. (1992). Ordered group reference priors with applications to a multinomial problem. *Biometrika* **79**, 25-37.

Bernardo, J.M. and Smith, A.F.M. (1994). *Bayesian Theory.* Whiley, Chichester.

Directorate for Financial, Fiscal and Enterprise Affairs Insurance Committee (DAFFE). (2001). Assessing the financial health of insurance undertakings: Solvency rules. *Organisation for Economic Co-operation and Development.*

Nieto-Barajas, L.E. and Walker, S.G. (2002). Markov beta and gamma processes for modelling hazard rates. *Scandinavian Journal of Statistics* **29**, 413-424.

Smith, A.F.M. & Roberts, G.O. (1993). Bayesian computations via the Gibbs sampler and related Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society, Series B* **55**, 3-23.

Tierney, L. (1994). Markov chains for exploring posterior distributions. *Annals of Statistics* **22**, 1701-1722.

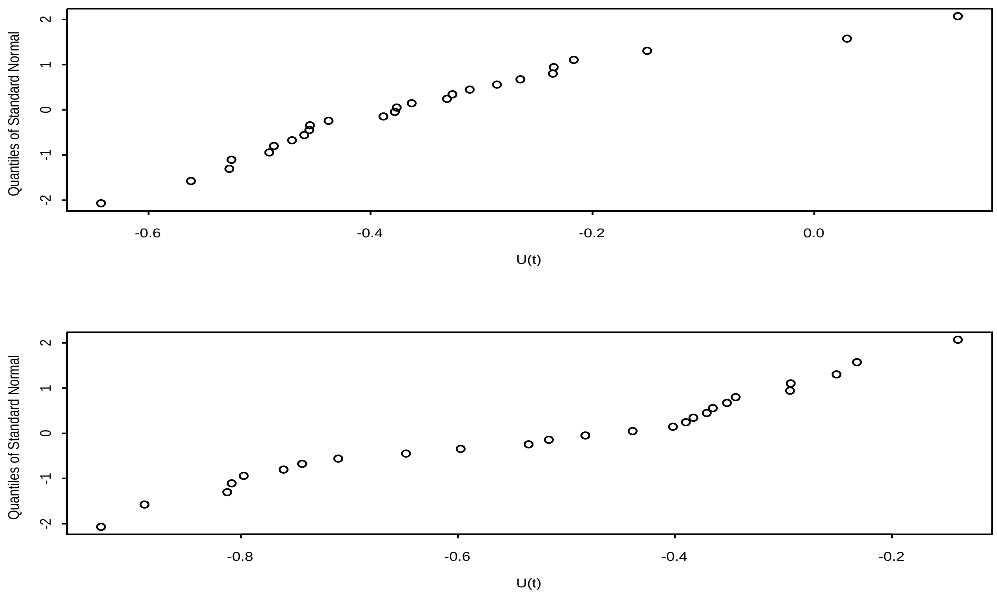

Figure 1: Normal probability plot of $U(t)$. Car insurance (top graph) and accident and disease insurance (bottom graph).

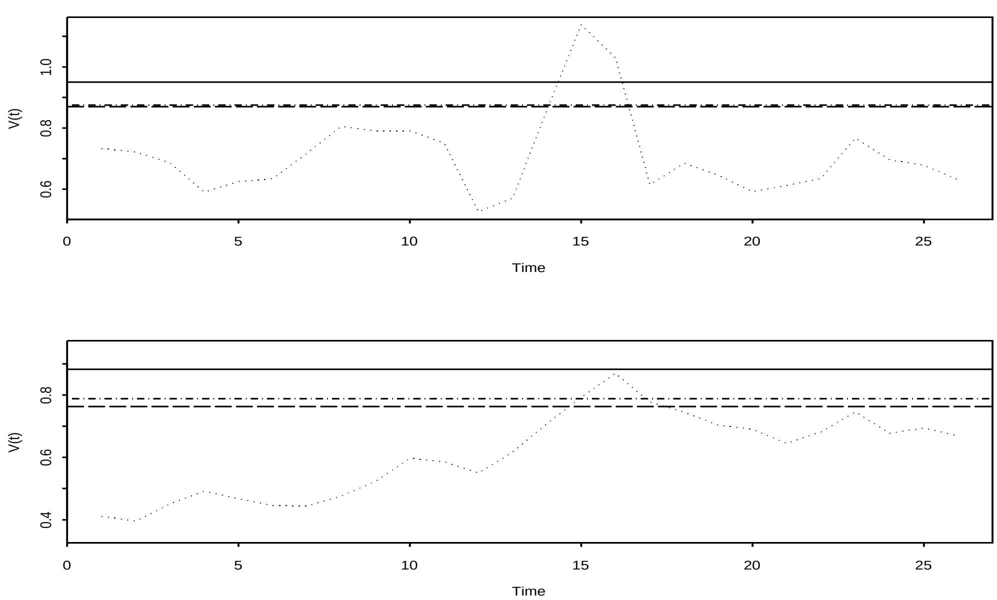

Figure 2: 95% Quantiles. Car insurance (top graph) and accident and disease insurance (bottom graph). (——) Independence case, (----) common ρ case, and (-· -·) different ρ_t case.

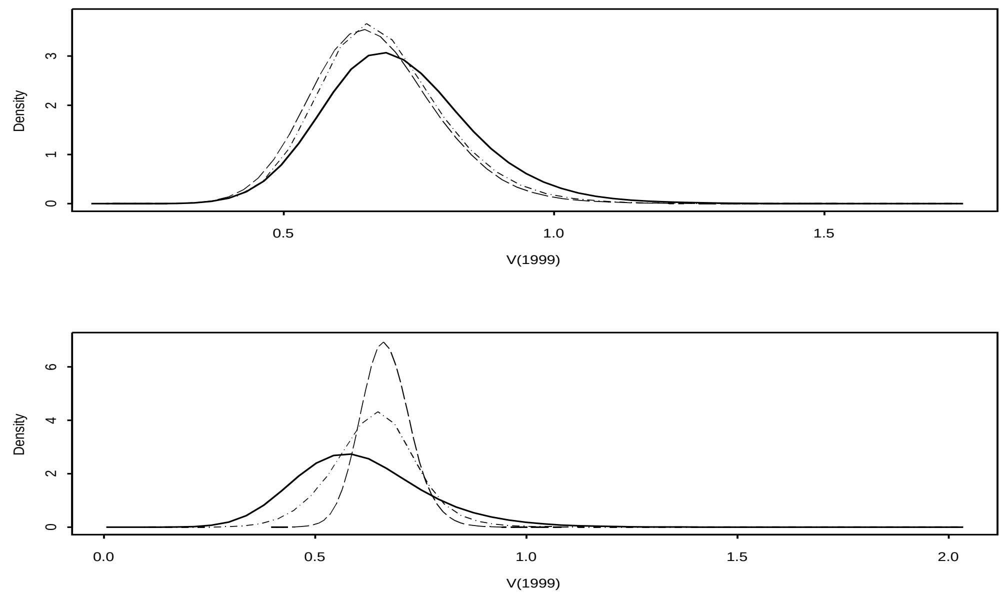

Figure 3: Predictive distribution of Y_{T+1}. Car insurance (top graph) and accident and disease insurance (bottom graph). (——) Independence case, (----) common ρ case, and (- · --·) different ρ_t case.

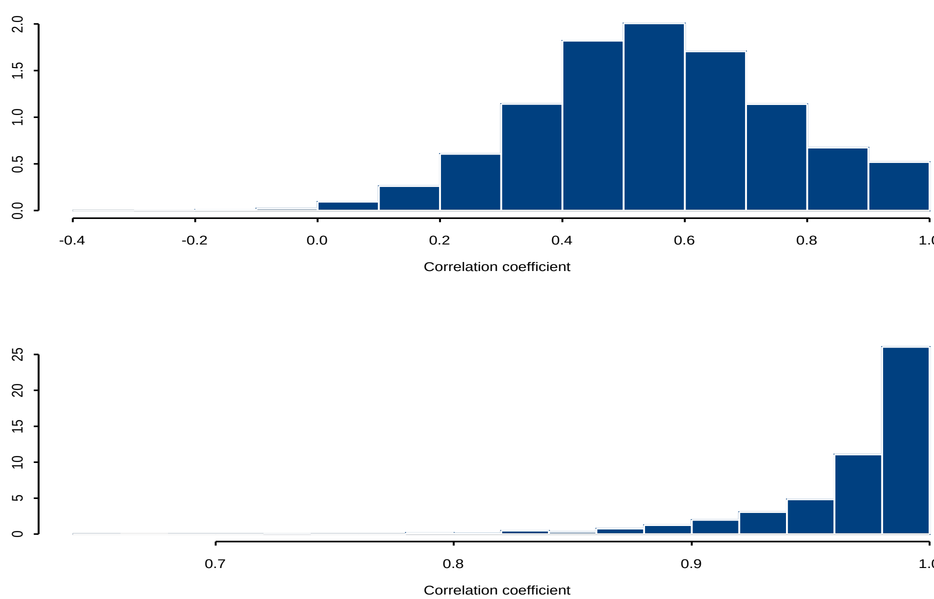

Figure 4: Histogram of the posterior distribution of the common ρ. Car insurance (top graph) and accident and disease insurance (bottom graph).

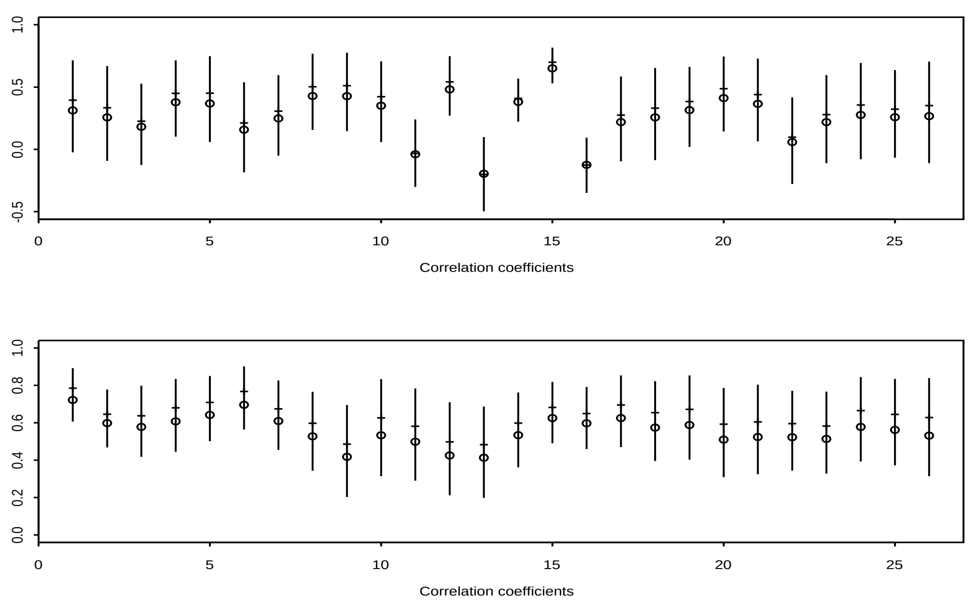

Figure 5: 95% Posterior bands of the different ρ_t's. Car insurance (top graph) and accident and disease insurance (bottom graph).