Implementing a Solvency II internal model: Bayesian stochastic reserving and parameter estimation

Marco Pirra[1], Salvatore Forte[2], and Matteo Ialenti[2]

[1] Department of Business Administration, University of Calabria
Ponte Bucci cubo 3 C - 87030, Rende (CS), Italy
Marco.Pirra@unical.it

[2] Department of Statistics, Sapienza University of Roma
Viale Regina Elena, 295 - Palazzina G - 00161, Roma, Italy
{Salvatore.Forte, Matteo.Ialenti}@uniroma1.it

Abstract. Often in non-life insurance claims reserves are the largest position on the liability side of the balance sheet. Therefore, the estimation of adequate claims reserves for a portfolio consisting of several lines of business is relevant for every non-life insurance company. The new solvency regulations require insurance companies to move to a market-consistent valuation of their liabilities (full balance sheet approach) and to prove the adequacy every year. Solvency II requires a view of the distribution of expected liabilities in one year: this implies considering the entire distribution of the profit/loss on reserves over a one year horizon for the proper assessment of the reserving risk. In the present paper we compare stochastic reserving models investigating what additional benefits a Bayesian approach can bring to a valuation. One of the most interesting is the skewness of actuarial parameters: we investigate prior distributions, which are a distinctly Bayesian feature, and analyze the most appropriate distribution shape or family for the parameters used in the models. Prior distributions allow actuaries to incorporate external information into their models, in a similar way to actuarial judgment being applied to traditional spreadsheet models. We investigate the ways in which prior distributions can be used, and their impact on the posterior distributions through the Metropolis-Hastings algorithm. Under new Solvency II developments insurance companies need to calculate a risk margin to cover possible shortfalls in their liability runoff. We derive an analytic formula for the risk margin which allows the comparison of the different proxies used in practice and we develop a flexible internal model that can be used for evaluating a specific risk profile. A case study on different liability datasets investigates the influence of the dimension on the results and gives a possible answer to some questions raised by the International Actuarial Association. Moreover, a backtesting process compares historical results to those produced by the current model in order to validate both the reasonableness and the implementation of the assumptions.

Keywords. Solvency II, internal models, Bayesian stochastic methods, reserve risk, uncertainty and ranges.

1 Introduction

The runoff of general insurance liabilities (outstanding loss liabilities) usually takes several years. Therefore, general insurance companies need to build appropriate reserves (provisions) for the runoff of the outstanding loss liabilities. These reserves need to be incessantly adjusted according to the latest information available. Under new solvency regulations, [6], general insurance companies have to protect against possible shortfalls in these reserves adjustments with risk bearing capital. In this spirit, this work provides a comprehensive discourse on multiperiod solvency considerations for a general insurance liability runoff and aims to give a possible answer to the questions raised by the International Actuarial Association, [10]. The discourse involves the description of the cost-of-capital approach in a multiperiod risk measure setting. In a cost-of-capital approach the insurance company needs to prove that it holds sufficient reserves firstly to pay for the insurance liabilities (claims reserves) and secondly to pay the costs of risk bearing capital (cost-of-capital margin or risk margin). Hence, at time 0, the insurer needs to hold risk-adjusted claims reserves that comprise best-estimate reserves for the outstanding loss liabilities and an additional margin for the coverage of the cash flow generated by the cost-of-capital loadings. Such risk-adjusted claims reserves are often called a market-consistent price for the runoff liabilities (in a marked-to-model approach), see e.g. [7], [8], [14]. Because the multiperiod cost-of-capital approach is rather involved, state-of-the-art solvency models consider a one-period measure together with a proxy for all later periods. Only high-quality internal models optimally reflecting the risk situation facing the company allow insurers to assess the level of risk capital required. This importantly involves measuring and evaluating reserve risk as a part of insurance risks. In literature there is a wide variety of methods for stochastic reserving such as the Mack method, [13], the Bootstrap method, [4], regression approaches, [3], Bayesian methods, [5], etc. All these approaches are based on an ultimo view, so that the uncertainty of full run-off of the liabilities is quantified. In contrast Solvency II requires the quantification of the one-year reserve risk. In addition the investment results, that have to be added to insurance results, are also based on a one-year view, which means that actually many internal models show an ultimo view for insurance results and the one-year view for investment results. So at the moment there is a discussion in academic literature and in insurance practice, how this one-year reserve risk can be quantified. This paper focuses on the Bayesian reserving methods and proposes a Bayesian model based on claim numbers and average cost per claim, following the approach otulined in [7], which can be applied in modelling reserve risk. Based on this approach we can quantify one-year risk capital and multi-year risk capital. The results of the method proposed are compared with the results of the approaches proposed by CEIOPS in [2].

2 Bayesian Methodology

Insurance is, by nature, a very uncertain subject. Insured events occur at random times and, particularly in the field of general insurance, the amounts of the claims are also random. The act of taking out insurance relieves the insured parties of some of the risk involved, passing it on to insurers, in return for a stable series of payments. The insurer must calculate the value of premium it should charge, which will be related to the total expenditure it is likely to have in fulfilling the conditions of the policies. In addition to this, the insurer must ensure it has sufficient funds, or reserves, in place to pay out claims when they occur. In order to do this, they need to learn about not only the average amount to be paid out in any one year (which would be sufficient to determine the basic premium amount), but also about the whole distribution of the aggregate claim for the year. In the general insurance market, insurers need to use the data gathered from previous years of experience to make predictions about future liabilities. The statistical analysis can be performed using Bayesian methodology, which is increasingly common in actuarial science. While some may question the appropriateness of prior experience in modeling, there are statisticians who believe that prior experience has a place in statistical analysis and should be formally recognized, even if it cannot be rigorously quantified. These statisticians are called Bayesians. In a Bayesian world, a modeler might start with some sense of how likely a parameter will take on a particular value from a prior model of the parameter, and then review the data available, changing the assessment of this likelihood to get a posterior model of the parameter. This is possible because of a fundamental Bayes' theorem, which can be summarized (see Klugman, [11], for an accurate description) introducing the following definitions:

- the prior distribution is a probability distribution over the space of possible parameter values. It is denoted $\pi(\theta)$ and represents our opinion concerning the relative chances that various values of θ are the true value;
- an improper prior distribution is one for which the probabilities (or pdf) are non-negative, but their sum (or integral) is infinite;
- the model distribution is the probability distribution for the data as collected, given a particular value for the parameter. Its pdf is denoted $f_{x|\theta}(x|\theta)$. Note that the expression is identical to the likelihood function, and may also be referred to as such;
- the joint distribution has pdf $f_{x,\Theta}(x|\theta)=f_{x|\Theta}(x|\theta)\pi(\theta)$;
- the marginal distribution of x has pdf $f_x(x)=\int f_{x|\Theta}(x|\theta)\pi(\theta)d\theta$;
- the posterior distribution is the conditional probability distribution of the parameters given the observed data, denoted as $\pi_{\theta|x}(\theta|x)$;

– the predictive distribution is the conditional probability distribution of a new observation y given the data x and is denoted $f_{Y|X}(y|x)$.

Given this background we can now state Bayes' theorem as follows. The posterior distribution can be computed as:

$$\pi_{\Theta|x}(\theta|x) = \frac{f_{x|\Theta}(x|\theta)\pi(\theta)}{\int f_{x|\Theta}(x|\theta)\pi(\theta)d\theta}, \tag{1}$$

while the predictive distribution can be computed as follows:

$$f_{Y|X}(y|x) = \int f_{Y|\Theta}(y|\theta)\pi_{\Theta|X}(\theta|x)d\theta \tag{2}$$

where $f_{Y|\Theta}(y|\theta)$ is the pdf of the new observation, given the parameter value. In both formulas the integrals are replaced by sums for a discrete distribution. Bayes' theorem, therefore, provides a way to let our prior assessment evolve with additional information, as well as to incorporate our stated uncertainty inherent in the estimation of the parameters directly into the assessment of possible future outcomes within the models framework. In this way Bayes' theorem gives a posterior model of the parameter, rather than an estimate of the parameter itself, as in the MLE. This principle allows for considerably more information, but also raises the question as to which single point on the model to use to represent the parameter. To address this issue, Bayesians may consider penalty functions, which involve selecting the value of the parameter, known as a Bayesian estimate of the parameter, that incurs the least expected penalty. One benefit of Bayesian estimates is that through a judicious choice of a penalty or loss function, one can reflect practical limitations and penalties inherent in misestimating a parameter. The importance of the Bayesian approach becomes apparent when we step back and look at actuarial and related problems from a broader perspective. At the heart of classical non-Bayesian statistical analysis is the concept of asymptotically, which suggests that a large enough number of observations from the same phenomenon will produce certain statistics with properties close to some convenient models; the more observations, generally, the closer the results will be to the ideal. The concept works well for repeatable experiments, such as the toss of a coin, but it requires a leap of faith when only a relatively small number of observations are available for an ever-changing environment. As noted above, Bayes' theorem does not require a sufficiently large number of observations and provides us with a useful model of the parameters, rather than the asymptotically normal model for MLE. However, Bayes' theorem does require a prior distribution of the parameter(s). And while it conveys how a model would evolve as more information becomes available, it does not provide a single point estimate for the parameters as does MLE. An analysts prior model may have a significant impact on the estimation of the parameters, particularly when empirical observed evidence is not plentiful. Thus, from a Bayesian point of view, the prior model is a powerful way for past experience to be brought to bear on a current problem. Indeed, Bayes' theorem gives a rigorous way to adapt that prior model for emerging facts. Although a prior model can be subjective, some studies, including Rebonato,[19] require it to be, in some sense, real and verifiable so that it can be tested in the marketplace.

3 Bayesian reserving models

Stochastic reserving methods have been a keen area of research for the global actuarial profession over the past decade (For reviews, see England and Verrall, [4] and Li [12]. One particular branch of stochastic methods are known as Bayesian methods. These methods can take a variety of forms, but share a common feature in that they take advantage of Bayesian statistical modelling techniques. Bayesian statistics is a branch of statistics that approaches statistical modelling from a slightly different perspective to classical statistics. Bayesian statistics makes use of two sources of information: the observed data (called the "likelihood"), and additional external information that may not necessarily be present in the observed data (this is called the "prior"). Both the likelihood and the prior are formulated as probability distributions. Bayesian statistics is in contrast to classical statistics, which is generally limited to using the observed data only. There are a number of benefits that Bayesian stochastic reserving models bring. The first benefit is that Bayesian modelling is flexible enough to build models that are similar to currently used reserving models. Bayesian models can easily incorporate GLM models, making Bayesian versions of a broad range of stochastic reserving models possible. Under certain conditions, the Bayesian version of a model will give the same central estimate as the spreadsheet version of the model. This removes much of the "black box" element that many stochastic reserving models suffer from, where the stochastic model gives a result that is not reconcilable to the result from a traditional method. Bayesian models can directly incorporate prior information, or information that is external to the observed data. This is in contrast to most stochastic reserving models, which cannot incorporate information that is external to the data being analysed. Bayesian models provide a formal framework for integrating actuarial judgment where an actuary does not consider the pure data alone to completely describe all of the information relevant to valuing the liabilities. In a sense, Bayesian models may be seen as a bridge between pure stochastic models and pure deterministic models. They allow actuaries to enhance and expand on the strengths of existing reserving approaches, without the need to start again from scratch with a purely stochastic model. Among the reasons why Bayesian models shall be preferred to other models are the results that can be obtained and the use of prior distributions. The results available are useful to understand the overall distribution of reserves, but also the distribution of each future payment (and thus potential values for each future payment). Having the full distribution of results makes it possible to use any percentile on the distribution. From a risk margins perspective, using a Bayesian model means you can use a single model to produce both the central estimate and risk margin. Bayesian models also give a distribution of all of the stochastic parameters in the model. You can look at any or all of these distributions, in order to check for things such as the reasonableness of assumptions, or what the key drivers of overall volatility are. In an actuarial context, the prior distribution allows you to incorporate information that is not in the data. Actuaries will already be familiar with this idea, as a significant part of any reserving exercise is the application of actuarial judgment. This can range from ignoring particular outliers from a dataset prior to model fitting, to directly adjusting development factors where there is reason to believe the future experience will differ from the historical observed experience, to selecting the prior ultimate loss ratio in a BF model. In a BF model, the model blends the actual experience as it emerges with the selected loss ratio, in a similar vein to a Bayesian model blending a prior distribution with the likelihood distribution. The external information could come from a range of sources, including:

- more extensive claims data (the particular model may be a subset of a broader class of business which may be partially useful in setting assumptions);
- analysis performed elsewhere, such as analysis done for pricing work;
- non-claims information, such as weather data (it would also be possible to incorporate weather data directly into the model);
- industry comparisons;
- judgment, particularly where things such as claims management processes or policy coverage change.

The impact on the final results of introducing a prior distribution depends on a number of factors. These include:

- the “strength” of the prior, that is to say the lower the variance of the prior is, the greater will be the impact on the results;
- the “strength” of the likelihood/data, that is to say the higher the sample variance/uncertainty of the likelihood, the greater the impact any prior will have on the results. The sample variance is a measure of how dispersed the data is;
- the number of data points, that is to say the bigger the dataset observed is, the lower the impact of the prior is. Any given prior will have a greater impact with only 5 data points compared to 500 data points.

There have been many papers that present a variety of Bayesian reserving models. For an easy to follow introduction to using WinBUGS to build Bayesian models, see Scollnik, [20]. For a selection of additional Bayesian models see Ntzoufras and Dellaportas, [18], Verrall, [21], [22], Meyers, [15], [16], [17].

4 The Bayesian Fisher Lange

In this paragraph a new Bayesian method based on claim numbers and average cost per claim is proposed by the authors. The stochastic method presented is an extension of the traditional deterministic one and an extension of the model proposed by the authors in GIIA, [7]. When the claim amounts paid or incurred are divided by the relevant number of claims, an average cost per claim results. This average cost can be projected, just as were the claim amounts themselves. Then, combined with a separate projection for the number of claims, it will yield the new estimate for the ultimate loss. The reserver can also examine the movement of the claim numbers and average costs as the accident years develop, and look for significant trends or discontinuities. A fuller view of the business can thus be obtained, perhaps leading to adjustment of the reserving figures, or showing where further investigation is needed. The method described in this paragraph is based on the definition of three variables that are crucial for the assessment of future payments, i.e.:

- the average cost per claim;
- the proportion of claims settled;
- the settlement speed.

Let $Y_{i,j}$ be the amount paid, properly deflated, $NP_{i,j}$ be the number of claims paid and $NR_{i,j}$ be the number of claim reserved, being i the accident year and j the development year. The average cost per claim is defined as follows:

$$AC_{i,j} = \frac{Y_{i,j}}{NP_{i,j}}, i + j \leq n + 1. \tag{3}$$

Assume the average costs per claim follow a gamma distribution with scale equal to θ_j^{AC} and shape equal to ω_j^{AC}, i.e.:

$$AC_{i,j} \sim \Gamma\left(\theta_j^{AC}, \omega_j^{AC}\right)^{1}. \tag{4}$$

In order to make the expected value of the average costs of development year j a stochastic variable, the parameter θ_j^{AC} is assumed to be drawn from a normal distribution as well:

$$\theta_j^{AC} \sim N\left(\mu.\theta_j^{AC}, \tau.\theta_j^{AC}\right)^{2}. \tag{5}$$

with mean $\mu.\theta_j^{AC}$ and variance $\tau.\theta_j^{AC}$. Distribution (4) represents the Bayesian Fisher Lange model, distribution (5) is a prior distribution , while $\mu.\theta_j^{AC}$,$\tau.\theta_j^{AC}$ and ω_j^{AC} are the hyperparameters of the model, expressed in the form of vectors of input.

Let $Rate_{i,j}$ represent the proportion of claims reported in the year i that are reserved in the development year j and will be paid in future different years, i.e.:

$$Rate_{i,j} = \frac{\sum_{k=j+1}^{n-i+1} NP_{i,k} + NR_{i,n-i+1}}{NR_{i,j}}, i + j \leq n. \tag{6}$$

This variable allows taking into account the possible re-opened claims and the closed without settlement ones as well. The probability model is similar to the

[1] In WinBUGS the gamma distribution with parameters a and b has the following moments: mean=a/b, variance=a/b^2.

[2] In WinBUGS the normal distribution with parameters a and b has the following moments: mean=a, variance=$1/b$.

previous one, (4). Assume that $Rate_{i,j}$ are drawn from a normal distribution:

$$Rate_{i,j} \sim N\left(\theta_j^{Rate}, \tau^{Rate}\right), \tag{7}$$

with the expected value θ_j^{Rate} that depends from the development year and is drawn from a normal distribution as well:

$$\theta_j^{Rate} \sim N\left(\mu.\theta_j^{Rate} = 1, \tau.\theta^{Rate}\right). \tag{8}$$

The hyperparameters τ^{Rate} and $\tau.\theta^{Rate}$ are expressed in the form of costant inputs.
The settlement speed $\nu_{i,j}$ represents the proportion of claims of the year i that are paid during the development year j:

$$\nu_{i,j} = \frac{NP_{i,j}}{\sum_{k=2}^{n-i+1} NP_{i,k} + NR_{i,n-i+1}}, i = 1, \ldots n-1; j = 2, \ldots n-i+1. \tag{9}$$

The probability model follows the same structure as for (4) and (7). Assume the settlement speeds $\nu_{i,j}$ follow a normal distribution:

$$\nu_{i,j} \sim N\left(\theta_j^{\nu}, \omega_j^{\nu}\right), \tag{10}$$

with the expected value being drawn from the following normal distribution:

$$\theta_j^{\nu} \sim N\left(\mu.\theta_j^{\nu}, \tau.\theta_j^{\nu}\right). \tag{11}$$

The hyperparameters ω_j^{ν}, $\mu.\theta_j^{\nu}$ and $\tau.\theta_j^{\nu}$ are expressed in the form of vectors of input.
The variables $Rate_{i,j}$ and $\nu_{i,j}$ are used for the recursive estimation of the number of future claims paid, i.e.:

$$NP_{i,j} = NR_{i,n-i+1} \cdot Rate_{i,n-i+1} \cdot \frac{\nu_{i,j}}{\sum_{k=n-i+1}^{n} \nu_{i,k}}, i+j > n+1. \tag{12}$$

The future payments $Y_{i,j}$ are obtained multiplying the number of future claims paid and the relative average cost conveniently projected considering inflation, i.e.:

$$Y_{i,j} = NP_{i,j} \cdot AC_{i,j} \cdot (1+ir)^{(i+j-n-1)}, i+j > n+1, \tag{13}$$

where ir represents the annual inflation rate for simplicity assumed constant for the entire reserve run-off horizon. Distribution (4), distribution (7) and distribution (10) represent the Bayesian Fisher Lange model, while distribution (5), distribution (8) and distribution (11) are prior distributions. The values of the input parameters are determined on actuarial knowledge and market experience: the information relates the costs and the settlement policies a company should keep traces on.

5 Solvency II and backtesting

In the Solvency II Directive framework, [6], the fair value of the technical provisions is defined as follows: “The value of technical provisions shall be equal to the sum of a best estimate and a risk margin; the best estimate shall be equal to the probability-weighted average of future cash flows, taking account of the time value of money (expected present value of future cash-flows), using the relevant risk-free interest rate term structure; the risk margin shall be such as to ensure that the value of the technical provisions is equivalent to the amount insurance and reinsurance undertakings would be expected to require in order to take over and meet the insurance and reinsurance obligations”. With regard to technical provisions, the act also requires insurers to have processes and procedures in order to insure that best estimates and the assumptions underlying the calculation of the best estimates are regularly compared against experience. In case the comparison identifies systematic deviations between the experience and the best estimate, the insurer shall make appropriate adjustments to the actuarial methods and/or the assumptions being made. The authors analysed the fair valuation of the technical provisions in a previous paper, [7].
As far as the backtesting procedures is concerned, firms that use VaR as a risk disclosure or risk management tool are facing growing pressure from internal and external parties such as senior management, regulators, auditors, investors, creditors, and credit rating agencies to provide estimates of the accuracy of the risk models being used. Users of VaR realized early that they must carry out a cost-benefit analysis with respect to the VaR implementation. A wide range of simplifying assumptions is usually used in VaR models (distributions of returns, historical data window defining the range of possible outcomes, etc.), and as the number of assumptions grows, the accuracy of the VaR estimates tends to decrease. As the use of VaR extends from pure risk measurement to risk control in areas such as VaR-based Stress Testing and capital allocation, it is essential that the risk numbers provide accurate information, and that someone in the organization is accountable for producing the best possible risk estimates. In order to ensure the accuracy of the forecasted risk numbers, risk managers should regularly backtest the risk models being used, and evaluate alternative models if the results are not entirely satisfactory. VaR models provide a framework to measure risk, and if a particular model does not perform its intended task properly, it should be refined or replaced, and the risk measurement process should continue. The traditional excuse given by many risk managers is that “VaR models only measure risk in normal market conditions” or “VaR models make too many wrong assumptions about market or portfolio behavior” or “VaR models are useless” should no longer be taken seriously, and risk managers should be accountable to implement the best possible framework to measure risk, even if it involves introducing subjective judgment into the risk calculations. It is always better to be approximately right than exactly wrong. How can the accuracy and performance of a VaR model be assessed? In order to answer this question, we first need to define what we mean by “accuracy”. By accuracy, we could mean not only how well the model measures a particular percentile of or the entire

profit-and-loss distribution, but also how well the model predicts the size and frequency of losses. Many standard backtests of VaR models compare the actual portfolio losses for a given horizon vs. the estimated VaR numbers. In its simplest form, the backtesting procedure consists of calculating the number or percentage of times that the actual portfolio returns fall outside the VaR estimate, and comparing that number to the confidence level used. For example, if the confidence level were 95%, we would expect portfolio returns to exceed the VaR numbers on about 5% of the days. Backtesting can be as much an art as a science. It is important to incorporate rigorous statistical tests with other visual and qualitative ones. The approach followed in the work in order to understand if the reserve predicted by the model matches the reserve held by the insurance company is to compare prior year development to model predictions, that is to say compare the probability distribution and the expected value of the first diagonal of the run-off triangle obtained by the exclusion of the last generation and the actual paid value written in the balance sheet. This type of back-testing should be a significant part of the validation process, although the test should not be limited to this.

6 Case study

This paragraphs shows the results on a sample case study. The initial data set is represented by the run-off triangle of incremental payments of an insurance company operating in the general liability LoB (Table 1), the triangle of the number of claims paid (Table 2) and the triangle of the number of claims reserved (Table 3). The aim is the analysis of the effects of the prior distributions on the outstanding claim reserve and its variability measure (coefficient of variation). The input values of the parameters of the prior distributions are reported in Table 4, Table 5 and Table 6; the choice is to set them in order to have a specific coefficient of variation, i.e. for $Rate_{i,j}$ a CV=10%, for θ_j^{Rate} a CV=50%, for $\nu_{i,j}$ a CV=5%, for θ_j^{ν} a CV=20%, for $AC_{i,j}$ a CV=10%, for θ_j^{AC} a CV=10%. Table 7 reports the Bayesian Fisher Lange outstanding claim reserve assessed through the model described in paragraph 4 on an undiscounted basis, while Table 8 reports the values of the discounted[3] Bayesian Fisher Lange outstanding claim reserve, as Solvency II defines the best estimate of the claim reserve as the expected present value of future cash-flows. For comparison purposes the outstanding claim reserve is determined through another stochastic reserving method, the Over Dispersed Poisson with Bootstrapping, described in [4]. Table 9 reports the ODP oustanding claim reserve on an undiscounted basis while Table 10 reports the discounted[3] ODP outstanding claim reserve. The expected value is different as the deterministic methodology underlying the stochastic model differs as well, but the gap (around 4%) is tolerable. As far as the variability is concerned the Bayesian Fisher Lange leads to higher values (6%-6.2%

[3] The discounted values are obtained considering the Interest Rate Term Structure 2007 given by CEIOPS in [2]. The interest rates do not consider any illiquidity premuim.

vs. 4.5%-4.3% as it has a double level of stochasticity. The ODP has a very high variability on older generations and lower values on recent generations: this is a typical characteristic of the ODP model and in general of the chain ladder methods. The Bayesian Fisher Lange gives homogeneous values on different generations due to its bayesian nature and construction: older generations have a restricted number of datapoints to compute the posterior distribution (high variability) and a restricted number of future values to project (low variability as they are influenced only by recent average costs and settlement speeds); recent generations have a larger number of datapoints to compute the posterior distribution (lower variability) but many more values to project (which are influenced by more average cost and settlement speed values).
Table 11 reports the values of the stability analysis on the parameter θ_j^{Rate}, carried out varying the input parameters $\mu.\theta_j^{Rate}$ and $\tau.\theta^{Rate}$ in the θ_j^{Rate} prior distribution. The results show a good stability of the model, both in terms of expected value and standard deviation of the outstanding claim reserve: even if the input parameter in the prior distribution changes significantly (passing from a value of 1 to a value of 1.5) the expected value remains pretty much the same. These results demonstrate the optimal convergence of the model to the posterior distribution, irrespective of the initial values.
Table 12 reports the values of the stability analysis on the parameter θ_j^{ν}, carried out varying the input parameters $\mu.\theta_j^{\nu}$ and $\tau.\theta_j^{\nu}$ in the θ_j^{ν} prior distribution. The results show a good stability of the model, both in terms of expected value and standard deviation of the outstanding claim reserve. The models tends to produce identical values varying the mean of the settlement speed expected value, $\mu.\theta_j^{\nu}$, while seems to be slightly more sensitive varying the coefficient of variation of the settlement speed expected value (that has a value of 20% in the basis scenario).
Table 13 reports the values of the stability analysis on the parameter θ_j^{AC}, carried out varying the input parameters $\mu.\theta_j^{AC}$ and $\tau.\theta^{AC}$ in the θ_j^{AC} prior distribution. The results show a good stability of the model both in terms of expected value and standard deviation of the outstanding claim reserve. Although the differences seem relevant, the model leads to a good convergence to the posterior distribution: even reducing the prior average costs of a percentage equal to 50% the reduction of the expected value of the outstanding claim reserve is limited to a 12%.
Figure 1, Figure 2 and Figure 3 show the prior distribution of some selected parameters while Figure 4, Figure 5 and Figure 6 their posterior distribution. The comparison of the figures demonstrates the optimal convergence of the model to the posterior distributions and confirms the considerations outlined in the stability analyses: for example the $AC_{12,9}$ passes from a prior expected value equal to 25k to a posterior expected value equal to 29k; the $Rate_{12,1}$ passes from a prior expected value equal to 100% to a posterior expected value equal to 91.71%; $\nu_{12,2}$ passes from a prior expected value equal to 75% to a posterior expected value equal to 80.1%.
Table 14, Table 15, Table 16 report the values of some sensitivities: the analyses

are carried out varying the coefficients of variation of the model in $Rate_{i,j}$, $\nu_{i,j}$, $AC_{i,j}$.
As outlined in previous paragraphs, in the Solvency II Directive framework the fair value of the technical provisions shall be equal to the sum of a best estimate and a risk margin. The different columns of Table 17 represent the different way of calculations adopted for the best estimate assessment:

- I (method) : Merz-Wuthrich, see [13];
- II (method) : ODP, see [4];
- III (method) : FL Bayes, see paragraph 4.

Two different approaches for the assessment of the reserve risk capital and the risk margin are compared[4]:

- the "Internal Model One Year Horizon" approach, [7];
- the "Standard QIS5" approach, [2];

The results show that the values both of the best estimate and the reserve risk capital are influenced by the model chosen: appropriate selection criteria and adequate backtesting procedures play a crucial role in the assessment. As far as the QIS5 is concerned the reserve risk capital expressed as a percentage of the best estimate can lead to a double level of prudential estimation (the higher the best estimate, the higher the reserve risk capital) or imprudential estimation (the lower the best estimate, the lower the riserve risk capital). The market value of σ seems too high and it is not related to the dmension of the portfolio: some studies (see e.g. [8]) demonstrate that the variability gets lower as the porftolio dimension gets higher. As already outlined in [8] the QIS 5 scheme leads to a double counting of the risk margin component.
As stated in previous paragraphs, Solvency II requires insurers to have processes and procedures in order to insure that best estimates, and the assumptions underlying the calculation of the best estimates, are regularly compared against experience. The approach followed in the work in order to understand if the reserve predicted by the model matches the reserve held by the insurance company is to compare prior year development to model predictions, that is to say compare the probability distribution and the expected value of the first diagonal of the run-off triangle obtained by the exclusion of the last generation and the actual paid value written in the balance sheet (in the case study the observed value is equal to 118,437,500 euros). Table 18 reports the results of the backtesting procedure: considering the expected value, the ODP model gets to a closer result, while in terms of p-value the Bayesian Fisher Lange is to be preferred. Figure 7 and Figure 8 complete the procedure through a graphical analysis. The most appropriate backtesting methodology and the choice among different stochastic reserving methods are still being debated in the actuarial literature.

[4] See [1], [7],[8] for a more detailed description.

7 Conclusions

The results of the case study presented seem to lead to the following conclusions:

- the prior distribution allows you to incorporate information that is not in the data. Actuaries will already be familiar with this idea, as a significant part of any reserving exercise is the application of actuarial judgment;
- the bayesian methodology combined with the Fisher Lange model has the advantage of explicitly taking into account the settlement and reserving policies of the insurer;
- the impact on the final results of introducing a prior distribution can be adequately explained;
- the assessment of the best estimate is much more influenced by the deterministic methodology underlying the stochastic model rather than by the probabilistic structure of the stochastic model itself;
- the variability measure (sigma) and the reserve risk capital are significantly affected by the probabilistic structure of the model and by the insurer dimensions;
- the QIS5 standard formula states that the risk capital is a percentage of the best estimate, different for each LoB. This approach could penalize prudential insurers and could lead the management to select the methodology for the claim reserve assessment that gives the lower result;
- the use of a unique sigma for all the insurance companies could lead to an overestimation both of the risk capital and the risk margin. A possible solution could be an entity specific sigma to be combined with the market wide sigma, as proposed with the "Undertaking-Specific Parameter" approach in [2]; yet USP are not easy to compute as they require input data that are not always available (net of reinsurance triangles) and have to be validated by the financial authority. A simpler solution could be the consideration of an adequate size factor;
- the choice of the internal model for the reserve risk assessment has a great importance; that is the reason a set of validation criteria should be defined and verified through a backtesting analysis.

The results presented and the conclusions exposed depend significantly on the datasets considered and on the insurance companies analyzed; the intention is to apply the methodologies to other insurers and verify the possibility to extend the conclusions to other case studies.

References

1. AISAM-ACME, 2007, Study on non-life long tail liabilities. Reserve risk and risk margin assessment under Solvency II. Available on www.amice-eu.org.
2. CEIOPS, 2010, QIS5 Technical Specifications. Available on www.ceiops.org.
3. Christofides S., 1990, Regression Models based on log-incremental payments. Claims Reserving Manual,Vol 2, Institute of Actuaries.

4. England P., Verrall R, 2002, Stochastic Claims Reserving in General Insurance. British Actuarial Journal 8, III, 443-544.
5. England P., Verrall R., 2006, Predictive Distributions of Outstanding Liabilities in General Insurance. Annals of Actuarial Science: 1, 221-270.
6. European Commission, 2009, Solvency II directive (2009/138/EC). Available on ec.europa.eu.
7. Forte S., Ialenti M., Pirra M., 2008, Bayesian Internal Models for the Reserve Risk Assessment, Giornale dell' Istituto Italiano degli Attuari, Volume LXXI N.1, 39-58.
8. Forte S., Ialenti M., Pirra M., 2010, A reserve risk model for a non-life insurance company. Mathematical and Statistical Methods for Actuarial Sciences and Finance. Springer.
9. IAIS, 2007, Guidance paper on the structure of regulatory capital requirements, Fort Lauderdale.
10. International Actuarial Association IAA, 2009, Measurement of liabilities for insurance contracts: current estimates and risk margins. Available on www.actuaries.org.
11. Klugman S. A., Panjer H. H., Willmot G. E., 2008, Loss Models: From Data to Decisions, 3rd Edition, Wiley and Sons Edition.
12. Li J., 2006, Comparison of Stochastic Reserving Methods. Australian Actuarial Journal, volume 12, number 4.
13. Mack T., 1993, Distribution-free Calculation of the Standard Error of Chain Ladder Reserve Estimates. ASTIN Bulletin 23, 214-225.
14. Merz M., Wuthrich M., 2008, Modelling The Claims Development Result For Solvency Purposes. Casualty Actuarial Society E-Forum.
15. Meyers G., 2007, Thinking Outside the Triangle. Paper presented to the 37th ASTIN Colloquium, Florida.
16. Meyers G., 2007, Estimating Predictive Distributions for Loss Reserve Models. Variance volume 1, issue 2.
17. Meyers G., 2009, Stochastic Loss Reserving with the Collective Risk Model. Variance, volume 3, issue 2.
18. Ntzoufras I., Dellaportas P., 2002, Bayesian modelling of outstanding liabilities incorporating claim count uncertainty. North American Actuarial Journal, volume 6 issue 1.
19. Rebonato R., 2007, Plight of the Fortune Tellers: Why We Need to Manage Financial Risk Differently: Princeton University Press.
20. Scollnik D. P. M.,2004, Bayesian Reserving Models Inspired by Chain Ladder methods and implemented using WINBUGS, Actuarial Research Clearing House 2004 issue 2.
21. Verrall R., 2004, A Bayesian Generalized Linear Model for the Bornhuetter Ferguson Method of Claims Reserving. North American Actuarial Journal, volume 8, number 3.
22. Verrall R., 2007, Obtaining Predictive Distributions for Reserves Which Incorporate Expert Opinion. Variance, volume 1, issue 1.
23. Wuthrich M.V., Buhlmann H., Furrer H., 2008, Market-Consistent Actuarial Valuation. Springer.

Table 1. Incremental Payments Triangle (Euro thousands)

Y[i,j]	Development year											
Accident year	0	1	2	3	4	5	6	7	8	9	10	11
1996	35,591	36,597	15,613	6,463	3,442	2,982	1,701	1,581	1,209	1,108	386	942
1997	39,987	45,166	16,834	7,368	3,636	3,061	1,181	1,378	951	606	603	
1998	47,252	50,189	16,222	7,576	3,796	1,613	1,596	1,452	1,163	732		
1999	50,556	55,657	19,246	7,026	3,303	2,513	2,704	1,513	1,124			
2000	55,178	56,896	19,207	6,881	3,209	3,666	1,651	1,438				
2001	62,901	60,083	22,337	8,827	4,951	3,441	2,867					
2002	62,058	62,522	24,496	12,592	7,221	4,174						
2003	58,046	62,151	26,134	10,286	5,926							
2004	60,402	61,726	22,913	11,074								
2005	65,771	63,291	23,288									
2006	73,282	67,212										
2007	75,484											

Table 2. Number of Claims Paid

NP[i,j] Accident year	Development year 0	1	2	3	4	5	6	7	8	9	10	11
1996	43,047	17,251	1,992	716	354	196	136	75	44	45	29	30
1997	44,350	17,154	1,882	691	267	174	70	61	40	31	26	
1998	46,261	17,281	1,915	551	249	96	64	42	27	27		
1999	46,304	17,045	1,835	631	211	106	85	56	39			
2000	46,067	16,776	1,961	534	234	140	107	59				
2001	48,970	15,757	1,996	705	347	225	200					
2002	46,871	15,359	2,577	930	495	365						
2003	42,741	15,312	2,429	957	375							
2004	39,141	13,435	2,391	805								
2005	37,952	12,652	2,020									
2006	38,402	13,857										
2007	38,244											

Table 3. Number of Claims Reserved

NR[i,j] Accident year	Development year 0	1	2	3	4	5	6	7	8	9	10	11
1996	24,391	5,054	1,724	915	562	367	231	162	129	95	74	51
1997	23,499	4,144	1,399	667	404	237	191	137	106	79	60	
1998	22,630	3,574	1,054	500	251	180	127	97	92	70		
1999	21,962	3,421	1,100	500	319	229	160	117	89			
2000	21,515	3,309	1,092	599	404	276	182	134				
2001	20,975	3,969	1,644	954	601	376	187					
2002	21,772	5,501	2,381	1,326	775	412						
2003	22,149	5,412	2,362	1,264	862							
2004	19,277	4,846	2,265	1,319								
2005	17,664	4,695	2,375									
2006	19,357	4,920										
2007	18,979											

Table 4. Input Parameters $Rate_{i,j}$

Parameter	Value
$\tau.\theta^{Rate}$	4
τ^{Rate}	100%

Table 5. Input Parameters $\nu_{i,j}$

j	$\mu.\theta_j^{\nu}$	$\tau.\theta_j^{\nu}$	ω_j^{ν}
0	1.000	25	400
1	0.750	44	711
2	0.100	2,500	40,000
3	0.030	27,778	444,444
4	0.015	111,111	1,777,778
5	0.010	250,000	4,000,000
6	0.005	1,000,000	16,000,000
7	0.003	2,777,778	44,444,444
8	0.002	6,250,000	100,000,000
9	0.001	25,000,000	400,000,000
10	0.001	25,000,000	400,000,000
11	0.001	25,000,000	400,000,000

Table 6. Input Parameters $AC_{i,j}$

j	$\mu.\theta_j^{AC}$	$\tau.\theta_j^{AC}$	ω_j^{AC}	E[Costs]
0	100	0.01	0.05000	2,000
1	100	0.01	0.02000	5,000
2	100	0.01	0.01000	10,000
3	100	0.01	0.00769	13,000
4	100	0.01	0.00667	15,000
5	100	0.01	0.00500	20,000
6	100	0.01	0.00500	20,000
7	100	0.01	0.00400	25,000
8	100	0.01	0.00400	25,000
9	100	0.01	0.00400	25,000
10	100	0.01	0.00400	25,000
11	100	0.01	0.00400	25,000

Table 7. Bayesian Fisher Lange Reserve Undiscounted (Euro thousands)

Year	Mean	StDev	CV
1997	1,907	329	17.2%
1998	2,047	287	14.0%
1999	2,640	329	12.5%
2000	4,515	499	11.0%
2001	6,334	681	10.8%
2002	12,700	1,337	10.5%
2003	23,950	2,539	10.6%
2004	30,820	3,350	10.9%
2005	43,960	5,053	11.5%
2006	61,910	8,056	13.0%
2007	108,700	14,460	13.3%
Total	299,400	18,090	6.0%

Table 8. Bayesian Fisher Lange Reserve Discounted (Euro thousands)

Year	Mean	StDev	CV
1997	1,821	314	17.2%
1998	1,902	266	14.0%
1999	2,403	299	12.4%
2000	4,042	446	11.0%
2001	5,595	602	10.8%
2002	11,080	1,171	10.6%
2003	20,790	2,217	10.7%
2004	26,650	2,916	10.9%
2005	38,320	4,446	11.6%
2006	54,990	7,219	13.1%
2007	98,620	13,280	13.5%
Total	266,200	16,400	6.2%

Table 9. ODP Reserve Undiscounted (Euro thousands)

Year	Mean	StDev	CV
997	1,070	566	52.9%
1998	1,748	707	40.4%
1999	2,902	883	30.4%
2000	4,347	1,029	23.7%
2001	6,828	1,276	18.7%
2002	9,819	1,508	15.4%
2003	13,143	1,676	12.8%
2004	18,236	1,951	10.7%
2005	29,375	2,472	8.4%
2006	58,779	3,727	6.3%
2007	140,472	7,987	5.7%
Total	286,719	12,827	4.5%

Table 10. ODP Reserve Discounted (Euro thousands)

Year	Mean	StDev	CV
1997	1,022	540	52.9%
1998	1,623	654	40.3%
1999	2,645	798	30.2%
2000	3,902	912	23.4%
2001	6,049	1,112	18.4%
2002	8,597	1,296	15.1%
2003	11,445	1,433	12.5%
2004	15,821	1,664	10.5%
2005	25,676	2,130	8.3%
2006	52,305	3,285	6.3%
2007	127,570	7,237	5.7%
Total	256,653	10,936	4.3%

Table 11. Parameter θ_j^{Rate}: Oustanding Claim Reserve Stability (Euro thousands)

Scenario	Mean OCR	CV OCR	ΔMean
$\mu.\theta_j^{Rate}$			
Base	299,400	6.04%	
+25%	299,800	6.04%	0.13%
+50%	300,200	6.03%	0.27%
-25%	299,100	6.05%	-0.10%
-50%	298,700	6.05%	-0.23%
$\tau.\theta^{Rate}$			
Base	299,400	6.04%	
+25%	299,400	6.04%	0.00%
+50%	299,400	6.04%	0.00%
-25%	299,500	6.04%	0.03%
-50%	299,500	6.04%	0.03%

Table 12. Parameter θ_j^{ν}: Oustanding Claim Reserve Stability (Euro thousands)

Scenario	Mean OCR	CV OCR	ΔMean
$\mu.\theta_j^{\nu}$			
Base	299,400	6.04%	
+25%	299,400	6.04%	0.00%
+50%	299,400	6.04%	0.00%
-25%	299,500	6.04%	0.03%
-50%	299,500	6.04%	0.03%
$\tau.\theta_j^{\nu}$			
Base	299,400	6.04%	
+25%	297,900	6.05%	-0.50%
+50%	297,100	6.06%	-0.77%
-25%	303,000	6.03%	1.20%
-50%	314,200	5.99%	4.94%

Table 13. Parameter θ_j^{AC}: Oustanding Claim Reserve Stability (Euro thousands)

Scenario	Mean OCR	CV OCR	ΔMean
$\mu.\theta_j^{AC}$			
Base	299,400	6.04%	
+25%	309,600	6.18%	3.41%
+50%	317,200	6.28%	5.95%
-25%	285,300	5.95%	-4.71%
-50%	263,000	5.86%	-12.16%
$\tau.\theta_j^{AC}$			
Base	299,400	6.04%	
+25%	299,600	6.07%	0.07%
+50%	299,800	6.04%	0.13%
-25%	299,500	6.03%	0.03%
-50%	300,400	6.09%	0.33%

Table 14. Parameter $Rate_{i,j}$: Oustanding Claim Reserve Sensitivity (Euro thousands)

Scenario	Mean OCR	CV OCR	ΔMean	ΔCV
Base CV = 10%				
Base	299,400	6.04%		
+25%	299,500	7.21%	0.0%	19.25%
+50%	299,500	8.41%	0.0%	39.26%
-25%	299,400	4.95%	0.0%	-18.08%
-50%	299,400	3.99%	0.0%	-33.94%

Table 15. Parameter $\nu_{i,j}$: Oustanding Claim Reserve Sensitivity (Euro thousands)

Scenario	Mean OCR	CV OCR	ΔMean	ΔCV
Base CV = 5%				
Base	299,400	6.04%		
+25%	302,000	6.04%	0.9%	-0.09%
+50%	305,300	6.03%	2.0%	-0.14%
-25%	297,500	6.05%	-0.6%	0.08%
-50%	296,200	6.05%	-1.1%	0.13%

Table 16. Parameter $AC_{i,j}$: Oustanding Claim Reserve Sensitivity (Euro thousands)

Scenario	Mean OCR	CV OCR	ΔMean	ΔCV
Base CV = 10%				
Base	299,400	6.04%		
+25%	299,700	6.48%	0.1%	7.24%
+50%	300,900	6.89%	0.5%	13.97%
-25%	298,800	5.69%	-0.2%	-5.84%
-50%	298,700	5.45%	-0.2%	-9.79%

Table 17. Solvency II (Euro thousands)

Values/Methods	I	II	III
Internal Model One Year Horizon			
Best Estimate	256,585	256,653	266,200
Risk Margin (%BE)	1.01%	1.20%	2.04%
Reserve Risk Capital (%BE)	6.30%	7.94%	14.13%
σ (1 year)	2.75%	3.44%	6.04%
QIS5 Market Wide			
Best Estimate	256,585	256,653	266,200
Risk Margin (%BE)	4.33%	4.33%	4.59%
Reserve Risk Capital (%BE)	31.85%	31.85%	31.85%
σ (1 year)	11.00%	11.00%	11.00%

Table 18. Backtesting (Euro thousands)

Methods	E[1st Diag]	p-value
BFL	125,919	25.80%
ODP	125,536	11.70%

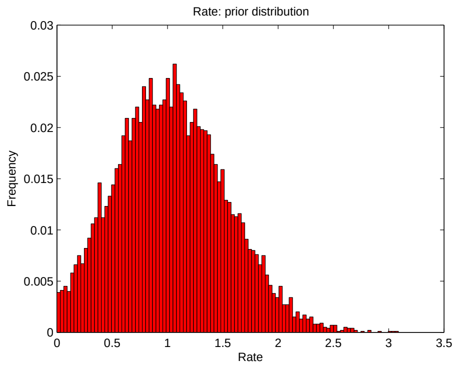

Fig. 1. *Rate*: Prior Distribution

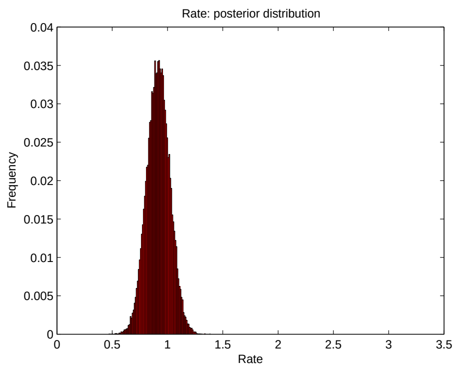

Fig. 2. *Rate*: Posterior Distribution

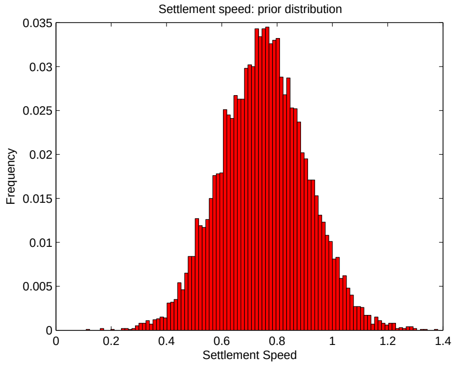

Fig. 3. ν: Prior Distribution

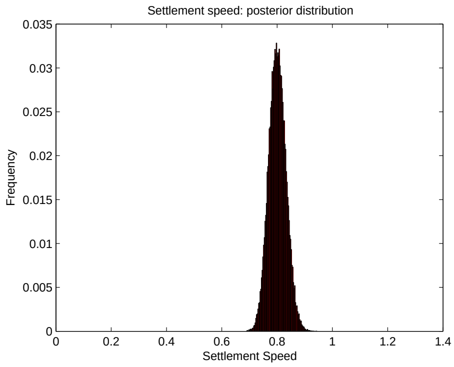

Fig. 4. ν: Posterior Distribution

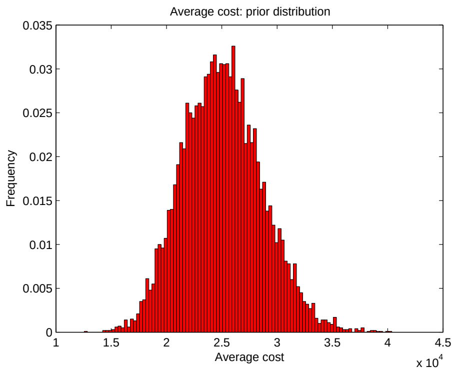

Fig. 5. AC: Prior Distribution

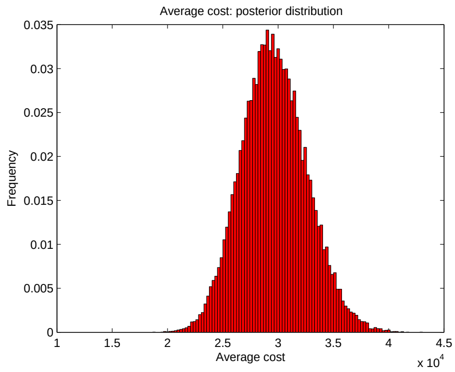

Fig. 6. *AC*: Posterior Distribution

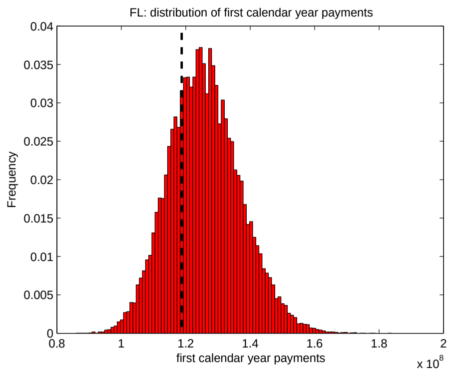

Fig. 7. Backtesting Fisher Lange

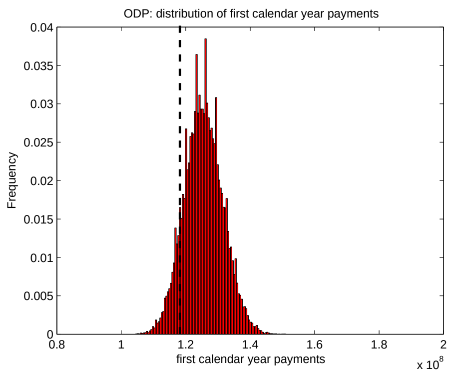

Fig. 8. Backtesting ODP