Bootstrapping the Cape Cod method: A working paper

By Nikita Harilal, Karen Muyengwa, Kovlin Perumal, Ronald Richman, Itumeleng Sefako, Rowald van der Walt

Presented at the Actuarial Society of South Africa's 2022 Convention
Cape Town 25–28 October 2022

ABSTRACT

Methods for estimating the reserve risk of incurred but not reported provisions are usually based on the assumption that these provisions were calculated using the chain ladder method. Thus, in practice, these methods relate to various ways of quantifying the prediction error of the chain ladder estimates, often the bootstrap procedures of England and Verrall (1999) or England and Verrall (2006) for the chain ladder method. On the other hand, IBNR estimates are often quantified using exposure-based methods such as the Bornhuetter-Ferguson or Cape Cod methods. In this working paper, we provide an initial view of how bootstrapping can be applied to the Cape Cod method. We define bootstrap procedures for the Cape Cod method, show how these can be applied and how the results compare to more established methods used in practice. We consider which of the methods are more appropriate for use in reserve risk estimation under Solvency II and accounting estimates in the context of IFRS 17, with a focus on stability and realism of the results. Finally, we provide `R` code to reproduce these methods on representative data.

KEYWORDS

Incurred But Not Reported (IBNR), bootstrapping, Cape Cod, chain ladder, Bornhuetter-Ferguson, reserving, IFRS 17, uncertainty estimation, reserve risk, model selection, premium risk

CONTACT DETAILS
Nikita Harilal, Old Mutual Insure; Email: nikita.harilal@ominsure.co.za
Karen Muyengwa, Old Mutual Insure, karen.muyengwa@ominsure.co.za
Kovlin Perumal, Old Mutual Insure, kovlin.perumal@ominsure.co.za
Ronald Richman, Old Mutual Insure, ronald.richman@ominsure.co.za
Itumeleng Sefako, Old Mutual Insure, itumeleng.sefako@ominsure.co.za
Rowald van der Walt Old Mutual Insure, rowald.vanderwalt@ominsure.co.za

1. INTRODUCTION

A well-known problem in the non-life insurance industry is the use of different methods for calculating the best estimate Incurred But Not Reported (IBNR) reserves of non-life[1] insurers and for estimating the risk that these reserves will be insufficient – see Dal Moro and Lo (2014) for a discussion. When calculating IBNR reserves, actuaries often use a combination of the chain ladder (CL), Bornhuetter-Ferguson (BF) (Bornhuetter & Ferguson, 1972) and Cape Cod (CC) (Bühlmann & Straub, 1983) methods, as well as other more specialised methods (Taylor, 2012; Mack, 2002). These methods use aggregated past claims data to project claims that have already occurred but will be reported in the future; in this work we focus only on these methods for aggregated data and do not consider individual claims reserving techniques which are not popular (yet) in practice. Besides these estimated reserves, under modern solvency regimes (e.g., Solvency II in Europe and Solvency Assessment and Management (SAM) in South Africa (Prudential Authority, 2020) and financial accounting standards (e.g., IFRS 17 (IASB, 2020)), insurers also need to hold risk capital or risk margins respectively, to cover cases when the (best-estimate) reserves turn out to be insufficient to meet the emerging claims. Usually, these capital or margins requirements are quantified with reference to the uncertainty of the IBNR reserve estimation process, considering the potential that either the parameters used to estimate the IBNR are themselves misestimated or randomness around the mean prediction of IBNR claims leads to higher than expected claims. The former component of the uncertainty is usually referred to as "parameter error" and the latter as "process error".

To quantify these sources of uncertainty, practitioners often rely on two main methods: the closed-form formula of Mack (1993) for the chain ladder method and the Overdispersed Poisson (ODP) Bootstrap of England & Verrall (1999) for the generalised linear model (GLM) formulation of the chain ladder method, as well as the so-called "Mack bootstrap" of England and Verrall (2006).

Importantly, these practical uncertainty quantification methods are formulated with

1 Non-life insurance goes by different names in different jurisdictions, for example, property and casualty insurance in the United States. We may also refer to non-life insurance as short-term insurance in line with the practice in the South African market.

reference to the chain ladder method, which itself may not have been used to estimate the IBNR reserves. In particular, for those accident or underwriting years that are less developed, actuaries usually prefer using *exposure-based* methods, for example the BF or CC methods, to estimate the remaining development, since these methods are less sensitive to the potentially volatile experience often seen to emerge in the early development periods. On the other hand, whereas the CL method could be used for the older years once sufficient claims development has taken place, nonetheless, it usually makes relatively little difference what method is used in these older years.

A rough practical solution to bridge between the IBNR estimates on the one hand and the uncertainty estimates on the other is, firstly, to derive an estimate of the IBNR reserves using a combination of reserving methods; secondly, to estimate the coefficient of variation (CoV) of the chain ladder using these uncertainty quantification methods and, finally, to use the chain ladder CoV to derive the standard deviation of the reserves, and thus derive risk margins. An even rougher solution, used when bootstrap methods have been utilised for the second step above, is to force the mean of the bootstrap simulation to equal the reserve estimated in the first step, using an additive or multiplicative scaling factor. Nonetheless, these attempts to "transfer" the uncertainty estimates from the chain ladder method to the estimated reserves have little theoretical justification and may expose the insurer to further reserve misestimation risk.

The problem of mismatched reserves and uncertainty estimates are exacerbated in a Solvency II environment, which requires a specialised definition of the uncertainty that corresponds to the prudential objective of Solvency II. Solvency II requires that risk capital, more formally called the Solvency Capital Requirement (SCR), be calculated over a one-year time horizon, meaning to say that capital only needs to be held to the extent of covering the risks associated with the experience of the insurer over the next year. This is different from the traditional "ultimate" view of reserve risk examined in the actuarial literature before Solvency II and its antecedent, the Swiss Solvency Test (SST). Merz and Wüthrich (2008) provide a clear definition of how this concept should be understood within a reserving context, by focusing on two elements of the reserving results that apply over a single year: the difference between actual and expected claims over the single year and the difference between IBNR held for subsequent development periods at the reserving date and the same quantity estimated at the end of the next year. The sum of these two amounts are called the Claims Development Result (CDR) and Merz and Wüthrich (2008) provide a seminal one-year reserve risk formula for the standard deviation of the CDR, which is the standard actuarial method for estimating the reserve risk requirement in the Solvency II legislation. Importantly, this formula is again stated in the context of the CL model of Mack (1993); thus, when reserves are calculated using alternative methods such as the BF or CC, the same problem of non-correspondence between the IBNR estimates and the uncertainty estimates arises. Until now, we have only mentioned the one-year reserve risk formula of Merz and Wüthrich (2008), which could be viewed as an adaptation of the formula of

Mack (1993) to a one-year time horizon; in a further analogy, Ohlsson and Lauzeningks (2009) discusses how to adapt the bootstrap methods of reserve risk quantitation to the one-year time view in a procedure called the "actuary in a box". Similarly, this is usually formulated in terms of the CL method, thus the same problem mentioned here applies.

The first contribution of this paper to propose an approach to resolving these issues, which is to define a bootstrap method for the CC method and its generalisation due to Gluck (1997), the Generalised Cape Cod (GCC) method. The GCC method is an exposure-based reserving method that uses a data-driven weighted average procedure to derive estimated loss ratios for each loss year in a triangle. The weight, or credibility, given to each loss year is fully determined by the choice of a single parameter, which is the exponential decay parameter within this model. Since no subjective process is involved in applying the GCC method besides the selection of this parameter, it becomes significantly easier to formulate a bootstrap procedure for the GCC method than other exposure-based reserving techniques. In this case, since the GCC IBNR estimates are usually robust and can be used directly as the booked reserves for the insurer, the corresponding uncertainty estimates derived using the bootstrap now relate directly to the insurer's reserves, thus solving the problem of non-correspondence of best estimates and risk margins discussed above.

Our second contribution is to propose a solution to a vexing practical problem encountered when applying the GCC method, which is how the exponential decay parameter, defined here as γ, is chosen. Besides for experience-based heuristics (Struzzieri et al., 1998), only a single method for selecting this parameter in a data-based manner is known to us: the scoring framework of Balona and Richman (2020). Here, we add another option, which is to select the exponential decay parameter that minimises the expected root mean squared error of prediction of the IBNR, estimated using the GCC bootstrap discussed in this paper, as shown in Section 4 below. In the subsequent section we also show how the one-year reserve risk and its run-off to ultimate can be computed using the GCC bootstrap.

The last contribution is to show how the reserve risk, which relates to the claims that have already occurred, i.e., the retrospective risk, can be connected to the premium risk, which relates to claims that have yet to occur, i.e., the prospective risk. This link is made using the distribution of the reserving loss ratios produced using the GCC bootstrap; we propose[2] to utilise these to estimate a distribution of the future expected claims which can be used both for premium risk capital in Solvency II and the risk adjustment for the liability for remaining coverage.

2 To be more precise, this proposal can be used when reserving is performed using accident year loss triangles. However, if underwriting year triangles are used, the loss ratios produced by the GCC method may not relate fully to retrospective risks, but also prospective risk.

1.1 Literature review

We briefly review some of the main sources relevant to this working paper. The CC method is due to Bühlmann and Straub (1983) (see Bühlmann (2016) for a more accessible version of the paper), who presented the method at a company conference organised for one of the large reinsurers that was held in Cape Cod, Florida in the United States (Bühlmann, 2018). The GCC method is presented in Gluck (1997), who added a time varying factor to the original CC method. Struzzieri et al. (1998) recommend the use of the GCC method as a best practice for reserving and discuss how the method can be adapted for different situations, for example, when reserving using frequency and severity approaches. Saluz (2015) discusses a stochastic model for the CC method that uses a novel definition of the loss development pattern based on both claims and premiums (which is different from the usual approach that relies only on claims to derive the development pattern; see the definition of β_j in Section 2.3 below) and provides a formula the Mean Square Error of Prediction (MSEP) for the CC method. To our knowledge, the prediction uncertainty of the GCC method has not been examined in detail in the literature, although Ohlsson and Lauzeningks (2009) mention the GCC method as a good candidate for use in reserve risk estimation in Solvency II.

1.2 Structure of the manuscript

The rest of this working paper is structured as follows. In Section 2, we define notation and discuss the technical background of the paper. In Section 3, we describe the proposed GCC bootstrap used in the rest of the paper. In Section 4, we demonstrate the proposals in the case of ultimate risk on a range of triangles and compare these results to the Mack and ODP bootstrap. We then show how the proposed model works in the one-year case and demonstrate how the run-off of the reserve risk can be derived. Finally, we discuss how prospective insurance risk (premium risk) can be estimated using the proposed method. Finally, Section 5 concludes with discussion and future proposed research. The code underlying the paper can be found at https://github.com/RonRichman/GCC_bootstrap.

2. NOTATION AND BACKGROUND

Here, we provide a brief overview of the main notation and methods used in the following sections, following Wüthrich and Merz (2008).

2.1 Loss triangle notation

Estimating IBNR reserves relies on aggregating claims according to two different dates: the first is the date used to assign the claim to a cohort and the second is the date used to assign the claim to transaction period. For the former, there are two main ways in which the date can be chosen, either as the period in which the claim occurred (accident period), or the period in which policy under which a claim is being made was underwritten (underwriting period). For simplicity, in what follows, we will refer to the accident period definition and denote these accident periods as $i \in [0, I]$. For the latter, the usual date is chosen as the amount of time elapsed between the period in which the claim transaction was recorded and the period in which the claim occurred; we denote these as $j \in [0, J]$. In what follows, for simplicity, we assume that $I = J$, i.e., that we have an equal number *n* of accident periods and development periods. We denote the total claims incurred or paid relating to accident period *i* up to development period *j* as $C_{i,j}$ and the incremental change in claims between two development periods as $X_{i,j} = C_{i,j} - C_{i,j-1}$ or, in period $j = 0$ simply as $X_{i,0} = C_{i,0}$. The main purpose of IBNR reserving techniques is to estimate the unknown ultimate losses, which are the total losses in the accident period after the maximum development has taken place, in development period *J*, denoted as $C_{i,J}$. Usually, this is done on the basis of claims known up to calendar year $k = i + j$, for claims where $i + j < k$. Since the ultimate losses are unknown, $C_{i,J}$ is estimated at time *k* as $\hat{C}^k_{i,J}$ and this estimate leads to reserves $R^k_{i,j} = \hat{C}^k_{i,J} - C_{i,j}$. Usually, a volume measure indicating the amount of risk taken on by the entity in period *i* is available, for example, earned premium and we denote the volume measures as π_i.

2.2 Chain ladder reserving techniques

The most well known reserving technique is the chain ladder (CL) method, which extrapolates the most recently observed claims $C_{i,j}$ to ultimate using a development pattern derived from the past experience of how claims develop in each development period. Mathematically,

$$\hat{C}^{CL}_{i,J} = \hat{F}_{j,J} C_{i,j} = \prod_{l=j}^{J} \hat{f}_l C_{i,j}, \tag{2.1}$$

where $\hat{C}^{CL}_{i,J}$ is the ultimate estimated using the CL method, $\hat{F}_{j,J}$ is the estimated development factor from development period *j* to ultimate (*J*) and $\hat{f}_k$ is the estimated development factor from development period *k* to $k+1$. The development factors $\hat{f}_k$ are usually estimated as

$$\hat{f}_j = \frac{\sum_{i=1}^{I-j} C_{i,j+1}}{\sum_{i=1}^{I-j} C_{i,j}}. \tag{2.2}$$

A different approach to applying the CL method is to formulate a Generalised Linear Model (GLM) which reproduces the CL development (England & Verrall, 1999). This formulation of the CL method relies on a cross-classified representation of the incremental losses $X_{i,j}$ as

$$X_{i,j} = \alpha_i + \beta_j, \tag{2.3}$$

where α_i represents the ultimate losses for accident period *i* and β_j is the proportion of these losses developing in period *j*. This model is fit on the assumption that the observed incremental losses follow an Over Dispersed Poisson (ODP) distribution:

$$X_{i,j} \sim ODP\left(\alpha_i \beta_j, \phi\right), \tag{2.4}$$

where ϕ is the overdispersion parameter that scales the variance of the distribution to the extent that the incremental claims $X_{i,j}$ are more or less variable than the Poisson distribution. Under these assumptions, England and Verrall (1999) show that the model producing estimates of ultimate losses

$$\hat{C}_{i,J}^{ODP} = C_{i,j} + \sum_{k=j+1}^{J} \hat{X}_{i,k}^{ODP} \tag{2.5}$$

reproduces the CL estimates of ultimate losses, $\hat{C}_{i,J}^{CL}$.

Since the current observed claim $C_{i,j}$ is often highly variable in the earlier development periods, the CL ultimates, which are proportional to the current claims, see (2.1), can similarly suffer from excessive variability. More precisely, the difference between estimates of ultimate claims produced in subsequent calendar years k and $k+1$ might be judged to be too large; this problem leads to alternative methods of reserving using exposure-based techniques.

2.3 Exposure-based reserving techniques

Exposure-based techniques, which include the Bornhuetter-Ferguson or Cape Cod methods, are those reserving methods that derive estimated ultimate losses using a measure of how much risk the insurer has been exposed to during an accident period *i*. This measure is usually the earned premium which we represent as π_i, although other metrics, such as time-based exposures, can be used. It is assumed that the proportion of the IBNR claims developed up to development period *j* is known, and can be represented as β_j, which is an estimate of the percentage of losses reported at development period *j*, and is usually estimated using the CL development pattern as the inverse of $\hat{F}_{j,J}$ (note that this interpretation of β_j can also be used for β_j in (2.4)). Thus, the remaining exposure to IBNR claims at development period *j* is $\pi_i\left(1-\beta_j\right)$, which are then multiplied by an expected Ultimate Loss Ratio (ULR) at time k for accident year *i*, $\widehat{ULR}_i^k$, to derive the expected ultimate losses as $\hat{C}_{i,J} = C_{i,j} + \left(1-\beta_j\right)\pi_i \widehat{ULR}_i^k$. Thus, there is no need to use the highly variable current claim amounts $C_{i,j}$ in deriving ultimate losses, except to the extent that the development patterns or expected ULRs are derived using these estimates. Replacing the expected ULR with the ULR estimated using the CL method, $\widehat{ULR}_i^{k,CL}$, reproduces the CL ultimates of (2.1).

The Bornhuetter-Ferguson (BF) method assumes that independent estimates of the expected ULRs are available, and these are, in practice, set using business plans, pricing information or using observed ULRs from more developed accident years. The methods described next provide a basis for setting expected ULRs based on the claims experience in a principled manner.

2.3.1 CAPE COD AND GENERALISED CAPE COD

The Cape Cod (CC) method of Bühlmann and Straub (1983) uses a weighted average of ULRs derived using the CL method as the estimate of the expected ULR for all accident years, where the weights are the extent of development for each accident year, β_j.

$$\widehat{ULR}^{k,CC} = \frac{\sum_{n=1}^{I} \hat{C}_{i,J}^{CL} \beta_{I-n}}{\sum_{n=1}^{I} \pi_n \beta_{I-n}} = \frac{\sum_{n=1}^{I} C_{n,I-n}}{\sum_{n=1}^{I} \pi_n \beta_{I-n}}, \tag{2.6}$$

where the simplification on the right-hand side of this equation relies on the CL assumptions that $C_{i,j} = \hat{C}_{i,J}^{CL} \beta_j$. The weights β_j can be seen as estimates of the credibility of the CL estimates in each accident year; this assumption is similar to the credibility assumption in the Benktander-Hovinen method (Benktander, 1976; Mack, 2000).

A constant ULR for all accident years is usually too restrictive an assumption in cases when the loss ratio experience varies considerably by accident years. In these instances, the generalisation of the CC method of Gluck (1997) can be used, which assigns greater credibility to the loss experience in years that are close to each other using an exponentially weighted average:

$$\widehat{ULR}(\gamma)_i^{k,GCC} = \frac{\sum_{n=1}^{I} \hat{C}_{i,J}^{CL} \beta_{I-n+1} \gamma^{abs(i-n)}}{\sum_{n=1}^{I} \pi_n \beta_{I-n+1} \gamma^{abs(i-n)}} = \frac{\sum_{n=1}^{I} C_{n,I-n+1} \gamma^{abs(i-n)}}{\sum_{n=1}^{I} \pi_n \beta_{I-n+1} \gamma^{abs(i-n)}}, \tag{2.7}$$

where the decay factor $\gamma \in (0,1)$ is added to the weights of the weighted average in (2.6). It can be seen that

$$\lim_{\gamma \to 1} \widehat{ULR}(\gamma)_i^{k,GCC} = \widehat{ULR}^{k,CC}$$

and

$$\lim_{\gamma \to 0} \widehat{ULR}(\gamma)_i^{k,GCC} = \widehat{ULR}_i^{k,CL},$$

meaning to say that the GCC method reproduces the CC and CL ULRs for values of γ equal to 1 and 0, respectively.

In practice, γ is often set heuristically based on experience, for example, Struzzieri et al. (1998) recommends a value of $\gamma = 0.75$. Recently, Balona and Richman (2020) have proposed using a search procedure to find the value of γ using a scoring approach that measures how well a proposed reserving method is able to produce out-of-sample predictions.

2.4 Bootstrapping estimates of IBNR

The CL, BF and CC methods produce point estimates of expected ultimate claims and IBNR reserves; however, for risk management and financial reporting purposes, it is also necessary to quantify the potential for the estimated IBNR reserves to be different from the true required reserves, $R_{i,j}^k$. This is usually performed by estimating the MSEP of the estimated reserves, see Wüthrich and Merz (2008) for full discussion. Here, we focus on the bootstrapping approach of England and Verrall (1999, 2002) which relies on the GLM formulation of the CL method in (2.3). In summary, the approach is firstly to calculate adjusted residuals based on the observed incremental claims $X_{i,j}$ and the predicted incremental claims $\hat{X}_{i,j}^{CL}$, secondly, to bootstrap these residuals to produce multiple new claims triangles, thirdly, to add stochastic noise to the triangles corresponding to the randomness of claims around the best estimates and finally, to reapply the CL method to each simulated triangle to derive simulated IBNR and ultimate claims. The resulting distribution of simulated IBNR is then used to estimate the MSEP of the estimated IBNR. The full procedure (which will be run at time *k*) is as follows.

ALGORITHM 1: ODP BOOTSTRAP ALGORITHM OF ENGLAND AND VERRALL (2002)

1. Derive the CL estimates of ultimate claims $\hat{C}_{i,J}^{CL}$ and development factors $\hat{f}_j$.

2. Estimate the expected cumulative claims $\hat{C}_{i,j}^{CL}$ and, using these, the expected incremental claims $\hat{X}_{i,j}^{CL}$ in the upper triangle, i.e. for $i+j<k$.

3. Derive the Pearson residuals $r_{i,j}^{k}=\dfrac{X_{i,j}-\hat{X}_{i,j}^{CL}}{\sqrt{|(\hat{X}_{i,j}^{CL}|}}$.

4. Derive the scale parameter $\phi=\dfrac{\sum_{i+j<k}(r_{i,j}^{k})^2}{0.5\,n(n+1)-2n+1}$, where $n=I=J$.

5. Derive bias-adjusted Pearson residuals $r_{i,j}^{k,adj}=\sqrt{\dfrac{n}{0.5\,n(n+1)-2n+1}}r_{i,j}^{k}$.

6. Perform *S* simulations using the following procedure:
   (a) Sample the adjusted residuals $r_{i,j}^{k,adj}$ with replacement *n* times to form a new set of residuals $\{r_{i,j}^{*,adj};i+j<k\}$.

   (b) Derive simulated incremental claims $\{X_{i,j}^{k,CL,*}:=X_{i,j}+sign\left(X_{i,j}\right)\,r_{i,j}^{*,adj}\sqrt{|(\hat{X}_{i,j}^{CL}|};i+j<k\}$.

   (c) Derive the simulated cumulative claims $\left\{C_{i,j}^{k,CL,*}:=\sum_{l=1}^{j};X_{i,l}^{k,CL,*};i+j<k\right\}$.

   (d) Using the simulated cumulative claims, re-fit the standard CL model to derive development factors $\hat{f}_j^{*}$.

   (e) Project the ultimate cumulative claims $\hat{C}_{i,J}^{k,CL,*}$ and derive the future incremental claims $\{\hat{X}_{i,j}^{k,CL,*}:=\hat{C}_{i,j+1}^{k,CL,*}-\hat{C}_{i,j}^{k,CL,*};i+j>=k\}$.

(f) Apply the process variance distribution[3] with a mean of $\hat{X}_{i,j}^{k,CL,*}$ and a variance of $\phi\,\hat{X}_{i,j}^{k,CL,*}$ to derive projected future incremental claims with stochastic noise, $\hat{X}_{i,j}^{k,CL,**}$.

(g) Using projected future incremental claims, $\hat{X}_{i,j}^{k,CL,**}$, derive estimates of the reserves $\hat{R}_{i,j}^{k,*}$ and ultimate claims $\hat{C}_{i,J}^{k,*}$ and store these.

7. The simulated reserves and ultimate claims are used as a distribution around the best estimate of the reserves and ultimate claims, $\hat{R}_{i,J}^{CL,k}$ and $\hat{C}_{i,J}^{CL,k}$, respectively.

This bootstrapping procedure has been implemented in the `R` language as the function `BootChainLadder` in the package `ChainLadder` (Carrato et al., 2018).

We remark that, when implementing this procedure, step 6d may produce highly unrealistic LDFs f_j for the later development periods when j is close to J. In these instances, single large incremental claims in these later development periods are often produced using the bootstrap, leading to large estimates from the LDF estimator (2.2). From an actuarial perspective, these simulated LDFs are highly unrealistic, since single large LDFs in old development periods are generally not observed. Furthermore, even if observed, these are usually not treated as credible when reserving and will be smoothed or otherwise excluded from reserving estimates. Moreover, we find that these simulations contribute significantly to the MSEP estimated using Algorithm 1. For this reason, we recommend to modify Algorithm 1 at step 6d by introducing a "rejection" step that returns the algorithm to step 6a if the LDFs beyond a certain development period are too "rough", as measured by the second difference between the LDFs in each development period. We introduce this step in Algorithm 3 for the GCC bootstrap below.

2.5 From static to dynamic reserving

Solvency II and SAM require risk capital to be held based on a one-year time horizon, which, in the context of IBNR reserving, refers to the potential adverse development occurring over the year following the estimation of the reserves at time k. Merz and Wüthrich (2008) defined this as the Claims Development Result (CDR), consisting of two components: the difference between actual and expected claims in calendar year $k+1$ and the difference in estimated future IBNR claims for accident periods $i \leq I$ as at time $k+1$ compared to time k. Thus, the CDR for accident year i, that is developed up to development period j in calendar year k is defined as

3 ODP or Gamma distributions

$$
\begin{aligned}
CDR_i^k &= \hat{C}_{i,J}^{k+1} - C_{i,J}^k \\
&= \hat{R}_{i,j+1}^{k+1} + C_{i,j+1} - \left(\hat{R}_{i,j}^k + C_{i,j}\right) \\
&= \sum_{l=j+1}^{J} \hat{X}_{i,l}^{k+1} + C_{i,j+1} - \left(\sum_{l=j}^{J} \hat{X}_{i,l}^k + C_{i,j}\right) \\
&= \hat{R}_{i,j+1}^{k+1} - \hat{R}_{i,j}^k + \left(X_{i,j+1} - \hat{X}_{i,j+1}^k\right) \\
&= \hat{R}_{i,j+1}^{k+1} - \hat{R}_{i,j}^k + AvE_{i,j}^k,
\end{aligned}
\tag{2.8}
$$

where $AvE_{i,j}^k$ is the actual versus expected (AvE) result of the incremental claims on the next diagonal of the triangle, i.e., the forecast error of the incremental claims in the next development period, where the forecast was made in calendar year k. It can be seen that the residuals defined in step 3 of Algorithm 1 are similar to this definition of the AvE, i.e., a procedure that consistently produces predictions with a small AvE will (eventually) lead to a smaller MSEP than those with a large AvE. We therefore expect that results of procedures that select IBNR methods based on minimising the AvE should correspond quite closely to procedures that select IBNR methods based on minimising the bootstrap error, and, indeed, we show this is the case in the example below.

It can also be shown that the expected value of the CDR is zero, i.e., if reality turns out in accordance with the CL assumptions underlying the CDR, then the CDR will be zero. To estimate solvency capital for Solvency II and SAM, it is necessary to derive the MSEP of the CDR, and this can be done with an adapted version of Algorithm 1, following the "actuary in the box" procedure of Ohlsson and Lauzeningks (2009). The key change is made to add several more steps after 6f of Algorithm 1.

ALGORITHM 2: ODP BOOTSTRAP ALGORITHM – CLAIMS DEVELOPMENT RESULT

1. Derive the CL estimates of ultimate claims $\hat{C}_{i,J}^{CL}$ and development factors $\hat{f}_j$.

2. Estimate the expected cumulative claims $\hat{C}_{i,j}^{CL}$ and, using these, the expected incremental claims $\hat{X}_{i,j}^{CL}$ in the upper triangle, i.e. for $i+j<k$.

3. Derive the Pearson residuals $r_{i,j}^k = \dfrac{X_{i,j} - \hat{X}_{i,j}^{CL}}{\sqrt{|(\hat{X}_{i,j}^{CL}|}}$.

4. Derive the scale parameter $\phi = \dfrac{\sum_{i+j<k} (r_{i,j}^k)^2}{0.5\, n(n+1) - 2n + 1}$, where $n = I = J$.

5. Derive bias-adjusted Pearson residuals $r_{i,j}^{k,adj} = \sqrt{\dfrac{n}{0.5\, n(n+1) - 2n + 1}} r_{i,j}^k$.

6. Perform S simulations using the following procedure:
   (a) Sample the adjusted residuals $r_{i,j}^{k,adj}$ with replacement n times to form a new set of residuals $\{r_{i,j}^{*,adj}; i+j<k\}$.

   (b) Derive simulated incremental claims $\{X_{i,j}^{k,CL,*} := X_{i,j} + r_{i,j}^{*,adj}\sqrt{|(\hat{X}_{i,j}^{CL}|}; i+j<k\}$.

   (c) Derive the simulated cumulative claims $\left\{C_{i,j}^{k,CL,*} := \sum_{l=1}^{j} X_{i,l}^{k,CL,*}; i+j<k\right\}$.

   (d) Using the simulated cumulative claims, re-fit the standard CL model to derive development facts $\hat{f}_j^*$.

   (e) Project the ultimate cumulative claims $C_{i,J}^{k,CL,*}$ and derive the future incremental claims $\{\hat{X}_{i,j}^{k,CL,*} := \hat{C}_{i,j+1}^{k,CL,*} - \hat{C}_{i,j}^{k,CL,*}; i+j>=k\}$.

   (f) Apply the process variance distribution with a mean of $\hat{X}_{i,j}^{k,CL,*}$ and a variance of $\phi\,\hat{X}_{i,j}^{k,CL,*}$ to derive projected future incremental claims with stochastic noise, $\hat{X}_{i,j}^{k,CL,**}$.

   (g) Set to zero all projected future incremental claims with stochastic noise $\hat{X}_{i,j}^{k,CL,**}$ where $i+j>k$.

   (h) Derive the simulated cumulative claims one calendar year ahead $\left\{\hat{C}_{i,j+1}^{k+1,CL,**} := C_{i,j} + \hat{X}_{i,j}^{k,CL,**}; i+j=k\right\}$.

   (i) Using the simulated cumulative claims one calendar year ahead, re-fit the standard CL model to derive development factors $\hat{f}_j^{**}$.

   (j) Project the ultimate cumulative claims $\hat{C}_{i,J}^{k+1,CL,**}$ and derive the CDR as $CDR_i^{k,CL,**} = \hat{C}_{i,J}^{k+1,CL,**} - \hat{C}_{i,J}^{k}$ and store this estimate .

7. The simulated estimates of the CDR form a distribution around the best estimate of the CDR, which is 0.

2.5.1 EXTENDING THE CDR

In the previous section, we have discussed the CDR over the next calendar period $k+1$. From a practical perspective, if reserves are calculated using a quarterly-quarterly triangle, then for Solvency II purposes, it is necessary to calculate the CDR over the next four calendar periods. In addition, to anticipate future risk capital needs, it is valuable to be able to estimate the run-off of the required reserve risk capital, for example, when calculating cost-of-capital based margins, it is necessary to have an estimate of the required capital at each future period in which this is required. For this reason, Merz and Wüthrich (2014) show how the full run-off of the MSEP can be estimated in the context of Mack's CL model (Mack, 1993) (note that, usually, in practical applications, an approximation is used for the run-off of capital requirements).

We therefore extend the definition of (2.8) to cater for multiple periods as follows:

$$CDR(j)_i^k = C_{i,J}^{k+j} - C_{i,J}^k, \tag{2.9}$$

where $j \leq J$ and where $CDR(J)_i^k = C_{i,J} - C_{i,J}^k$, i.e., the ultimate claims after the full run-off has been observed less the estimates in calendar period *k*. Algorithm 2 can likewise be extended in a simple manner to estimate the MSEP of $CDR(l)$ by modifying step 6(i)i to set to zero only those projected future incremental claims with stochastic noise where $i+j>k+l-1$. This leads to a highly efficient approach to deriving the reserve run-off, as we describe next.

2.5.2 EFFICIENT SIMULATION OF THE CDR AND THE RESERVE RISK RUN-OFF

Running Algorithm 2 for each value of $l \in [I-i+1, J]$ for accident period *i* is computationally onerous, especially for large triangles with longer development periods *J*. To reduce the computational burden, it can be seen that Algorithm 2 is the same as Algorithm 1 up to step 6, after which some of the diagonals of the simulated triangle, where $i+j>k$, are successively set to zero to derive the correct progression of simulated future diagonals to estimate $CDR(j)$. Instead of running Algorithm 2 multiple times, one can simply run one simulation step of Algorithm 1, which provides the full simulated lower triangle, and then successively set the simulated future diagonals to zero for each value of $CDR(j)$ that is required. In this way, only one simulated triangle is needed to derive a simulated estimate of the ultimate risk and each value of $CDR(j)$.

2.6 Summary

This section started with the CL method, discussed exposure-based methods, and then focused on bootstrapping for the ultimate, one-year and multi-year reserve risk. In what follows, we adapt the bootstrap presented here for the GCC method.

3. GCC BOOTSTRAP

The main idea of the Cape Cod bootstrap is to replace all the steps in Algorithms 1–2 that relate to the CL method with corresponding steps that relate to the GCC method. Since the weight parameter γ in (2.7) is unknown, we run the algorithm below multiple times, updating the value of γ for each run of the algorithm on a grid spanning [0,1] in increments of 0.1 to estimate a value of the MSEP and other quantities for each value of γ. The main difference between the algorithm presented here and Algorithm 2 above is that the CL method has now been replaced with the GCC method.

ALGORITHM 3: GCC BOOTSTRAP ALGORITHM FOR A GIVEN VALUE OF γ

1. Derive the GCC estimates of ultimate claims $\hat{C}_{i,J}^{GCC}$ for the value of γ and development factors $\hat{f}_j$.

2. Estimate the expected cumulative claims $\hat{C}_{i,j}^{GCC}$ and, using these, the expected incremental claims $\hat{X}_{i,j}^{GCC}$ in the upper triangle, i.e. for $i+j<k$.

3. Derive the Pearson residuals $r_{i,j}^{k}=\dfrac{X_{i,j}-\hat{X}_{i,j}^{GCC}}{\sqrt{|(\hat{X}_{i,j}^{GCC}|}}$.

4. Derive the scale parameter $\phi=\dfrac{\sum_{i+j<k}(r_{i,j}^{k})^2}{0.5\;n(n+1)-2n+1}$, where $n=I\;J$.

5. Derive bias-adjusted Pearson residuals $r_{i,j}^{k,adj}=\sqrt{\dfrac{n}{0.5\;n(n+1)-2n+1}}r_{i,j}^{k}$.

6. Perform S simulations using the following procedure:

   (a) Sample the adjusted residuals $r_{i,j}^{k,adj}$ with replacement n times to form a new set of residuals $\{r_{i,j}^{*,adj};i+j<k\}$.

   (b) Derive simulated incremental claims
   $\{X_{i,j}^{k,GCC,*}:=X_{i,j}+sign\left(\hat{X}_{i,j}^{GCC}\right)\,r_{i,j}^{*,adj}\sqrt{|(\hat{X}_{i,j}^{GCC}|};i+j<k\}\cdot$

   (c) Derive the simulated cumulative claims $\left\{\hat{C}_{i,j}^{k,GCC,*}:=\sum_{l=1}^{j}\hat{X}_{i,l}^{k,GCC,*};i+j<k\right\}$.

   (d) Using the simulated cumulative claims, re-fit the GCC model, first by deriving development factors $\hat{f}_j^{*}$, then by re-estimating the GCC expected ULR using (2.7).

   (e) If the second differences of the simulated LDFs $\hat{f}_j^{*}$ beyond a selected development period are larger than a threshold t, return to step 6a.

   (f) Project the ultimate cumulative claims $\hat{C}_{i,J}^{k,GCC,*}$ and derive the future incremental claims $\{\hat{X}_{i,j}^{k,GCC,*}:=\hat{C}_{i,j+1}^{k,GCC,*}-\hat{C}_{i,j}^{k,GCC,*};i+j>=k\}$.

   (g) Apply the process variance distribution with a mean of $\hat{X}_{i,j}^{k,GCC,*}$ and a variance of $\phi\;\hat{X}_{i,j}^{k,GCC,*}$ to derive projected future incremental claims with stochastic noise, $\hat{X}_{i,j}^{k,GCC,**}$.

(h) Using projected future incremental claims, $\hat{X}_{i,j}^{k,GCC,**}$, derive estimates of the reserves $\hat{R}_{i,j}^{k,*}$ and ultimate claims $\hat{C}_{i,J}^{k,*}$ and store these.

(i) For each $l \in [1, L]$ perform the following procedure:

i. Set to zero all projected future incremental claims with stochastic noise $\hat{X}_{i,j}^{k,GCC,**}$ where $i + j > k + l$.

ii. Derive the simulated cumulative claims l calendar years ahead

$$\left\{ \hat{C}_{i,j+l}^{k+l,GCC,**} := C_{i,j} + \sum_{m=1}^{l} \hat{X}_{i,j+m}^{k,GCC,*} ; i + j = k + l \right\}.$$

iii. Using the observed cumulative claims and the simulated cumulative claims l calendar year ahead, re-fit the standard LDFs from the CL model to derive development factors $\hat{f}_j^{**}$.

iv. Project the ultimate cumulative claims $C_{i,J}^{k+l,GCC,**}$ and derive the CDR as $CDR(l)_i^{k,GCC,**} = \hat{C}_{i,J}^{k+l,GCC,**} - C_{i,J}^{k}$ and store this estimate.

7. The simulated reserves, ultimate claims and CDRs are used as a distribution around the best estimate of the reserves and ultimate claims, $\hat{R}_{i,J}^{GCC,k}$, $\hat{C}_{i,J}^{GCC,k}$ and 0, respectively.

We remark that Steps 4 and 5 contain an approximation, in the sense that the number of parameters used in the GCC method are fewer than those in the original CL method, due to the smoothing of the loss ratios produced by the GCC method. The number of parameters for the GCC method will vary as the weight parameter γ changes: at low values of γ, this will be close to the CL method, and at high values of γ, these will be lower. Since the more significant part of the bias adjustment is the number of observations $0.5n(n+1)$ and not the number of parameters, which are $2n$ for the CL method, we accept this approximation here and leave refinements for future work.

We suggest, as an heuristic, to use the values of MSEP produced by running Algorithm 3 to select the optimal value of γ i.e. $\arg\min_{\gamma}$MSEP(γ), where MSEP(γ) is the value of the MSEP produced by running Algorithm 3, setting the weighting parameter equal to γ. Since the MSEP estimates an expectation of the out-of-sample performance of the GCC method, the parameter producing the minimum estimate of the MSEP will, from this perspective, produce an optimal estimate of IBNR. We attempt to validate this heuristic in the example that follows.

4. EXAMPLE

We provide an example of applying the proposed GCC Bootstrap to a motor property damage triangle from a large and stable portfolio. The triangle is constructed with both quarterly accident and development periods and consists of claims paid as well as changes in Outstanding Claims Reserves (OCR), i.e., this is an incurred triangle. To preserve confidentiality of the company providing the triangle, both the claims and earned premium figures have been divided by random real numbers and, furthermore, the historical period covered by the triangle, which does not align to the publication date of this working paper, is not disclosed. This triangle is presented in Figure 1.

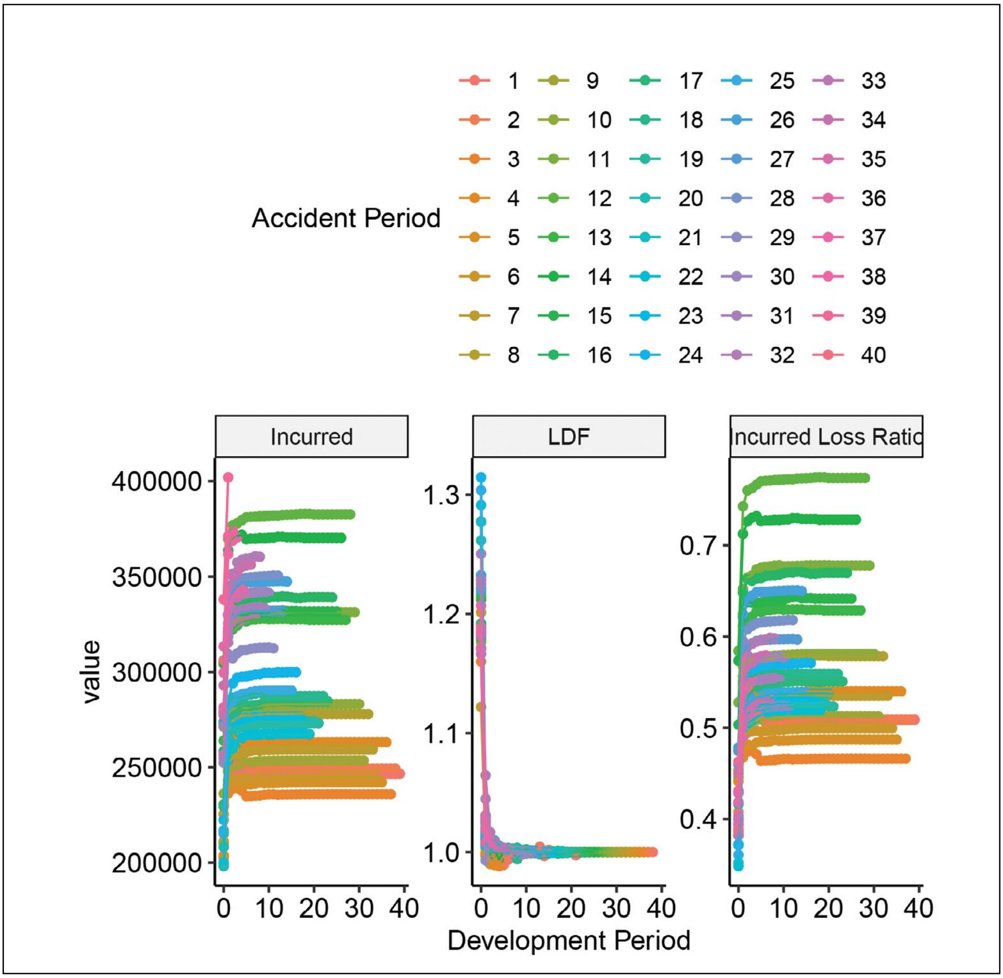

FIGURE 1 Incurred Claims, Loss Development Factors, and Incurred Loss Ratios for the example motor property triangle. All figures have been scaled by random variables such that the development structure of the triangle has been maintained, but the absolute figures are not meaningful.

The values of the required IBNR reserves produced by applying the CL, GCC and CC methods are shown in Table 1, noting that setting the values of γ to 0 and 1 reproduce the CL and CC methods, respectively. It can be observed that the value of the required IBNR is largest for values of γ in the range 80%–100%.

TABLE 1 Best-estimate IBNR for different values of the weight parameter, γ. The CL method is reproduced for $\gamma = 0$ and the CC method for $\gamma = 1$.

Decay	GCC BE IBNR
0.00%	98,960
10.00%	99,785
15.00%	100,149
20.00%	100,479
25.00%	100,775
30.00%	101,035
35.00%	101,259
40.00%	101,447
45.00%	101,601
50.00%	101,723
55.00%	101,819
60.00%	101,895
65.00%	101,962
70.00%	102,036
75.00%	102,138
80.00%	102,300
85.00%	102,578
90.00%	103,047
95.00%	103,685
100.00%	104,056

4.1 Estimating the ultimate risk in the GCC model

We begin by showing in Table 2 the values of the root of the MSEP (rMSEP) (i.e. the standard deviation of the IBNR reserves) calculated using Mack's method (Mack, 1993) and the bootstrap method assuming an over-dispersed Poisson distribution, as well as a gamma process distribution (England & Verrall, 2002). In this case, the values produced using Mack's method show a coefficient of variation (CoV) of about 15% and the values produced with the bootstrap method are slightly lower at about 13.5%. Mack's method does not produce a distribution of IBNR values, thus, we have not provided a value for the percentiles in Table 2.

Table 2 Estimates of the rMSEP, coefficient of variation and 75th percentile for the motor triangle, estimated using Mack's methods and two variations of the bootstrap method.

Method	Mean IBNR estimate	rMSEP	Coefficient of variation	75th percentile
ODP	98,810	13,339	13.50%	107,654
Mack	98,960	15,015	15.17%	NA
Gamma	98,668	13,411	13.59%	107,270

We compare these results to those produced using Algorithm 3 for the GCC method, which are shown in Table 3. For γ equal to zero, Algorithm 3 produces a similar estimate of the rMSEP and CoV, which is directly comparable to those values produced by the ODP method in Table 2, but is slightly higher, perhaps due to applying the suggested approach in Step 6e of Algorithm 3. For higher values of γ, a relatively stable progression of rMSEP and CoV values can be observed, that are lower than those produced by the CL case.

The former quantity is shown in Figure 2 where it can be observed that the minimum value of the rMSEP is produced when γ is set to 95%, with a similar value of the rMSEP produced when γ is set equal 35%. Similar low values of the rMSEP are produced for values of γ set to 60% and 80%. Seemingly, this indicates that the CL method produces an inflated value of the rMSEP and is less suitable than the GCC method, for this triangle. We would suggest the selection of one of these values of γ for reserving purposes.

A nice consequence of aligning the method used to derive the best estimate IBNR and the bootstrapped distribution of the IBNR can be observed by comparing the IBNR estimates in Tables 1 and 3. It can be seen that these estimates are quite close for the different values of γ, whereas this is not usually the case if, for example, the GCC method was applied for the best estimate and the ODP bootstrap was applied to derive the bootstrapped distribution, which, for this triangle, would produce estimates of the IBNR of 103,685 and 98,938, respectively.

Figure 3 shows empirical kernel density estimates of the IBNR distribution derived for values of $\gamma \in [0, 0.3, 0.6, 0.7, 0.75, 0.95, 1]$. It can be observed that among all these values of γ, the CL method produces a distribution shifted to the left, i.e., the mean of the distribution is different, and with significant weight placed on lower values of the distribution. The distribution produced by the CC method is somewhat shifted to the right. The distributions for other values of γ lie between these extremes. Among the distributions shown for the GCC method, the distribution produced with a value of γ set to 95% is the most peaked and the thinnest tailed. Thus, the choice of the weight parameter may have significant influence on the amount of risk capital needed for reserve risk; indeed, Figure 4 shows that the CL and GCC methods with low values of γ tend to estimate higher reserve risk capital requirements (on an ultimate basis) than the GCC method with higher values of γ and the original CC method.

The GCC bootstrap results have indicated a value of the weight parameter of 95%.

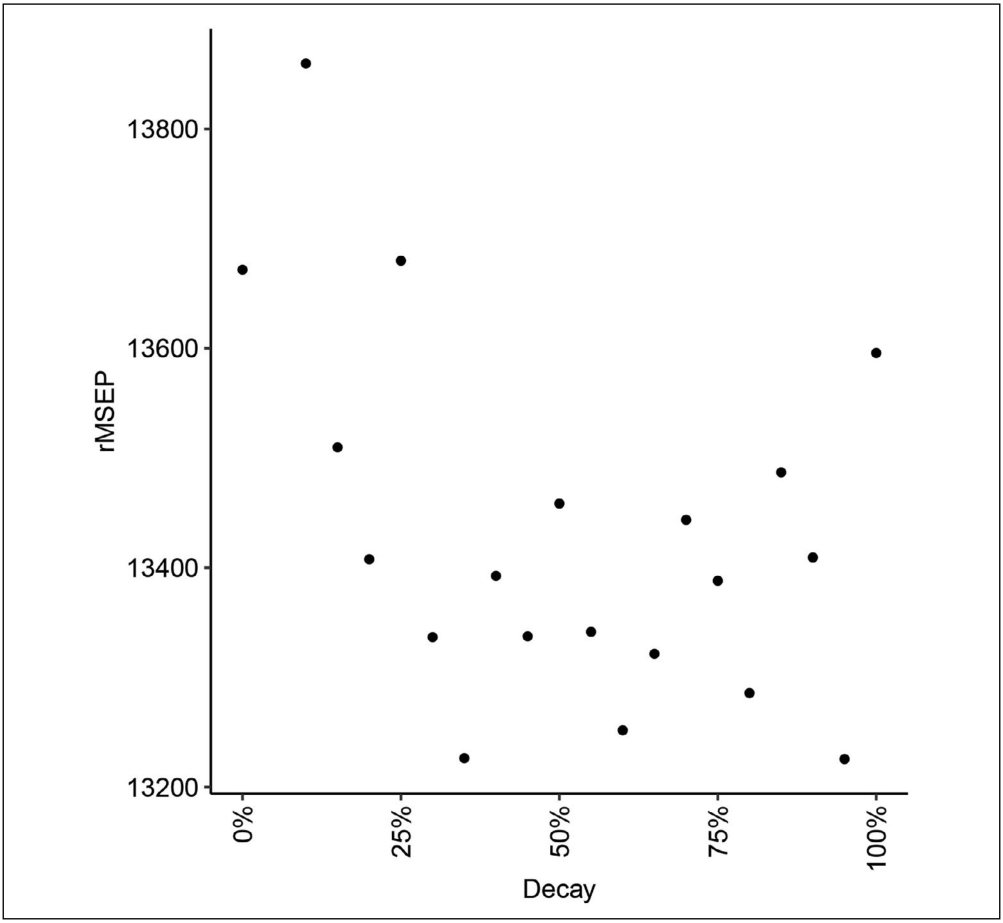

FIGURE 2 Estimates of the rMSEP for different values of the decay parameter, γ

This indication can be compared to the results produced using the method of Balona and Richman (2020). In that work, the loss development factors were calculated for using only some of the available accident years and allowed for high and low values of these factors to be excluded from the calculation i.e. the method for estimating the LDFs discussed there is significantly more complex than the more simple methods used here, see (2.2). Using the tryangle package of Balona (2021) to implement the method of Balona and Richman (2020), the optimal parameter γ was re-estimated on these simpler assumptions that the LDFs should use all accident years of data and that high and low LDFs are not dropped from the calculation. In this case, the optimal value of γ was found to be 80% using both the CDR and AvE as scoring objectives and 80% using the AvE as a scoring objective. These results, which indicate one of the lower values of the RMSEP produced using the GCC bootstrap, indicate that the heuristic proposed here produces similar results to a method using a different, yet related, principle.

Table 3 Estimates of the rMSEP, coefficient of variation and 75th percentile for the motor triangle, estimated using Algorithm 3.

Decay	Mean IBNR estimate	rMSEP	Coefficient of variation	75th percentile
0.00%	98,938	13,672	13.82%	107,584
10.00%	100,618	13,860	13.77%	109,487
15.00%	101,040	13,510	13.37%	109,796
20.00%	101,334	13,408	13.23%	110,019
25.00%	101,528	13,680	13.47%	110,445
30.00%	102,266	13,337	13.04%	110,899
35.00%	102,271	13,226	12.93%	110,774
40.00%	101,904	13,393	13.14%	110,857
45.00%	102,374	13,337	13.03%	111,143
50.00%	102,306	13,458	13.16%	110,947
55.00%	102,388	13,342	13.03%	111,066
60.00%	102,370	13,252	12.94%	111,106
65.00%	102,467	13,322	13.00%	111,094
70.00%	102,130	13,444	13.16%	111,059
75.00%	102,385	13,388	13.08%	111,119
80.00%	102,811	13,286	12.92%	111,478
85.00%	103,158	13,487	13.07%	112,204
90.00%	103,772	13,409	12.92%	112,447
95.00%	103,950	13,226	12.72%	112,716
100.00%	104,402	13,596	13.02%	113,306

4.2 Estimating the one-year risk in the GCC model

The results above were on an ultimate basis, however, as discussed, Solvency II and SAM require that reserve risk capital be estimated on a one-year basis. We now show results based on the remaining steps of Algorithm 3. To provide a baseline comparison, the full reserve-risk run-off for Mack's CL model was estimated using the methods in Merz and Wüthrich (2014) as implemented in the function CDR in the package ChainLadder (Carrato et al., 2018) and is shown in Table 4.

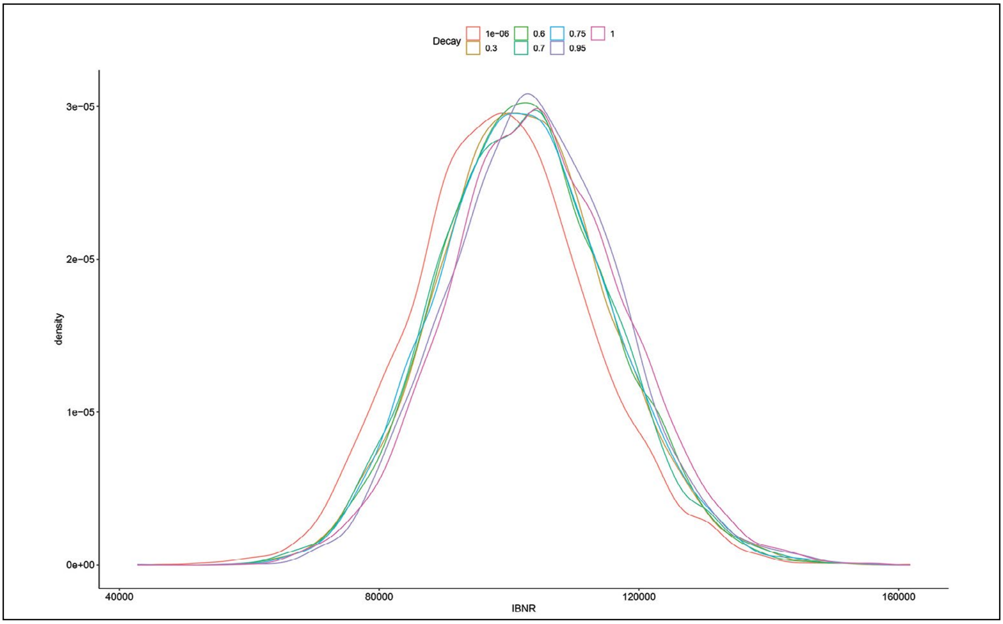

FIGURE 3 Density of the simulated distribution of the IBNR for values of $\gamma \in [0, 0.3, 0.6, 0.7, 0.75, 0.95, 1]$

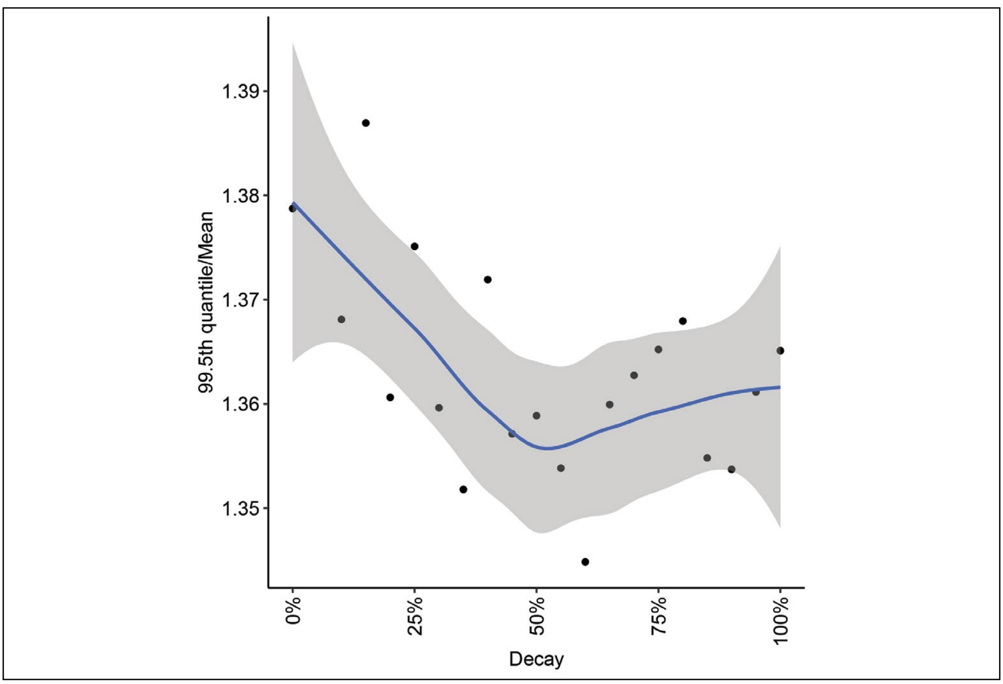

FIGURE 4 99.5th quantile of the simulated distribution of IBNR for different values of γ divided by the mean IBNR, smoothed using LOESS regression line.

Table 4 Reserve risk run-off in Mack's CL model

Calendar quarter	Cumulative standard deviation of the CDR	Percentage of ultimate reached
1	13,038	86.83%
2	14,136	94.15%
3	14,428	96.09%
4	14,607	97.28%
5	14,724	98.06%
6	14,796	98.54%
7	14,843	98.86%
8	14,880	99.10%
9	14,912	99.31%
10	14,933	99.45%
11	14,949	99.56%
12	14,964	99.66%
13	14,977	99.74%
14	14,988	99.82%
15	14,995	99.87%
16	15,000	99.90%
17	15,003	99.92%
18	15,006	99.94%
19	15,009	99.96%
20	15,011	99.97%
21	15,013	99.98%
22	15,014	99.99%
23	15,014	100.00%

We compare these results to those produced by the GCC method with γ set to 95%. The GCC bootstrap reaches ultimate faster than the CL model, but, on the other hand, less of the reserve risk has run-off after a single year. Thus, the choice of reserving method can influence the amount of reserve risk capital needed for Solvency II and SAM.

In Table 6 we compare the reserve risk parameters derived using the GCC bootstrap, the Merz-Wüthrich method and the prescribed parameter for this line of business in SAM. While there is broad consistency between the first two of these methods, with a minimum parameter produced using the GCC bootstrap and γ set to 95%, the SAM parameter is significantly lower. This is because the SAM parameter is calibrated to be applied both to IBNR reserves as well as OCR, i.e., the normalisation of the standard deviation of the CDR has been performed using a different base (note this is an incurred triangle, thus only IBNR is estimated). For the SAM parameter to produce a similar CDR as the other methods, there would roughly need to be a ratio of OCR to IBNR of 1:1, which seems unlikely for this line of business.

TABLE 5 Reserve risk run-off in the GCC model for $\gamma = 0.95$

Calendar quarter	Cumulative standard deviation of the CDR	Percentage of ultimate reached
1	10,860	82.10%
2	11,804	89.24%
3	12,255	92.65%
4	12,520	94.66%
5	12,729	96.24%
6	12,833	97.02%
7	12,917	97.66%
8	12,957	97.96%
9	13,025	98.48%
10	13,109	99.11%
11	13,166	99.54%
12	13,178	99.63%
13	13,197	99.77%
14	13,214	99.90%
15	13,244	100.13%*

*Note that this slight amount above 100% is due to simulation noise.

TABLE 6 Reserve risk parameters using the GCC bootstrap, the Merz-Wüthrich method and SAM

Method	Decay	Mean IBNR	Percentage of ultimate	Standard deviation of CDR	Sigma
GCC Bootstrap	0.00%	98,938	95.04%	12,994	13.1%
GCC Bootstrap	30.00%	102,266	95.19%	12,696	12.4%
GCC Bootstrap	60.00%	102,370	94.79%	12,563	12.3%
GCC Bootstrap	70.00%	102,130	94.29%	12,677	12.4%
GCC Bootstrap	75.00%	102,385	95.29%	12,759	12.5%
GCC Bootstrap	95.00%	103,950	94.66%	12,520	12.0%
GCC Bootstrap	100.00%	104,402	94.23%	12,813	12.3%
Merz-Wüthrich	0.00%	98,960	94.66%	12,520	12.7%
SAM	0.00%	98,960	94.66%	5,938	6.0%

4.3 Premium risk considerations

Having addressed the risk arising from claims arising from loss events in the past, i.e., retrospective risk, we now focus on the risk arising from claims arising from events occurring in the future, i.e., prospective risks, which is called "premium risk" in Solvency II and SAM. Figure 5 shows various different loss ratios estimated using the GCC bootstrap, which are the means and percentiles of the ULRs for each accident quarter $\widehat{ULR} = \frac{\hat{C}_{i,J}^{k,CC}}{\pi_i}$

and the estimated GCC loss ratios, $\widehat{ULR}^{CC}$. These are shown for several different values of the weight parameter γ. The mean GCC ULR varies quite dramatically as the weight parameter is changed, with more smoothing as γ is set closer to 1; similarly, the quantiles of the distribution of the GCC ULR behave in a similar manner. On the other hand, the ULR for each accident period, derived using the estimated ultimates $\hat{C}_{i,J}^{k,CC} = C_{i,J} + \hat{R}_{i,j}^{k,CC}$ are much less affected by changing the weight parameter, and indeed, only the most recent accident period is uncertain due to the extent of the uncertainty of the IBNR reserves, $\hat{R}_{i,j}^{k,CC}$, i.e., the total uncertainty of the ultimates is significantly lower than that of the reserves, since the claims in the most recent period are a constant.

Since the GCC ULR is estimated based on the limited claims data available, these may suffer from some parameter misestimation risk, which is what is quantified by the estimated quantiles of the GCC ULR. Figure 5 shows that the potential misestimation risk is largest for small values of γ and becomes negligible for the selected value of $\gamma = 0.95$. On the other hand, the mean GCC ULR tracks the ULRs of each accident year much more closely for low values of γ than higher values, meaning to say, that low values of γ will seemingly reduce the potential forecast error produced by using the GCC ULR as an estimate of the accident year ULR. This is because low values of γ produce estimates of the ULR that are more adapted to the experience in each accident year. We remark that the potential forecast error is close to the quantity targeted by the estimators of the premium risk using the formula given in, for example, the SAM regulatory standards (Prudential Authority, 2020) to estimate premium risk, on the assumption that the forecast loss ratio is a constant. Of course, since the premium risk relates to experience that has not yet been observed, more correct ways of quantifying the premium risk are, firstly, by allowing the forecast to vary each year, and also, by estimating how well the GCC ULR forecasts the experience in the next year. Finally, a simple way of quantifying premium risk is simply to estimate the standard deviation of the ULRs over time; this is also similar to the quantity targeted in SAM. Estimates for each of these definitions of premium risk are shown in Table 7, and are comparable to the estimate of premium risk for this line of business in SAM of 6.3%.

Table 7 Estimates of premium risk parameters

Decay	Parameter estimation error	Forecast error, current year	Forecast error, next quarter	Standard deviation of ULR
0%	4.4%	0.1%	5.1%	6.5%
30%	3.1%	3.0%	3.2%	6.5%
60%	2.0%	4.3%	3.5%	6.5%
70%	1.8%	5.0%	4.0%	6.5%
75%	1.6%	5.3%	4.2%	6.5%
95%	0.8%	6.4%	4.9%	6.5%
100%	0.8%	6.4%	4.9%	6.5%

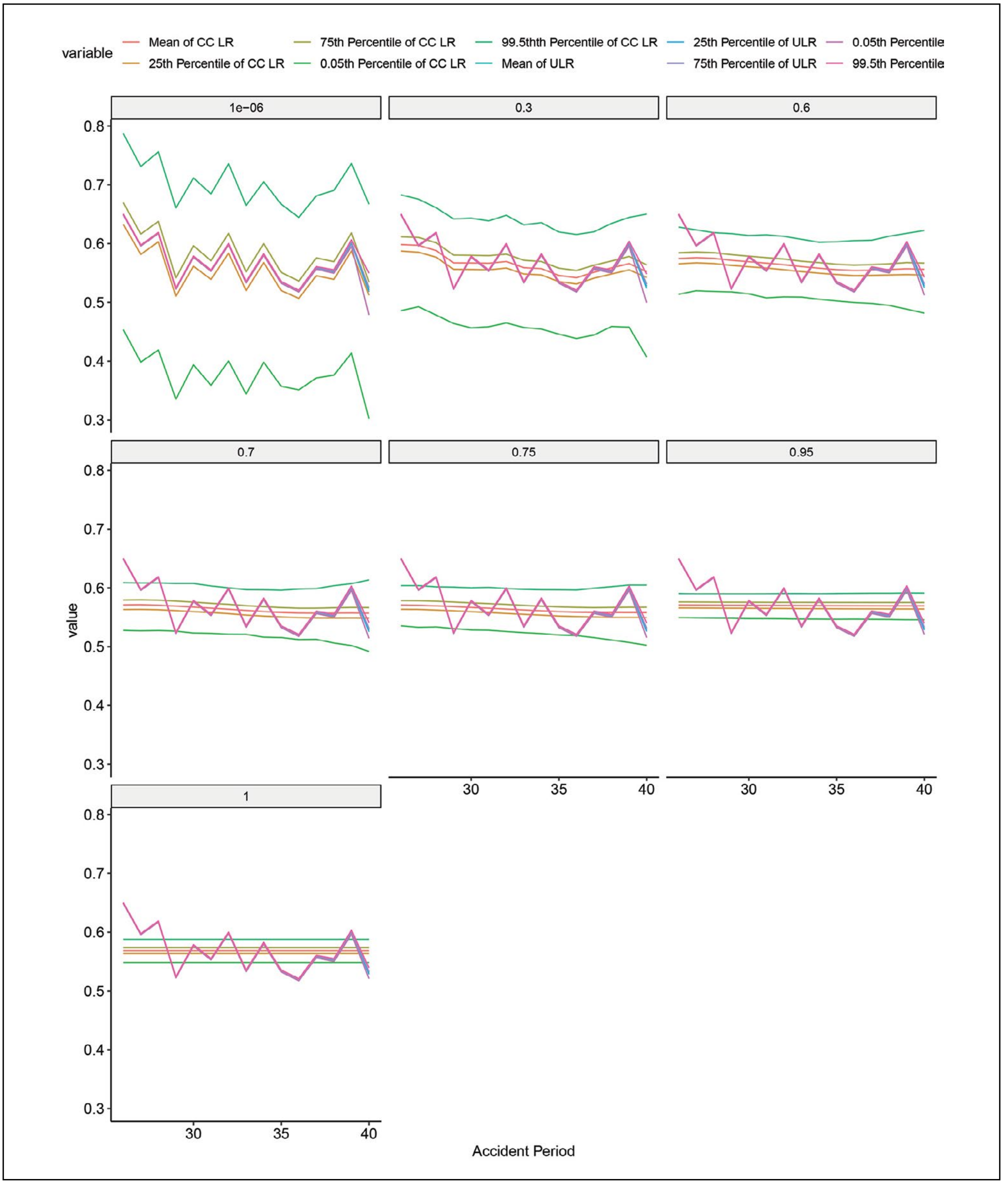

FIGURE 5 Means and percentiles of the ULRs for each accident year and estimated GCC loss ratios

5. DISCUSSION AND CONCLUSION

A bootstrap method for estimating the rMSEP of IBNR reserves has been provided in this working paper and we have demonstrated that more closely aligning the method used to produce the best estimates of IBNR and the rMSEP reduces the number of arbitrary adjustments commonly made in practical IBNR reserving. We have suggested a method for selecting the weight parameter of the GCC method and discussed that other, similar methods, produce comparable outcomes. Finally, we have shown that several important outcomes of the reserving process – best estimates of IBNR, rMSEP, extent of the reserve risk run-off, and overall reserve distribution – depend on the choice of reserving method and that significant risk of misestimation exists if an assumption is made that the CL method produces acceptable proxies for each of these. This has been shown in the context of a very short-tailed triangle; it is likely that for long-tailed triangles, these differences will be significantly more pronounced. Given the requirements for accurate reserve risk estimation in Solvency II and SAM, it appears that, at least as a check, this assumption should be verified by applying methods that align the best estimate to the estimates of the distribution of the IBNR. In a similar manner, reserve estimation work performed for IFRS 17, which requires that risk margins (adjustments) be estimated for claims reserves, will likely benefit from the methods proposed here. For the short-tailed triangle considered here, the difference between the one-year and ultimate views of reserve risk were only different by about 5%, and in similar instances to this, it seems reasonable that the ultimate reserve risk parameters can be used to approximate the one-year view since the difference is immaterial. On the other hand, for longer-tail business, this is unlikely to be the case. There are several avenues for future research arising from this work. The most important of these is to justify rigorously the heuristic proposed for selecting γ on the basis of the rMSEP. Another extension is to improve the approximation of the scale parameter ϕ depending on the effective number of parameters used in the GCC method. An interesting follow-up would be to consider the various definitions of premium risk in more detail. Finally, a larger scale study of the outcomes of the proposed GCC bootstrap would be useful to understand the strengths and weaknesses of the method.

ACKNOWLEDGEMENTS

We are very grateful to Professor Mario Wüthrich for a discussion of some of the topics addressed in this working paper and for his insights into using the rMSEP for selection of reserve method parameters, and to Pieter Fouche for his thorough review of the working paper.

REFERENCES

Balona, C (2021). tryangle. GitHub.

Balona, C & Richman, R (2020). The actuary and IBNR techniques: A machine learning approach. *SSRN Electronic Journal.*

Benktander, G (1976). An approach to credibility in calculating IBNR for casualty excess reinsurance. *Actuarial Review*, **3**(7): 7–31.

Bornhuetter, RL & Ferguson, RE (1972). The Actuary and IBNR. *Proceedings of the Casualty Actuarial Society*, 59, 181–195.

Bühlmann, H (2016). Historical origin of the Cape Cod claims reserving method. *SSRN Electronic Journal* .

Bühlmann, H (2018). Personal communication. *In* Insurance data science conference, 2018. Zürich.

Bühlmann, H & Straub, E (1983). Estimation of IBNR reserves by the methods chain ladder, Cape Cod and complementary loss ratio (Tech. Rep.). Unpublished.

Carrato, A, Concina, F, Gesmann, M, Murphy, D, Wüthrich, M & Zhang, W (2018). ChainLadder: Statistical methods and models for claims reserving in general insurance.

Dal Moro, E & Lo, J (2014). An industry question: The ultimate and one-year reserving uncertainty for different non-life reserving methodologies. *Astin Bulletin*, **44**(03): 495–499.

England, P & Verrall, R (1999). Analytic and bootstrap estimates of prediction errors in claims reserving. *Insurance: Mathematics and Economics*, **25**(3): 281–293.

England, P & Verrall, R (2006). Predictive distributions of outstanding liabilities in general insurance. *Annals of Actuarial Science*, **1**(2): 221–270.

England, P & Verrall, RJ (2002). Stochastic claims reserving in general insurance. *British Actuarial Journal*, **8**(03): 443–518.

Gluck, SM (1997). Balancing development and trend in loss reserve analysis. *Proceedings of the Casualty Actuarial Society*, 84, 482–532.

IASB (2020). IFRS – IFRS 17 Insurance contracts. International Accounting Standards Board.

Mack, T (1993). Distribution-free calculation of the standard error of chain ladder reserve estimates. *ASTIN Bulletin*, **23**(2): 213–225.

Mack, T (2000). Credible claims reserves: The Benktander method. *Astin Bulletin*, **30**(02): 333–347.

Mack, T (2002). *Schadenversicherungsmathematik 2. Auflage*. Schriftenreihe Angewandte Versicherungsmathematik, DGVM.

Merz, M & Wüthrich, MV (2008). Modelling the claims development result for solvency purposes. CAS E-Forum, 542–568.

Merz, M & Wüthrich, MV (2014). Claims run-off uncertainty: The full picture. *SSRN Electronic Journal.*

Ohlsson, E & Lauzeningks, J (2009). The one-year non-life insurance risk. *Insurance: Mathematics and Economics*, **45**(2): 203–208.

Prudential Authority (2020). Prudential standards – South African Reserve Bank.

Saluz, A (2015). Prediction uncertainties in the Cape Cod reserving method. *Annals of Actuarial Science*, **9**(2): 239–263.

Struzzieri, P, Hussian, PR & Plaza, TP (1998). Using best practices to determine a best reserve estimate. *CAS Forum*, **10121**(212): 353–413.

Taylor, G (2012). *Loss reserving: an actuarial perspective* (Vol. 21). Berlin, Heidelberg: Springer Science & Business Media.

Wüthrich, MV & Merz, M (2008). *Stochastic Claims Reserving Methods in Insurance* (Vol. 435). Hoboken: John Wiley & Sons.