Zeitschrift: Mitteilungen / Schweizerische Aktuarvereinigung = Bulletin / Association Suisse des Actuaires = Bulletin / Swiss Association of Actuaries

Band: - (1996)

Heft: 2

Artikel: An application of the bootstrap method to Bühlmann's classical credibility model

Autor: Giersbergen, Noud van / Dannenburg, Dennis

DOI: https://doi.org/10.5169/seals-551272

NOUD VAN GIERSBERGEN and DENNIS DANNENBURG, Amsterdam

An application of the bootstrap method to Bühlmann's classical credibility model

1 Introduction

After having constructed a theory and transformed it into a mathematical model one usually wants to estimate and test some of the parameters of interest by means of an econometric model. The part of econometrics which deals with the estimation of parameters and testing of hypotheses is called statistical inference.

As an example, suppose that the data $x = (x_1, \ldots, x_p)$ is considered to be a realization of a random sample $X = (X_1, \ldots, X_p)$ from $F(x; \eta)$, where $F(x; \eta)$ denotes some distribution function. The parameter of interest is η, e.g. the mean or some other functional of $F(\cdot)$; it is taken to be a scalar for sake of simplicity. Since we do not observe this parameter value, we are forced to estimate it. Let $\widehat{\eta}_p = h(X)$ denote an estimator of η, which is constructed with the sole aim of providing us with the 'most representative value' of η in the parameter space. Given that the estimator $\widehat{\eta}_p$ is a random variable any formalisation of what we mean by a 'representative value' must be in terms of the distribution of $\widehat{\eta}_p$, say $f_p(\widehat{\eta}_p)$, which is called the *finite sample distribution*. An obvious property to require a 'good' estimator $\widehat{\eta}_p$ of η to satisfy is that $f_p(\widehat{\eta}_p)$ is centred around η. We say that an estimator $\widehat{\eta}_p$ is *unbiased* for η if $E(\widehat{\eta}_p) = \eta$, where E denotes the expectation operator under the finite sample distribution. What this means is that if we calculate $\widehat{\eta}_p$ for each sample and repeat this process infinitely many times, the average of all these estimates will be equal to η. If $E(\widehat{\eta}_p) \neq \eta$, then $\widehat{\eta}_p$ is said to be biased and we refer to $E(\widehat{\eta}_p) - \eta$ as the *bias*. An estimator $\widehat{\eta}_p$ is said to be (weakly) consistent, when $\widehat{\eta}_p$ converges in probability to η, i.e. $\widehat{\eta}_p \xrightarrow{p} \eta$. If an estimator is consistent, the finite sample distribution collapses at the population value η as the number of observations becomes large. Note that a consistent estimator can still be biased. Besides an accurate point estimate, one often wants to obtain a set which covers the true parameter η with high probability. More formally, in interval estimation we construct two functions $g_1(X_1, \ldots, X_p)$ and $g_2(X_1, \ldots, X_p)$ of the sample such that

$$\text{Prob}[\eta \in (g_1, g_2)] = (1 - \alpha)$$

for some given probability α; usually α is small. The interval (g_1, g_2) is called the $(1 - \alpha) \cdot 100\%$ confidence interval. Since η is a parameter, the probability statement is a statement about g_1 and g_2 and not about η. In most of the situations, the construction of the functions g_1 and g_2 is far from trivial. Therefore, we usually employ (large sample) approximations of these functions, which may be accurate or not.

There exists a close relationship between confidence intervals and testing a hypothesis. Suppose that we want to test a hypothesis at a 5 % significance level. Then we have to construct a 95 % confidence interval for the parameter under consideration and see whether the hypothesized value is in the interval. If it is, we do not reject the hypothesis. Otherwise, we reject the hypothesis.

In this paper we focus on the so-called *credibility factor* and its estimator in the classical Bühlmann model. This factor is important for insurance companies, because it is a measure for the heterogeneity in their portfolios. In case it is high, the individual premium of a contract depends highly on the mean over time of the individual contract. If the factor is low, the individual premium will highly depend on the overall mean of the portfolio, which contains all the contracts.

Little is known about the finite sample distribution of a commonly used estimator of the credibility factor, in spite of its widespread usage. All we do know is that this estimator is biased in general. Because the large sample approximations of the finite sample distribution do not take the finite sample characteristics (e.g. bias) into account, it is expected that these approximations are inaccurate. The bootstrap, which was invented by Efron in 1979 and especially designed to mimic the finite sample characteristics of a distribution, claims to provide a better approximation. This claim has been verified for the sample mean in a number of studies, see *inter alia* Singh (1981) and Abramovitch and Singh (1985). The bootstrap confidence intervals can be constructed in several ways. Hence, the aim of this paper is to investigate how well these various bootstrap procedures perform in a particular case of Bühlmann's model. Since theoretical justification using Edgeworth expansions is rather cumbersome, we use Monte Carlo experiments to asses the accuracy of the bootstrap approximations.

The structure of the paper is as follows. In section 2 we define Bühlmann's model, its parameters and their estimators. Section 3 introduces the bootstrap, looks at its asymptotic justification, and defines various bootstrap methods for constructing bootstrap confidence intervals. The Monte Carlo design and results are given in section 4. The last section summarises the results and gives some directions for further research.

2 Bühlmann's Credibility Model

2.1 The model

Consider a full set of observations for the number of claims concerning a period of n years for a portfolio consisting of p contracts. Let these observations be realisations of random variables X_{jt}, $(j = 1, \ldots, p;\ t = 1, \ldots, n)$. Further, define the non-observable random variable Θ_j as the so-called *risk-parameter* belonging to the j-th contract $(j = 1, \ldots, p)$. The classical Bühlmann model postulates the following assumptions for $j = 1, \ldots, p$.

(i) The vectors $(X_{j1}, \ldots, X_{jn}, \Theta_j)'$ and $(X_{k1}, \ldots, X_{kn}, \Theta_k)'$ are stochastically independent for $j \neq k$ $(k = 1, \ldots, p)$.
(ii) For some function $\mu(\cdot)$, $E[X_{jt} \mid \Theta_j] = \mu(\Theta_j)$.
(iii) For some function $\sigma^2(\cdot)$, $\text{Cov}[X_{jt}, X_{js} \mid \Theta_j] = \delta_{st}\sigma^2(\Theta_j)$, where $\delta_{st} = 1$ if $s = t$ and zero otherwise.
(iv) The random variables Θ_j are identically distributed.

Define the following *structural parameters*:

$$m = E[\mu(\Theta_j)]\,, \quad s^2 = E[\sigma^2(\Theta_j)] \quad \text{and} \quad a = \text{Var}[\mu(\Theta_j)]$$

Bühlmann (1967, 1969) shows that the linearised credibility premium in the present model is a weighted average of the estimated individual expected claim size and the expected claim size for the whole portfolio: $P_j = zX_{jw} + (1 - z)m$, where $X_{jw} = 1/n \sum_t X_{jt}$ denotes the sample mean over time of the j-th contract. Here

$$z = \frac{an}{an + s^2}$$

is the *credibility factor*, which only takes values between zero and unity. z is a measure of the heterogeneity in the portfolio for the insurer.

2.2 *Estimators of the structural parameters and the credibility factor*

The mean m is unbiased estimated by the overall mean, denoted by $X_{ww} = 1/p \sum_j X_{jw}$, and s^2 is unbiased estimated by

$$\widehat{s}^2 = \frac{1}{p(n-1)} \sum_{j=1}^{p} \sum_{t=1}^{n} (X_{jt} - X_{jw})^2.$$

An unbiased estimator for a is

$$\widehat{a} = \frac{1}{p-1} \sum_{j=1}^{p} (X_{jw} - X_{ww})^2 - \widehat{s}^2/n\,.$$

From this, a frequently used consistent estimator for z follows:

$$\widehat{z} = \frac{\widehat{a}n}{\widehat{a}n + \widehat{s}^2}\,.$$

Although $\widehat{z}$ is a consistent estimator of z, it is clearly biased in finite samples because $E[f(X)] \neq f(E[X])$ for any (non-trivial) non-linear function $f(\cdot)$. This bias does not have any impact at the top-level of the portfolio, because the total premium income is equal to the number of contracts times the overall mean, i.e. pX_{ww}, for any z. The estimated value of z does, however, have influence on the individual premiums.

3 The Bootstrap Methodology

3.1 *Introduction*

The non-parametric bootstrap, which is applied in this paper, requires not only that the vectors $(X_{j1}, \ldots, X_{jn}, \Theta_j)'$ and $(X_{k1}, \ldots, X_{kn}, \Theta_k)'$ are stochastically independent for $j \neq k$, but also identically distributed. The non-parametric bootstrap may be used when it is hard (or impossible) to find the exact stochastic properties of an estimator without specifying the underlying distribution; see Efron and Tibshirani (1993) and Hall (1992) for recent comprehensive overviews. The idea of the bootstrap method is as follows. Let $X = (X_1, \dots, X_p) \in \mathbb{R}^{p \times n}$ be a random sample of size p from a population with distribution F and let $T(X_1, \dots, X_p; F)$ be the specified random variable of interest, e.g. $z = z(\cdot, F)$, possibly depending on the unknown distribution F. Let $\widehat{F}_p$ be the empirical distribution function of $X_1, \dots, X_p$, i.e. the distribution giving probability mass $1/p$ to each observed contract. The bootstrap method aims to approximate the distribution of $T(X_1, \dots, X_p; F)$ under F by that of $T(X_1^*, \dots, X_p^*; \widehat{F}_p)$ under $\widehat{F}_p$, where $X_1^*, \dots, X_p^*$ denotes a random sample of size p from $\widehat{F}_p$. Although the latter distribution cannot usually be calculated explicitly, it is always possible to approximate it very easily by Monte Carlo simulation, since $\widehat{F}_p$ is available from the sample. So, we can distinguish the following steps for the non-parametric bootstrap:

(i) simulate artificially a random sample $X^* = (X_1^*, \dots, X_p^*)$ from the empirical distribution function $\widehat{F}_p$.

(ii) evaluate T at the bootstrap sample to obtain the bootstrap version of the statistic $T^* = T(X_1^*, \dots, X_p^*; \widehat{F}_p)$.

(iii) repeat steps (i) and (ii) a large number of times, say B, in order to get B realisations of T^*; $T_i^* = T(X_i^*; \widehat{F}_p)$, $i = 1, \dots, B$

Finally, a histogram (or any other estimate of the distribution of T^*) is obtained from T_i^*, $i = 1, \dots, B$; this is the *approximation* to the distribution of T^* which in its turn is the bootstrap *estimator* of the unknown distribution of T.

3.2 *Consistency of the bootstrap*

From Hesselager (1988), it follows that under certain mild regularity conditions the sampling distribution of z can be approximated arbitrarily close by the bootstrap distribution, when the number of contracts p is large enough. In other words, the distribution function

$$G_p(x) = P[\sqrt{p}(\widehat{z} - z) \leq x]$$

is approximated by the conditional distribution

$$G_p^*(x) = P[\sqrt{p}(\widehat{z}^* - \widehat{z}) \leq x \mid X_1, \dots, X_p],$$

such that

$$P\Big[\lim_{p\to\infty}\sup_x \mid G_p(x) - G_p^*(x) \mid = 0\Big] = 1\,.$$

In this case we say that the bootstrap 'works', since the asymptotic conditional bootstrap distribution is equal to the asymptotic unconditional distribution.

Of course, one could also resample both contracts and observations. Then, the observations are resampled in time conditional on a contract that has been drawn. This latter procedure is called *compound* bootstrapping in Hesselager (1988), whereas the former procedure is called *simple* bootstrapping.

Intuitively, one might expect that compound bootstrapping leads to more accurate results. However, Hesselager (1988) shows that compound bootstrapping is not consistent in the number of contracts p only. For consistency, both the number of contracts and the number of years should go to infinity. In practice, n is often small, whereas p can be very large. A pilot study showed that the compound bootstrap gives poor approximations to the finite sample distribution indeed. Therefore, we focus only on the finite sample properties of the simple bootstrap in the Monte Carlo study.

3.3 Various bootstrap estimators

We denote an estimator of a confidence interval for a particular single parameter, say θ, by

$$\hat{I}(\theta) = [\hat{I}_L(\theta), \hat{I}_U(\theta)]\,,$$

where it is made clear that the confidence interval is used for inference with respect to the parameter θ. So, for Bühlmann's model we have $\theta \in \{m, a, s^2, z\}$. Since $z \in [0,1]$, it is convenient and reasonable to take $\hat{I}_L(z) = \max(0, \hat{I}_L(z))$ and $\hat{I}_U(z) = \min(1, \hat{I}_U(z))$.

The first method which we will call the *standard normal method,* is based on bootstrapping the standard deviation of an estimator. More formally, if we assume asymptotic normality for $\widehat{z}$, i.e. $\sqrt{p}(\widehat{z} - z) \rightarrow_d N(0, \sigma^2_{\widehat{z}})$, then the interval given by $[\widehat{z} + \phi_\alpha \sigma_{\widehat{z}}/\sqrt{p}, \widehat{z} + \phi_{1-\alpha}\sigma_{\widehat{z}}/\sqrt{p}]$ yields a $(1-2\alpha)$ confidence interval for z with equal nominal tail probabilities; here ϕ_α denotes the α-th quantile of the standard normal distribution. In general, we do not have an analytical expression for the estimator of $\sigma^2_{\widehat{z}}$. Fortunately, we can always obtain a bootstrap approximation of this estimator

$$\widehat{\sigma}^2_{\widehat{z}^*} = 1/(B-1)\sum_{j=1}^{B}\left[\widehat{z}^*_j - \frac{1}{B}\sum_{i=1}^{B}\widehat{z}^*_i\right]^2$$

The bootstrap approximation for the confidence interval based on the standard normal method is given by

$$\hat{I}^{N^*}(z) \equiv [\hat{I}_L^{N^*}(z), \hat{I}_U^{N^*}(z)] = [\widehat{z} + \phi_\alpha \widehat{\sigma}_{\widehat{z}^*}/\sqrt{p}, \widehat{z} + \phi_{1-\alpha}\widehat{\sigma}_{\widehat{z}^*}/\sqrt{p}]\,.$$

A drawback of this interval is the implicit assumption that the bootstrap distribution is symmetric, so when the finite sample distribution is asymmetric, inference based on this method can be inaccurate.

We will now discuss two other methods to construct asymmetric bootstrap confidence intervals. First Efron's (1979) *percentile method,* which defines the following interval

$$\hat{I}^{P^*}(z) \equiv [\hat{I}_L^{P^*}(z), \hat{I}_U^{P^*}(z)] = [\widehat{z}^*_\alpha, \widehat{z}^*_{1-\alpha}]\,,$$

where $\widehat{z}^*_\alpha$ denotes the α-th quantile from the bootstrap sample, i.e. $\widehat{z}^*_\alpha = \widehat{z}^*_{[\alpha(B+1)]}$ with $\alpha(B+1)$ an integer and $\{\widehat{z}^*_{[i]}; i = 1, \ldots, B\}$ are the corresponding ascending order statistics. Note that this confidence interval will not be symmetric around the point estimate $\widehat{z}$ in general.

The second method proposed by Efron (1981) is called the *bias-corrected* (BC) *percentile method.* This method makes an explicit correction for the median bias of the bootstrap distribution. Let $F^*(t) = P[T^* < t] = \#(T^*_i < t)/B$ denote the approximation of the bootstrap distribution, where $\#(c)$ is the number of occurrences of eventuality c. The confidence interval is given by

$$\begin{aligned}\hat{I}^{BC^*}(z) &\equiv [\hat{I}_L^{BC^*}(z), \hat{I}_U^{BC^*}(z)]\\ &= [F^{*-1}(\Phi(2z_0 + \phi_\alpha)), F^{*-1}(\Phi(2z_0 + \phi_{1-\alpha}))],\end{aligned}$$

where $z_0 = \Phi^{-1}(F^*(\widehat{z}))$, and Φ is the cdf of the standard normal distribution. Note that $F^*(\widehat{z})$ gives the fraction of times that the bootstrap realisations is smaller than the initial estimate $\widehat{z}$. If there is no median bias, then

$F^*(\widehat{z}) = 0.5$ and $z_0 = 0$. In this case the confidence interval based on the bias-corrected percentile method reduces to the confidence interval based on the percentile method. So, the factors $\Phi(2z_0 + \phi_\alpha)$ and $\Phi(2z_0 + \phi_{1-\alpha})$ may be interpreted as a correction for the indices α and $(1-\alpha)$ respectively.

3.4 Rejection frequencies

Let $\widehat{\alpha}(\hat{I}^*_L)$ denote the fraction of times that the lower bound of the interval is greater than the true value. In these cases, the bootstrap confidence interval does not contain the true value. Let $\widehat{\alpha}(\hat{I}^*_U)$ be similarly defined with respect to the upper bound of the interval. So, we have

$$\widehat{\alpha}(\hat{I}^*_L(z)) = \frac{\#(\hat{I}^*_L(z) > z)}{R} \quad \text{and} \quad \widehat{\alpha}(\hat{I}^*_U(z)) = \frac{\#(\hat{I}^*_U(z) < z)}{R},$$

where $\#(c)$ indicates the number of occurrences of eventuality c over R Monte Carlo replications. For the $(1-2\alpha)$ bootstrap confidence interval with equal nominal tail probability, $\widehat{\alpha}(\hat{I}^*_L)$ and $\widehat{\alpha}(\hat{I}^*_U)$ should converge to the nominal size α. If we define $\widehat{\alpha}(\hat{I}^*) \equiv \widehat{\alpha}(\hat{I}^*_L) + \widehat{\alpha}(\hat{I}^*_U)$, then $(1 - \widehat{\alpha}(\hat{I}^*))$ is a Monte Carlo estimator for the actual confidence coefficient of interval $\hat{I}^*$. An estimator for the standard deviation of $\widehat{\alpha}(\cdot)$ over R Monte Carlo replications is given by

$$\sqrt{\widehat{\alpha}(\cdot)[1 - \widehat{\alpha}(\cdot)]/R}\,.$$

4 Monte Carlo Results

Consider the following data generating process (DGP):

$$\begin{aligned} \Theta_j &\sim \mathit{Gamma}(\beta, \lambda) \qquad & j = 1, \ldots, p \\ X_{jt} \mid \Theta_j &\sim \mathit{Poisson}(\Theta_j) \qquad & t = 1, \ldots, n\,. \end{aligned} \tag{1}$$

So, we have assumed a gamma distribution for the risk-parameter Θ_j with parameters β and λ, i.e. $E(\Theta_j) = \beta/\lambda$ and $\mathrm{Var}(\Theta_j) = \beta/\lambda^2$. The number of claims for the j-th contract in period t conditional on Θ_j, i.e. $X_{jt} \mid \Theta_j$, is assumed to be Poisson distributed with parameter equal to Θ_j. It can be shown that in this particular case the unconditional distribution of X_{jt}, is Negative Binomial$(\beta, \lambda/(1+\lambda))$ and that the true credibility factor is given by

$$z = \frac{n}{n+\lambda}\,.$$

This DGP has been chosen in accordance to Dannenburg (1994) and Lemaire (1985) has used this model specification to estimate claim frequencies for automobile-insurance.
The actual performance of the different confidence intervals is investigated by Monte Carlo experiments. We generated 4000 $(= R)$ samples according to the DGP defined by (1) for the parameter values $\beta = 10$, $\lambda = 10$, $p = 50$, 100 and $n = 5$, 20. We choose $\alpha = 0.2$, 0.1, 0.05. The number of bootstrap replications (B) is taken to be 1999 to guaranty that $\alpha(1+B)$ is an integer. The minimum number of bootstrap replications required for obtaining accurate confidence intervals is $B = 1000$, so the amount of replications we have used should be sufficient. All simulations were carried out on a 486-Personal Computer using the matrix programming language Gauss 2.2. Random numbers were obtained by using its 'RNDU' function. Table 1 shows the Monte Carlo results.
From this table, we see that the actual rejection frequencies slowly converge towards the nominal rejection probabilities as the number of contracts, p, increases. There is also some improvement when the nominal significant level α is allowed to increase from 0.05 to 0.20, especially for the standard normal method. Apparently, the bootstrap approximation of the finite sample distribution is better at the center than at the tails. However, no systematic improvement seems to be present if we enlarge the number of periods, n, from 5 to 20. The percentile method yields confidence intervals which are extremely inaccurate, i.e. the reject frequencies of the intervals can be two times as large/small as the nominal rejection probabilities. The bias-corrected percentile method performs better, which is due to the correction of the bias. However, most of the actual rejection frequencies are still significantly different from the nominal rejection probabilities. Inference based on the standard normal method performs best in all cases considered. This is surprising because a symmetric interval around a biased estimator is not expected to perform well. In case $p = 100$, most of the actual reject frequencies are not significantly different than the nominal

Table 1. Monte Carlo results for the rejection frequencies of $(1 - 2\alpha)$ nominal confidence intervals in model (1) based on: the standard normal, the percentile and the bias-corrected percentile method respectively ($R = 4000$, $B = 1999$, $\beta = 10$, $\lambda = 10$).

			Standard normal		Percentile		BC-percentile	
α	p	n	$\widehat{\alpha}(\hat{I}_L^*)$	$\widehat{\alpha}(\hat{I}_U^*)$	$\widehat{\alpha}(\hat{I}_L^*)$	$\widehat{\alpha}(\hat{I}_U^*)$	$\widehat{\alpha}(\hat{I}_L^*)$	$\widehat{\alpha}(\hat{I}_U^*)$
0.05	50	5	0.055	0.032*	0.028*	0.114*	0.046	0.082*
		20	0.051	0.033*	0.021*	0.131*	0.044	0.089*
	100	5	0.058	0.044	0.033*	0.093*	0.050	0.075*
		20	0.054	0.043*	0.029*	0.094*	0.045	0.072*
0.10	50	5	0.099	0.081*	0.061*	0.176*	0.105	0.131*
		20	0.093	0.082*	0.055*	0.188*	0.104	0.143*
	100	5	0.101	0.092	0.072*	0.154*	0.100	0.122*
		20	0.101	0.088*	0.068*	0.160*	0.102	0.121*
0.20	50	5	0.182	0.192	0.145*	0.285*	0.213*	0.222*
		20	0.184	0.192	0.138*	0.295*	0.220*	0.225*
	100	5	0.188	0.199	0.156*	0.266*	0.205	0.213*
		20	0.195	0.199	0.158*	0.269*	0.217*	0.208

$*$ = the approximate 95 % confidence interval for the estimated confidence coefficient does not contain the nominal confidence coefficient.

reject probabilities, so in our Monte Carlo this inference method mimics the behaviour of an exact inference procedure.

To shed some light over the different performance of the various bootstrap intervals, we focus on the Monte Carlo results for $p = 50$ and $n = 10$, which are not reported in Table 1. For a graphical comparison of the estimated bootstrap confidence intervals, we estimated the regression functions

$$E[\hat{I}^*(z) \mid \widehat{z}] = \delta_0 + \sum_{i=1}^{4} \delta_i \widehat{z}^i \tag{2}$$

for $\hat{I}^*(z) \in \{\hat{I}^{N^*}(z), \hat{I}^{P^*}(z), \hat{I}^{BC^*}(z)\}$. These regression functions are flexible enough to capture the salient features of these intervals. Figure 1 shows the estimated systematic components of the regression functions. Since the bounds of the interval $\hat{I}^{P^*}(z)$ are the 5 % and 95 % percentiles of the bootstrap distribution, we observe that the bootstrap distribution

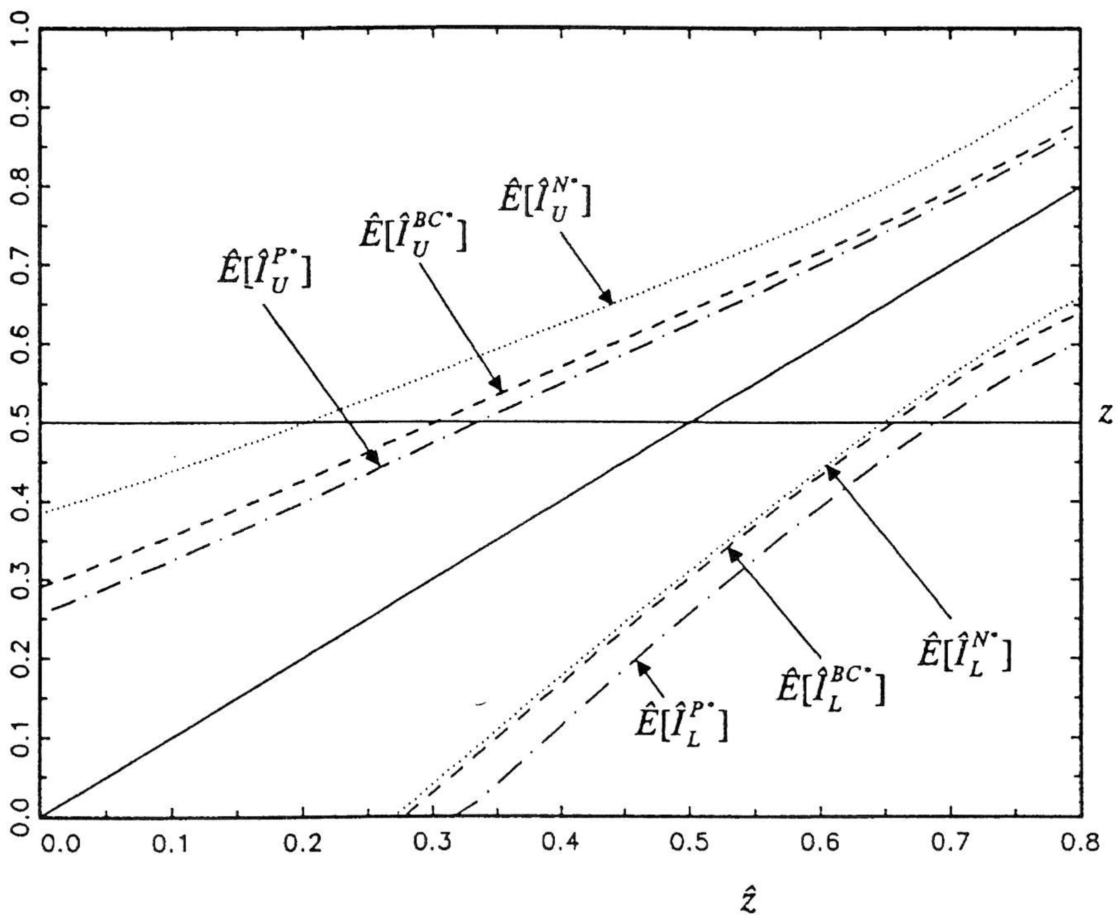

Figure 1. Estimated mean based on formula (2) of various bootstrap confidence bounds in model (1) obtained from 4000 Monte Carlo simulations ($p = 50$, $n = 10$).

is negatively skewed in general. This skewness is mainly due to the fact that $\hat{z}$ is bounded from above by 1 and it causes the percentile method to perform very inaccurate. The bias-corrected percentile method is able to make a correction for the bias. This corresponds to the fact that the bounds are shifted upwards with respect to the bounds of the confidence intervals based on the percentile method. However, due to the negative skewness of the bootstrap distribution, the upper bound has shifted less than the lower bound. Therefore, $\widehat{\alpha}(\hat{I}_U^{BC^*})$ is still significantly greater than its nominal size α. Figure 1 also shows that the length of the intervals decreases as $\hat{z}$ increases. So, the characteristics of the bootstrap distribution vary with respect to $\hat{z}$, e.g. the standard deviation tend to decrease as $\hat{z}$ increases. We tentatively conclude that this phenomenon enables the standard normal method to produce accurate inference even in the presence of skewness.

5 Conclusions

In this paper, we have investigated the small sample performance of various nonparametric bootstrap methods for constructing confidence intervals for the estimator of the credibility factor in the classical Bühlmann model through Monte Carlo experiments. The bootstrap methods are: the standard normal method, the percentile method and the bias-corrected percentile method. The general conclusion is that inference based on the percentile method performs worst. This is due to the negative skewness of the actual and bootstrap distribution of the estimator of the credibility factor. The bias-corrected percentile and the normal method give rise to much more accurate inference. The accuracy of the former method corresponds to the ability of this method to correct for the bias. The latter method emerges as the most promising inference procedure, since in our Monte Carlo experiments most of its actual confidence coefficients were not significantly different from the nominal confidence coefficients for a sample size consisting of 100 contracts. This result is somewhat surprising, since the distribution of the estimator is negatively skewed and the bootstrap confidence intervals based on the normal method are symmetric. However, this phenomenon can be explained by the fact that the characteristics of the bootstrap distribution changes with the realisation of the estimator.

Hopefully, the present study is a step towards a more general resampling based inference method. Issues concerning the sensitivity of the results to the Monte Carlo design, questions like whether we can use the bootstrap to construct accurate confidence intervals for the individual premiums in finite samples, and how well performs Hall's (1988) recommended percentile-t method, which in this case should be based on a iterated bootstrap, all deserve further investigation.

References

Abramovitch, L. and Singh, K. (1985), "Edgeworth Corrected Pivotal Statistics and the Bootstrap", *The Annals of Statistics,* 13, 116–132.

Bühlmann, H. (1967), "Experience Rating and Credibility", *ASTIN-Bulletin,* 4, 199–207.

Bühlmann, H. (1969), "Experience Rating and Credibility", *ASTIN-Bulletin,* 5, 157–165.

Dannenburg, D.R. (1994), "Some Results on the Estimation of the Credibility Factor in the Classical Bühlmann Model", *Insurance: Mathematics and Economics,* 14, 39–50.

Efron, B. (1979), "Bootstrap Methods: Another look at the Jackknife", *The Annals of Statistics,* 7, 1–26.

Efron, B. (1981), "Nonparametric Standard Errors and Confidence Intervals" (with discussion), *Journal of the American Statistical Association,* 82, 171–200.

Efron, B. and Tibshirani, R.J. (1993), *An Introduction to the Bootstrap,* Chapman & Hall.

Hall, P. (1988), "Theoretical Comparison of Bootstrap Confidence Intervals" (with discussion), *The Annals of Statistics,* 16, 927–985.

Hall, P. (1992), *The Bootstrap and Edgeworth Expansion,* Springer-Verlag, New York.

Hesselager, O. (1988), "On the Application of Bootstrap in some Empirical Linear Bayes Estimation Problems", working-paper University of Copenhagen.

Lemaire, J. (1985), *Automobile Insurance: Actuarial Models,* Kluwer-Nijhoff.

Singh, K. (1981), "On the Asymptotic Accuracy of Efron's Bootstrap", *The Annals of Statistics,* 6, 1187–1195.

N. van Giersbergen
D. Dannenburg
Department of Actuarial Science & Econometrics
University of Amsterdam
Roetersstraat 11
NL-1018 WB Amsterdam

Zusammenfassung

Im Artikel werden verschiedene Bootstrap-Methoden für Konfidenzintervalle für den Schätzer des Credibility-Faktors im klassischen Credibility-Modell von Bühlmann näher untersucht. Nach einer kurzen Einführung in das Modell von Bühlmann wird die Idee des Bootstrap-Ansatzes vorgestellt. Anschliessend geben die Autoren drei Methoden an, um Bootstrap-Konfidenzintervalle für den Credibility-Faktor des Bühlmann-Modells zu bestimmen. Die Methoden werden anhand von Monte Carlo-Simulationen illustriert.

Abstract

In the paper, various bootstrap methods for constructing confidence intervals for the estimator of the credibility factor in Bühlmann's classical credibility model are investigated. After a short introduction to the model of Bühlmann, the concept of bootstrapping is presented. In the sequel, the authors give three methods to define bootstrap confidence intervals for the credibility factor. The methods are illustrated by Monte Carlo experiments.

Résumé

Le présent papier traite de la construction d'intervalles de confiance de l'estimateur du facteur de crédibilité du modèle classique de Bühlmann a l'aide de diverses applications de la méthode du bootstrap. Après une brève introduction du modèle de Bühlmann le concept du bootstrap est présenté. Par la suite les auteurs donnent trois méthodes permettant de définir des intervalles de confiance bootstrap du facteur de crédibilité. Ces méthodes sont illustrées à l'aide d'expériences de Monte Carlo.