



UNIVERSITÀ
DEGLI STUDI
DI UDINE

Università degli studi di Udine

A note on simultaneous calibrated prediction intervals for time series

Original

Availability:

This version is available <http://hdl.handle.net/11390/1186142> since 2020-06-19T12:00:47Z

Publisher:

Published

DOI:10.1007/s10260-020-00526-6

Terms of use:

The institutional repository of the University of Udine (<http://air.uniud.it>) is provided by ARIC services. The aim is to enable open access to all the world.

Publisher copyright

(Article begins on next page)

A note on simultaneous calibrated prediction intervals for time series

Giovanni Fonseca

Dipartimento di Scienze Economiche e Statistiche,

Università di Udine

via Tomadini 30/A, I-33100 Udine, Italy.

E-mail: giovanni.fonseca@uniud.it

Federica Giummolè

Dipartimento di Scienze Ambientali, Informatica e Statistica,

Università Ca' Foscari Venezia

via Torino 155, I-30172 Venice, Italy.

E-mail: giummole@unive.it

Paolo Vidoni*

Dipartimento di Scienze Economiche e Statistiche,

Università di Udine

via Tomadini 30/A, I-33100 Udine, Italy.

E-mail: paolo.vidoni@uniud.it

Abstract

This paper deals with simultaneous prediction for time series models. In particular, it presents a simple procedure which gives well-calibrated simultaneous prediction intervals with coverage probability close to the target nominal value. Although the exact computation of the proposed intervals is usually not feasible, an approximation can be easily attained by means of a suitable bootstrap simulation procedure. This new predictive solution is much simpler to compute than those ones already proposed in the literature, based on asymptotic calculations. Applications of the bootstrap calibrated procedure to AR, MA and ARCH models are presented.

Keywords: bootstrap; coverage probability; Monte Carlo simulation; prediction region; simultaneous prediction intervals; time series

*Corresponding author

1 Introduction

In the statistical analysis of time series, a key problem concerns prediction of future values. Although, in the literature, great attention has been received by pointwise predictive solutions, in this paper we deal with the notion of prediction intervals, which explicitly takes account of the uncertainty related to the forecasting procedure. In particular, we assume a parametric statistical model and we follow the frequentist viewpoint, with the aim of constructing prediction intervals having good coverage accuracy.

It is well-known that the estimative or plug-in solution, though simple to derive, is usually not adequate. In fact, it does not properly take account of the sampling variability of the estimated parameters, so that the (conditional) coverage probability of the estimative prediction intervals may substantially differ from the nominal value.

Improved prediction intervals based on complicated asymptotic corrections have been proposed in a general framework by Barndorff-Nielsen and Cox (1996) and, for the case of time series models, by Giummolè and Vidoni (2010) and Vidoni (2004). A calibrating approach has been suggested by Beran (1990) and applied, for example, by Hall et al. (1999), using a suitable bootstrap procedure. Simulation-based prediction intervals for autoregressive processes are considered by Kabaila and Syuhada (2007). Finally, there is an extensive literature on non-parametric bootstrap prediction intervals for autoregressive time series (see, for example, Clements and Kim, 2007; and references therein).

Besides the specification of prediction intervals for a single future observation of a time series, it is certainly of interest, from both the theoretical and the applied point of view, to define a collection of prediction intervals for a set of future observations. In this more challenging framework, the aim is to define a joint prediction region that contains the entire future sequence of realisations with the desired probability. Although the specification of a multivariate prediction region may be quite general, we restrict our attention to joint regions of rectangular form, which are usually considered for forecasting future paths of time series observations (see Alpuim, 1997; Nolan and Ravishanker, 2009; Ravishanker et al., 1991; Wolf and Wunderli, 2015).

In this paper we apply, in the context of time series prediction, a simple procedure based on results in Fonseca et al. (2014), which gives well-calibrated simultaneous prediction intervals with coverage probability equal or close to the target nominal value. Although the exact computation of the proposed intervals is usually not feasible, this can be easily approximated by means of a suitable bootstrap simulation procedure. This new predictive solution is second-order equivalent to those ones based on asymptotic calculations, but it turns out to be much simpler to compute. Applications of the bootstrap calibrated procedure for prediction within AR, MA and ARCH models are presented.

2 Simultaneous calibrated prediction intervals

Given a discrete-time stochastic process $\{Y_t\}_{t \geq 1}$, we assume that the random vector $Y = (Y_1, \dots, Y_n)$, $n > 1$, is observable, while $Z = (Z_1, \dots, Z_m) = (Y_{n+1}, \dots, Y_{n+m})$, $m \geq 1$, is a random vector corresponding to a future or not yet available m -dimensional sequence of observations. Y and Z are continuous random vectors and we assume that $g(z|y; \theta)$ and $G(z|y; \theta)$, the conditional density and distribution function of Z given $Y = y$, are specified up to a d -dimensional unknown parameter $\theta \in \Theta \subseteq \mathbb{R}^d$. In the presence of a transitive statistic U (see, for example, Barndorff-Nielsen and Cox, 1996), y is substituted by the observed value u of U .

Given an observed sample $y = (y_1, \dots, y_n)$, a system of simultaneous α -prediction limits for vector Z is a set of functions $c_\alpha^j(y)$, $j = 1, \dots, m$, such that, exactly or approximately,

$$P_{Y,Z}\{Z_j \leq c_\alpha^j(Y), j = 1, \dots, m; \theta\} = \alpha, \quad (1)$$

for every $\theta \in \Theta$ and for any fixed $\alpha \in (0, 1)$. In the presence of a finite dimensional transitive statistics, we usually consider the conditional coverage probability

$$P_{Y,Z|U}\{Z_j \leq c_\alpha^j(Y), j = 1, \dots, m | U = u; \theta\} = \alpha. \quad (2)$$

Obviously, a system of prediction limits satisfying (2) also satisfies condition (1).

An α -level joint prediction region of rectangular form is readily obtained by specifying two suitable systems of lower and upper simultaneous prediction limits. As we can see, for instance, in Alpuim (1997), Nolan and Ravishanker (2009), Ravishanker et al. (1991) and Wolf and Wunderli (2015), rectangular prediction regions are usually defined, in a time series context, using simultaneous prediction limits for Z defined as

$$z_{j,\alpha} = z_{j,\alpha}(Y; \theta) = P_j + h_\alpha(\theta) se_j(\theta), \quad j = 1, \dots, m, \quad (3)$$

evaluated at $\theta = \hat{\theta}$, where $\hat{\theta} = \hat{\theta}(Y)$ is the maximum likelihood estimator for θ or an asymptotically equivalent alternative. Here $P_j = P_j(Y; \theta)$ is a suitable unbiased point predictor for Z_j , such that $E_{Z_j|Y}(Z_j - P_j|Y = y; \theta) = 0$, with conditional prediction standard error $se_j(\theta) = \sqrt{V_{Z_j|Y}(Z_j - P_j|Y = y; \theta)}$. Moreover, $h_\alpha(\theta)$ is a quantity satisfying

$$P_{Z|Y}\{\mathcal{E}_j \leq h_\alpha(\theta), j = 1, \dots, m|Y = y; \theta\} = F\{h_\alpha(\theta), \dots, h_\alpha(\theta)|y; \theta\} = \alpha,$$

with $\mathcal{E}_j = (Z_j - P_j)/se_j(\theta)$, $j = 1, \dots, m$, the standardised forecast errors, with joint distribution function $F(e_1, \dots, e_m|y; \theta)$, conditional on $Y = y$.

For stationary linear models, we usually consider the optimal predictors given by $P_j = E_{Z_j|Y}(Z_j|Y; \theta)$, $j = 1, \dots, m$. Indeed, with this choice for the point forecasts, provided that we have a linear or a Gaussian process, the vector of the standardized forecasts errors $(\mathcal{E}_1, \dots, \mathcal{E}_m)$ is independent of Y .

Alternative systems of simultaneous α -prediction limits involve the specification of a sequence of marginal prediction limits for each future component Z_j , $j = 1, \dots, m$, or a projection on the axes of \mathbb{R}^m of a suitable multivariate prediction region for Z , such as the Scheffé joint prediction region (see, for example, Kim, 1999; Wolf and Wunderli, 2015). However, both solutions present a coverage probability which may substantially differ from the target nominal value α . For this reason, it is better to define systems of simultaneous prediction limits which are designed to be of rectangular form to begin with and which account for the dependencies among the components of vector Z . Wolf and Wunderli (2015) introduce a system of simultaneous prediction limits similar to (3), having a coverage probability asymptotically equal to α . Anyway, their method applies to

large samples, when the coverage error associated to estimative simultaneous prediction limits $\hat{z}_{j,\alpha} = z_{j,\alpha}(Y; \hat{\theta})$, $j = 1, \dots, m$, can be disregarded. Our aim here is to calibrate the estimative solution, thus providing simultaneous prediction limits with coverage probability closer to the nominal value, even for a small or moderate sample size.

In order to compute the prediction limits specified in (3), we need a vector of unbiased point predictors $P = (P_1, \dots, P_m)$, the associated vector of prediction standard errors $se(\theta) = \{se_1(\theta), \dots, se_m(\theta)\}$ and the quantity $h_\alpha(\theta) = \varphi^{-1}(\alpha|y; \theta)$, where $\varphi^{-1}(\cdot|y; \theta)$ is the inverse of function $\varphi(x|y; \theta) = F(x, \dots, x|y; \theta)$, which corresponds to the conditional distribution function $F(e_1, \dots, e_m|y; \theta)$ constrained to $\{(e_1, \dots, e_m) \in \mathbb{R}^m | e_1 = \dots = e_m = x\}$. It can be useful noticing that $\varphi(x|y; \theta)$ is the distribution function of the maximum of the standardised forecast errors, conditioned on $Y = y$. As suggested in Wolf and Wunderli (2015), we can consider the generalised family-wise error rate, k -FWE, instead of (1) for specifying the coverage of simultaneous prediction limits. In that case, α is the probability that at most k future observations lay above the corresponding prediction limit. Then the multiplier $h_\alpha(\theta)$ in (3) must be substituted with the α -quantile of the $k + 1$ largest standardised forecast error, that corresponds to the $(m - k)$ -th order statistic, $\mathcal{E}_{(m-k)}$. In this paper, we consider the coverage probability defined in (1), which is the limit case of k -FWE obtained for $k = 0$.

Our proposal consists of calibrating the multiplier $h_\alpha(\theta)$ in (3) using the following simple procedure, borrowed from Fonseca et al. (2014).

The (unconditional) coverage probability of the estimative simultaneous prediction limits $\hat{z}_{j,\alpha} = z_{j,\alpha}(Y; \hat{\theta})$, $j = 1, \dots, m$, corresponds to

$$\begin{aligned} P_{Y,Z}\{Z_j \leq \hat{z}_{j,\alpha}, j = 1, \dots, m; \theta\} &= E_Y\{P_{Z|Y}(Z_j \leq \hat{z}_{j,\alpha}, j = 1, \dots, m|Y; \theta); \theta\} \\ &= E_Y[P_{Z|Y}\{\mathcal{E}_j \leq (\hat{z}_{j,\alpha} - P_j)/se_j(\theta), j = 1, \dots, m|Y; \theta\}; \theta] \\ &= E_Y[F\{a_1 + h_\alpha(\hat{\theta})b_1, \dots, a_m + h_\alpha(\hat{\theta})b_m|Y; \theta\}; \theta] = D(\alpha, \theta), \end{aligned}$$

where $a_j = a_j(Y, \theta) = (\hat{P}_j - P_j)/se_j(\theta)$, with $\hat{P}_j = P_j(Y; \hat{\theta})$, and $b_j = b_j(Y, \theta) = se_j(\hat{\theta})/se_j(\theta)$, $j = 1, \dots, m$. Unfortunately, as already mentioned, the quantity $D(\alpha, \theta)$

may differ from the target value α by a term of order $O(n^{-1})$, which can be substantial for a small sample size n .

Following the calibrating procedure proposed in Fonseca et al. (2014) for univariate prediction limits, we may consider function

$$\varphi_c(x|y; \hat{\theta}, \theta) = D\{F(x, \dots, x|y; \hat{\theta}), \theta\} = D\{\varphi(x|y; \hat{\theta}), \theta\} \quad (4)$$

instead of $\varphi(x|y; \hat{\theta})$, in order to specify the quantity

$$h_\alpha^c(\hat{\theta}, \theta) = \varphi_c^{-1}(\alpha|y; \hat{\theta}, \theta) = \varphi^{-1}\{D^{-1}(\alpha, \theta)|y; \hat{\theta}\} = h_{D^{-1}(\alpha, \theta)}(\hat{\theta}), \quad (5)$$

with $\varphi_c^{-1}(\cdot|y; \hat{\theta}, \theta)$ and $D^{-1}(\cdot, \theta)$ the inverse functions of $\varphi_c(\cdot|y; \hat{\theta}, \theta)$ and $D(\cdot, \theta)$, respectively. It is easy to show that the calibrated simultaneous prediction limits thus obtained, namely

$$z_{j, \alpha}^c(Y; \hat{\theta}, \theta) = \hat{P}_j + h_\alpha^c(\hat{\theta}, \theta) se_j(\hat{\theta}), \quad j = 1, \dots, m,$$

present a coverage probability equal to the target nominal value α . Indeed the specification of quantities $h_\alpha^c(\hat{\theta}, \theta)$ from (5) determines simultaneous prediction limits satisfying relation (1) exactly for all $\alpha \in (0, 1)$.

As mentioned before, function (4) depends on the unknown θ and it is in fact not useful. Provided that the coverage probability $D(\alpha, \theta)$ may be calculated, at least to order $O(n^{-1})$, a useful surrogate for (4) is the corresponding plug-in estimator $\varphi_c(x|y; \hat{\theta}, \hat{\theta})$. The associated α -level quantile $h_\alpha^c(\hat{\theta}, \hat{\theta})$ enables the specification of calibrated simultaneous prediction limits with a coverage error term reduced to order $o(n^{-1})$.

Whenever a closed form (approximate) expression for $D(\alpha, \theta)$ is not available, we may consider a suitable parametric bootstrap estimator for function $\varphi_c(x|y; \hat{\theta}, \theta)$. Let y^b , $b = 1, \dots, B$, be parametric bootstrap samples generated from the estimative distribution of the data and let $\hat{\theta}^b$, $b = 1, \dots, B$, be the corresponding maximum likelihood estimates. Since $D(\alpha, \theta)$ is an expectation, we define the following bootstrap estimator for (4):

$$\varphi_c^b(x|y; \hat{\theta}) = \frac{1}{B} \sum_{b=1}^B F\{\hat{a}_1^b + h_\alpha(\hat{\theta}^b) \hat{b}_1^b, \dots, \hat{a}_m^b + h_\alpha(\hat{\theta}^b) \hat{b}_m^b | y; \hat{\theta}\} |_{\alpha = \varphi(x|y; \hat{\theta})},$$

where $\hat{a}_j^b = (\hat{P}_j^b - \hat{P}_j)/se_j(\hat{\theta})$, with $\hat{P}_j^b = P_j(Y; \hat{\theta}^b)$, and $\hat{b}_j^b = se_j(\hat{\theta}^b)/se_j(\hat{\theta})$, $j = 1, \dots, m$. In this case, the associated α -level quantile permits the definition of a system of simultaneous prediction limits with coverage probability equal to α , apart from an error term depending on the efficiency of the bootstrap procedure.

In the presence of a transitive statistic U , we may define a similar calibrating procedure improving the conditional coverage probability (2). In this case, whenever a closed form (approximate) expression for the conditional coverage probability of the estimative solution is not available, we have to consider parametric bootstrap samples generated from the estimative distribution of the data given $U = u$.

3 Examples and simulations

In this section we present some examples of application of the proposed method to autoregressive, moving average and autoregressive conditional heteroschedastic Gaussian processes. Assuming the normal distribution for the innovations implies that the distribution of the standardised forecast errors, $F(e_1, \dots, e_m|y; \theta)$, is Gaussian. As already pointed out, $\varphi(x|y; \theta)$ is the distribution function of the maximum of the standardised forecast errors and it is in general not explicitly available. However, in the case of Gaussian processes, suitable functions for calculating the values of $\varphi(x|y; \theta)$ and its quantiles are available in most of the commonly used statistical packages. Moreover, for $m = 2$, it is a known result that the distribution of the maximum of a bivariate normal vector has a skew normal distribution, see Arellano-Valle and Genton (2008) and Nadarajah and Kotz (2008). These considerations may help fastening the simulations.

3.1 Autoregressive models

Let $\{Y_t\}_{t \geq 1}$ be a first-order Gaussian autoregressive process with

$$Y_t = \mu + \rho(Y_{t-1} - \mu) + \epsilon_t, \quad t \geq 1,$$

where μ and ρ are unknown parameters and $\{\epsilon_t\}_{t \geq 1}$ is a sequence of independent Gaussian random variables with zero mean and unknown variance σ^2 . We assume $|\rho| < 1$ so that the process is stationary and $Y_0 = y_0$ known.

The observable random vector is $Y = (Y_1, \dots, Y_n)$ and the next m realisations of the process are $Z = (Y_{n+1}, \dots, Y_{n+m})$. The conditional distribution of Z given $Y = y$ is m -dimensional Gaussian with mean vector $\mu^{Z|Y} = (\mu_{n+1}, \dots, \mu_{n+m})^T$, where $\mu_{n+1} = (1 - \rho)\mu + \rho y_n$ and $\mu_{n+j} = (1 - \rho)\mu + \rho\mu_{n+j-1}$, $j = 2, \dots, m$, and variance-covariance matrix $\Sigma^{Z|Y} = (\sigma_{ij})$, where $\sigma_{ij} = \sigma^2 \rho^{|i-j|}$, $i \neq j$, and $\sigma_{ii} = \sigma^2 \sum_{k=1}^i \rho^{k-1}$, $i, j = 1, \dots, m$.

We take the conditional expectations as point predictors, $P_j = \mu_{n+j}$, $j = 1, \dots, m$. Thus, the standardised forecast errors are independent of Y with a m -dimensional Gaussian distribution with zero mean vector and variance-covariance matrix equal to the correlation matrix of Z .

Then, $\varphi(x|y; \theta)$ is the distribution function of the maximum of a m -dimensional normal vector. The corresponding quantiles are provided by suitable functions through numerical approximation. For $m = 2$, $\varphi(x|y; \theta)$ is the distribution function of a skew normal random variable with skewness parameter $\sqrt{1 + \rho^2} - \rho$ and its quantiles are readily available.

A simulation study shows the performance of the proposed solution in comparison with the estimative one. Here, $U = Y_n$ is a transitive statistic and we evaluate prediction limits by means of their coverage probability conditioned on the observed value y_n of Y_n . For computing the conditional coverage probabilities we use the simulation technique presented in Kabaila (1999), keeping the last observed value y_n fixed and $y_0 = 0$. The parameter μ is assumed to be known and equal to 0 and the sample size is $n = 20$. The results are collected in Table 1 for $m = 2$ and different values of $y_n = -1, 0, 1$ and in Table 2 for $m = 1, 2, 5$ and $y_n = 1$.

TABLE 1 AND 2 HERE

It can be seen that the bootstrap calibrated simultaneous prediction limits have coverage probability that closely approximates the target value and remarkably improves on

the estimative one. The results do not seem to depend much on the value of the transitive statistic y_n . Instead, they depend on the value of the autoregressive coefficient ρ : both the estimative and the calibrated limits work worse when ρ is closer to the frontier of stationarity. Moreover, it can be noticed that, as m increases, the improvement on the estimative is less evident, in particular for higher target coverage. This is probably due to the particular form of the considered intervals and to the fact that we calibrate a single multiplier for all the components of the future sample.

Figure 1 represents a trajectory of a AR(1) process with $\rho = 0.5$ and $\sigma^2 = 1$, together with the point predictions, the estimative and the bootstrap calibrated prediction bands of level $\alpha = 0.9$, for $m = 5$ future observations. Prediction is based on $n = 20$ past observations with 500 bootstrap replications. It is important noticing that the resulting prediction intervals are not centered on the point prediction, being the distribution of the maximum of the standardised forecast errors strongly skewed. Of course, this becomes more evident as the number of future observations increases. To partially overcome this problem, we have used quantiles of level 0.01 and 0.91 for calculating the lower and the upper limits, respectively. It can be seen that the lower bootstrap calibrated prediction limit is very close to the estimative one. Instead the correction becomes substantial in the upper limit.

FIGURE 1 HERE

3.2 Moving average models

Let $\{Y_t\}_{t \geq 1}$ be a first-order Gaussian moving average process where

$$Y_t = \mu + \epsilon_t + \rho\epsilon_{t-1}, \quad t \geq 1,$$

with $\epsilon_t \sim N(0, \sigma^2)$, $t \geq 0$, independent Gaussian random variables. We want to predict $Z = (Y_{n+1}, \dots, Y_{n+m})$ on the basis of an observed sample $y = (y_1, \dots, y_n)$ from $Y = (Y_1, \dots, Y_n)$. We assume $|\rho| < 1$ to ensure invertibility. The conditional distribution of Z given Y is $Z|Y \sim N(\mu^{Z|Y}, \Sigma^{Z|Y})$, where $\mu^{Z|Y} = (\mu_{n+1}, \dots, \mu_{n+m})^T$, with $\mu_{n+1} =$

$\mu + \sum_{i=1}^n (-1)^{i+1} \rho^i (Y_{n+1-i} - \mu)$, and $\mu_{n+j} = \mu$, $j = 2, \dots, m$, and $\Sigma^{Z|Y} = (\sigma_{ij})$, with $\sigma_{11} = \sigma^2$, $\sigma_{jj} = \sigma^2(1 + \rho^2)$, $j = 2, \dots, m$, and $\sigma_{ij} = \sigma^2 \rho$, if $|i - j| = 1$, and $\sigma_{ij} = 0$ otherwise.

A simulation study shows the performance of the bootstrap calibrated predictive solution in comparison with the estimative solution. Coverage probabilities for estimative and bootstrap calibrated prediction limits of levels $\alpha = 0.9, 0.95$ are calculated, assuming the parameter μ to be known and equal 0 and the sample size $n = 20$. The results are collected in Table 3 and confirm the superiority of the bootstrap calibrated prediction limits over the estimative ones.

TABLE 3 HERE

Figure 2 represents a trajectory of a MA(1) process with $\rho = 0.5$ and $\sigma^2 = 1$, together with the point predictions, the estimative and the bootstrap calibrated prediction bands of level $\alpha = 0.9$, for $m = 5$ future observations. Prediction is based on $n = 20$ past observations with 500 bootstrap replications. As already observed in the AR(1) example, the resulting prediction intervals are not centered on the point prediction, being the distribution of the maximum of the standardised forecast errors strongly skewed. This is why we have used quantiles of level 0.01 and 0.91 for calculating the lower and the upper limits, respectively.

FIGURE 2 HERE

3.3 ARCH models

Finally, we present an example of non-linear model. Let $\{Y_t\}_{t \geq 1}$ be a first-order autoregressive conditional heteroschedastic Gaussian process with

$$Y_t = \sqrt{\beta + \gamma Y_{t-1}^2} \epsilon_t, \quad t \geq 1,$$

where β and γ are unknown parameters and $\{\epsilon_t\}_{t \geq 1}$ is a sequence of independent standard Gaussian random variables. We assume $\beta > 0$ and $\gamma \in [0, 3.56]$ to ensure strict stationarity. The unknown parameter is $\theta = (\beta, \gamma)$ and likelihood inference is conditioned on

$Y_0 = y_0$, with y_0 known. The observable random vector is $Y = (Y_1, \dots, Y_n)$ and the next m realizations of the process are $Z = (Y_{n+1}, \dots, Y_{n+m})$. The conditional distribution of Z given $Y = y$ is m -dimensional Gaussian with zero mean vector and diagonal variance-covariance matrix $\Sigma^{Z|Y} = \text{diag}(\sigma_{jj})$, $\sigma_{jj} = \beta \sum_{k=1}^j \gamma^{k-1} + \gamma^j y_n^2$, $j = 1, \dots, m$. Indeed, Y_n is a transitive statistic and we evaluate prediction limits by means of their coverage probability conditioned on the observed value y_n of Y_n .

We take $P_j = E[Y_{n+j}|Y_n] = 0$ as a point predictor for Z_j , $j = 1, \dots, m$. The standardized forecast errors, $\mathcal{E}_j = Z_j/se_j(\theta)$, are independent of Y_n , with $se_j(\theta)^2 = E(Z_j^2|Y_n = y_n; \theta) = \sigma_{jj}$, $j = 1, \dots, m$. Thus, the conditional distribution of the standardised forecast errors is standard normal, i.e. $F(e_1, \dots, e_m|y; \theta) = \prod_{j=1}^m \Phi(e_j)$, where $\Phi(\cdot)$ denotes the univariate standard normal distribution function. It is important noticing that F does not depend on the observed value of the transitive statistic y_n , nor on the unknown parameter θ . Furthermore, the quantity $h_\alpha(\theta)$ is nothing but the quantile of level $\alpha^{1/m}$ of a standard normal distribution, since $F(h_\alpha(\theta), \dots, h_\alpha(\theta)|y; \theta) = \Phi(h_\alpha(\theta))^m = \alpha$.

A simulation study shows the performance of the bootstrap calibrated predictive solution in comparison with the estimative solution. The bootstrap samples are generated keeping the observed value of the transitive statistic fixed to y_n . Conditional coverage probabilities for estimative and bootstrap calibrated prediction limits of different levels are calculated by means of the simulation technique presented in Kabaila (1999). In this case the sample size is taken to be $n = 50$ for more stable estimates of the parameters in the Monte Carlo replications. As a consequence, the improvement on the estimative solution is less evident. Anyway, the results collected in Table 4 still confirm the superiority of the bootstrap calibrated prediction limits over the estimative ones.

TABLE 4 HERE

Figure 3 represents a trajectory of a ARCH(1) process with $\beta = 0.5$, $\gamma = 1$, together with the point predictions, the estimative and the bootstrap calibrated prediction bands of level $\alpha = 0.9$, for $m = 5$ future observations. Prediction is based on $n = 50$ past

observations with 500 bootstrap replications. Again, we have used quantiles of level 0.01 and 0.91 for calculating the lower and the upper limits, respectively. In fact, with this choice the resulting prediction intervals include the point prediction, that is constantly equal to 0. It is important noticing that for this model the standardised forecast errors have independent standard normal distributions. Thus, for instance, the quantile of level 0.05 of their maximum is positive, giving rise to positive lower prediction bounds. The effect is milder when the standardised forecast errors are dependent, as in AR and MA models. This is, of course, a drawback of this kind of prediction regions, and it is even more evident when the number of future observations increases.

FIGURE 3 HERE

4 Final remarks

In this paper we have proposed a method for obtaining rectangular prediction regions with a fixed joint coverage probability. The proposed method has been applied to autoregressive, moving average and autoregressive conditional heteroschedastic Gaussian processes, showing the improvement with respect to the usual estimative approach.

For non-Gaussian processes (see for instance Alpuim, 1997; Nolan and Ravishanker, 2009), whenever function $F(e_1, \dots, e_m|y; \theta)$ is not known, a suitable bootstrap estimate has to be considered. Similarly, it is possible to approximate $\varphi(x|y; \theta)$ and its quantiles. This, together with the calibration, implies a double bootstrap procedure that can be easily carried out in the applications. Nonetheless, when many bootstrap replications are required, the computations can be time demanding.

The proposed method calibrates the estimative quantiles of the function $\varphi(x|y; \theta)$ and it is particularly effective when the size of the observed sample is small. In the presence of a large sample size, the effect of the calibration is less evident and the coverage associated to the estimative quantiles is already very precise. Indeed, Wolf and Wunderli (2015) obtain good results without need of calibration, using sample sizes $n = 100, 400$.

As already mentioned, the notion of coverage probability can be extended by considering the generalised family-wise error rate, k -FWE, see Wolf and Wunderli (2015). In that case, α is the probability that at most k future observations lay above the corresponding prediction limit. Our proposal can be easily extended to this more general setting by simply substituting $\varphi(x|y; \theta)$ and its quantiles with the distribution function and the quantiles of the $(m - k)$ -th order statistic for the standardized forecast errors, $\mathcal{E}_{(m-k)}$. Anyway, we have not considered this extension here because with a small sample size the number m of future observations is also small and the notion of k -FWE is not sensible anymore. Moreover, as we have seen in the examples, when the number of future observations increases the distribution $\varphi(x|y; \theta)$ becomes more skewed and the simultaneous prediction intervals are not centred on the point predictors. To overcome this drawback, we suggest choosing non symmetric levels for lower and upper quantiles. Using k -FWE is, of course, another way for obtaining lower and upper prediction limits that are more symmetric with respect to the point predictor, being the distribution of the $(m - k)$ -th order statistic less skewed than that of the maximum. Finally, it is possible to obtain prediction bands that are centred on the point predictor by considering marginal prediction intervals for each future component, conditioned on the observed values. Anyway, the coverage of the resulting multivariate prediction region can not be easily controlled without increasing the coverage and thus the length of each marginal interval.

Acknowledgements

This research was partially supported by the University of Udine under the PRID2017 research grants and by the Italian Ministry for University and Research under the PRIN2015 grant No.2015EASZFS_003. Simulation studies were performed on SCSCF, a multiprocessor cluster system owned by Università Ca' Foscari Venezia. This is a post-peer-review, pre-copyedit of an article published in Statistical Methods & Applications. The final authenticated version is available online at <https://link.springer.com/article/10.1007/s10260-020-00526-6>.

References

- Alpuim MT (1997) One-sided simultaneous prediction intervals for AR(1) and MA(1) processes with exponential innovations. *J Forecasting* 16:19–35.
- Arellano-Valle RB, Genton MG (2008) On the exact distribution of the maximum of absolutely continuous dependent random variables. *Stat Probabil Lett* 78:27–35.
- Barndorff-Nielsen OE, Cox DR (1996) Prediction and asymptotics. *Bernoulli* 2:319–340.
- Beran R (1990) Calibrating prediction regions. *J Amer Statist Assoc* 85:715–723.
- Clements MP, Kim JH (2007) Bootstrap prediction intervals for autoregressive time series. *Comput Statist Data Anal* 51:3580–3594.
- Fonseca G, Giummolè F, Vidoni P (2014) Calibrating predictive distributions. *J Stat Comput Sim* 84:373–383.
- Giummolè F, Vidoni P (2010) Improved prediction limits for a general class of Gaussian models. *J Time Series Anal* 31:483–493.
- Hall P, Peng L, Tajvidi N (1999) On prediction intervals based on predictive likelihood or bootstrap methods. *Biometrika* 86:871–880.
- Kabaila P (1999) An efficient simulation method for the computation of a class of conditional expectations. *Aust N Z J Stat* 41:331–336.
- Kabaila P, Syuhada K (2007) Improved prediction limits for AR(p) and ARCH(p) processes. *J Time Series Anal* 29:213–223.
- Kim J (1999) Asymptotic and bootstrap prediction regions for vector autoregressions. *Int J Forecasting* 15:393–403.

Nadarajah S, Kotz S (2008) Exact distribution of the Max/Min of two Gaussian random variables. *IEEE T VLSI Syst* 16:210–212.

Nolan JP, Ravishanker N (2009) Simultaneous prediction intervals for ARMA processes with stable innovations. *J Forecasting* 28:235–246.

Ravishanker N, Wu LSY, Glaz J (1991) Multiple prediction intervals for time series: comparison of simultaneous and marginal intervals. *J Forecasting* 10:445–463.

Vidoni P (2004) Improved prediction intervals for stochastic process models. *J Time Series Anal* 25:137–154.

Wolf M, Wunderli D (2015) Bootstrap joint prediction regions. *J Time Series Anal* 36:352–376.

α	y_n	ρ	Estimative	Bootstrap
0.90	-1	0.5	0.884	0.901
		0.8	0.890	0.900
	0	0.5	0.873	0.901
		0.8	0.868	0.898
	1	0.5	0.871	0.899
		0.8	0.848	0.891
0.95	-1	0.5	0.932	0.945
		0.8	0.936	0.922
	0	0.5	0.923	0.948
		0.8	0.921	0.948
	1	0.5	0.924	0.948
		0.8	0.908	0.944

Table 1: AR(1) Gaussian model. Conditional coverage probabilities for estimative and bootstrap calibrated prediction limits of level $\alpha = 0.9, 0.95$, conditioned on $y_n = -1, 0, 1$; $\mu = 0$ known, $\rho = 0.5, 0.8$, $\sigma^2 = 1$, $y_0 = 0$, $n = 20$ and $m = 2$. Estimation is based on 1,000 Monte Carlo replications. Bootstrap procedure is based on 500 bootstrap samples. Estimated standard errors are always smaller than 0.01.

m	α	ρ	Estimative	Bootstrap
1	0.90	0.5	0.875	0.894
		0.8	0.862	0.864
	0.95	0.5	0.925	0.947
		0.8	0.924	0.948
2	0.90	0.5	0.871	0.899
		0.8	0.848	0.891
	0.95	0.5	0.924	0.948
		0.8	0.908	0.944
5	0.90	0.5	0.862	0.897
		0.8	0.809	0.861
	0.95	0.5	0.915	0.922
		0.8	0.877	0.915

Table 2: AR(1) Gaussian model. Conditional coverage probabilities for estimative and bootstrap calibrated prediction limits of level $\alpha = 0.9, 0.95$, conditioned on $y_n = 1$; $\mu = 0$ known, $\rho = 0.5, 0.8$, $\sigma^2 = 1$, $y_0 = 0$, $n = 20$ and $m = 1, 2, 5$. Estimation is based on 1,000 Monte Carlo replications. Bootstrap procedure is based on 500 bootstrap samples. Estimated standard errors are always smaller than 0.01.

m	α	ρ	Estimative	Bootstrap
1	0.9	0.5	0.871	0.893
		0.8	0.861	0.873
	0.95	0.5	0.924	0.944
		0.8	0.917	0.925
2	0.9	0.5	0.862	0.895
		0.8	0.861	0.892
	0.95	0.5	0.918	0.939
		0.8	0.918	0.933
5	0.9	0.5	0.854	0.901
		0.8	0.851	0.898
	0.95	0.5	0.912	0.933
		0.8	0.904	0.923

Table 3: MA(1) Gaussian model. Coverage probabilities for estimative and bootstrap calibrated prediction limits of level $\alpha = 0.9, 0.95$; $\mu = 0$ known, $\rho = 0.5, 0.8$, $\sigma^2 = 1$, $y_0 = 0$, $n = 20$ and $m = 1, 2, 5$. Estimation is based on 1,000 Monte Carlo replications. Bootstrap procedure is based on 500 bootstrap samples. Estimated standard errors are always smaller than 0.01.

m	α	β	γ	Estimative	Bootstrap
1	0.9	0.5	1	0.893	0.899
		0.2	0.9	0.886	0.894
		1	2	0.899	0.894
	0.95	0.5	1	0.942	0.949
		0.2	0.9	0.939	0.947
		1	2	0.948	0.941
2	0.9	0.5	1	0.877	0.894
		0.2	0.9	0.874	0.891
		1	2	0.881	0.887
	0.95	0.5	1	0.928	0.933
		0.2	0.9	0.922	0.934
		1	2	0.930	0.937
5	0.9	0.5	1	0.828	0.864
		0.2	0.9	0.822	0.840
		1	2	0.780	0.856
	0.95	0.5	1	0.885	0.909
		0.2	0.9	0.875	0.887
		1	2	0.849	0.901

Table 4: ARCH(1) Gaussian model. Conditional coverage probabilities for estimative and bootstrap calibrated prediction limits of level $\alpha = 0.9, 0.95$, conditioned on $y_n = 1$; $\beta = 0.5, 0.2, 1$, $\gamma = 1, 0.9, 2$, $y_0 = 0$, $n = 50$ and $m = 1, 2, 5$. Estimation is based on 1,000 Monte Carlo replications. Bootstrap procedure is based on 500 bootstrap samples. Estimated standard errors are always smaller than 0.01.

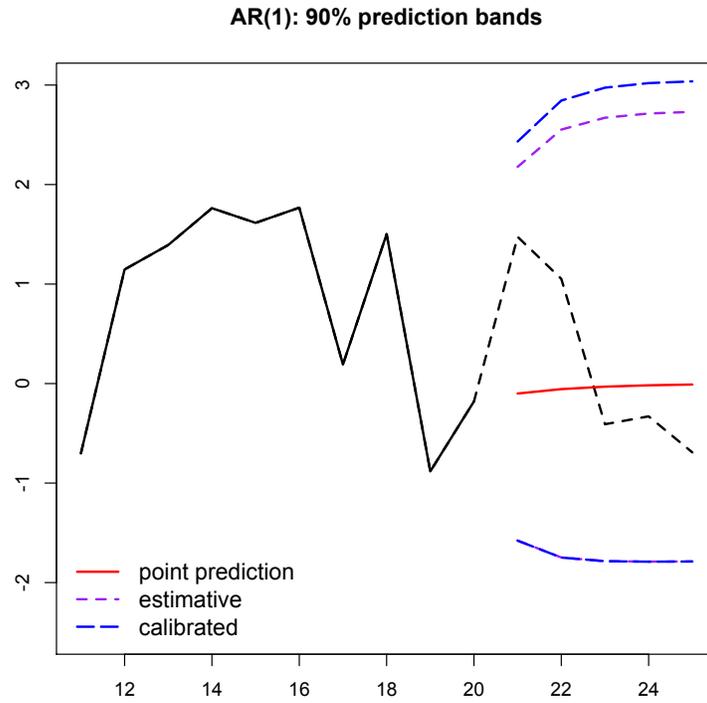


Figure 1: AR(1) Gaussian model. Plot of a trajectory of the process with point predictors, estimative and bootstrap calibrated prediction bands of level $\alpha = 0.9$, for $m = 5$ future observations.

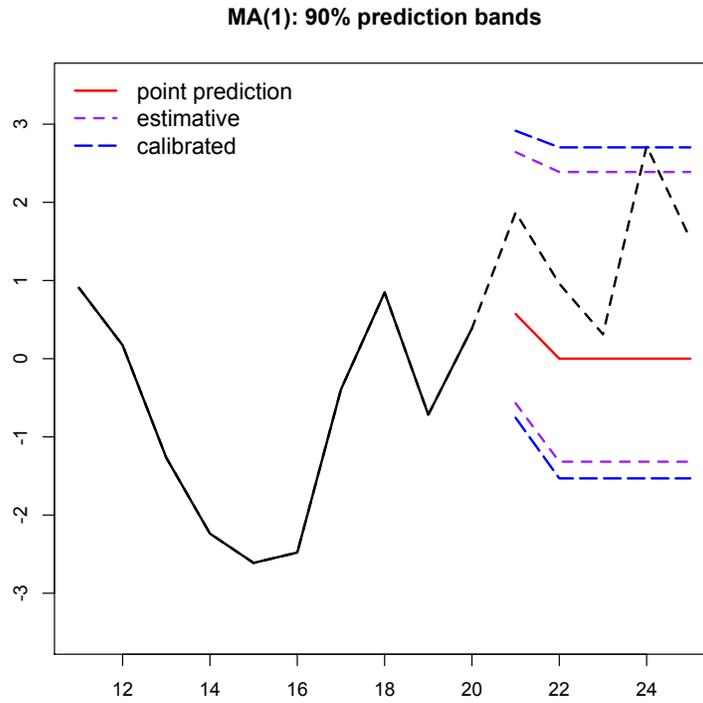


Figure 2: MA(1) Gaussian model. Plot of a trajectory of the process with point predictors, estimative and bootstrap calibrated prediction bands of level $\alpha = 0.9$, for $m = 5$ future observations.

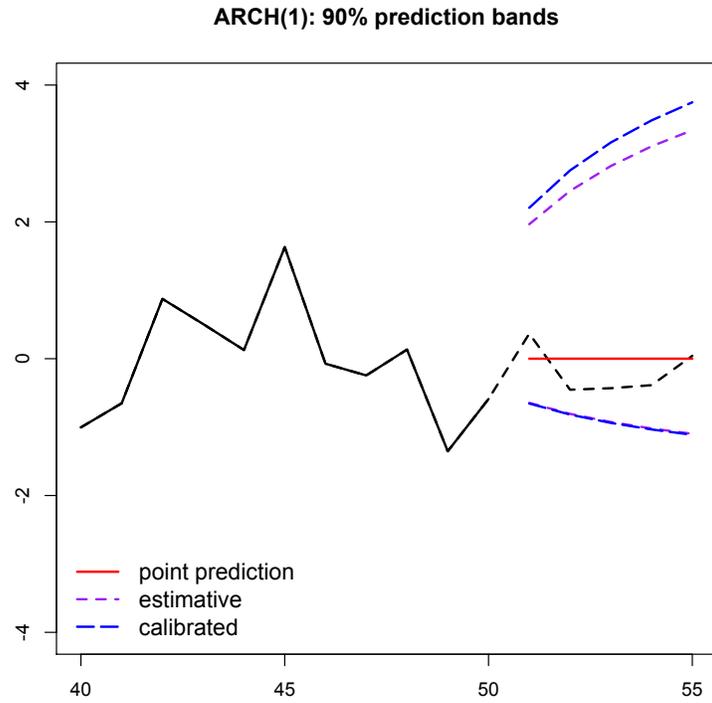


Figure 3: ARCH(1) Gaussian model. Plot of a trajectory of the process with point predictors, estimative and bootstrap calibrated prediction bands of level $\alpha = 0.9$, for $m = 5$ future observations.