Bootstrap calibrated predictive distributions for time series

Giovanni Fonseca[1], Federica Giummolè[2] and Paolo Vidoni[1]

[1] Dipartimento di Scienze Economiche e Statistiche
Università degli Studi di Udine
via Treppo 18, I-33100 Udine, Italy
(e-mail: `giovanni.fonseca@uniud.it`, `paolo.vidoni@uniud.it`)
[2] Dipartimento di Scienze Ambientali, Informatica e Statistica
Università Ca' Foscari Venezia
San Giobbe, Cannaregio 873, I-30121 Venice, Italy
(e-mail: `giummole@unive.it`)

ABSTRACT. This paper deals with prediction for time series models and, in particular, it presents a simple procedure which gives well-calibrated predictive distributions, generalizing the calibrating approach proposed by Beran (1990). The associated prediction intervals have coverage probability equal or close to the target nominal value. Although the exact computation of the proposed distribution is usually not feasible, it 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 to AR, ARCH and MA models are presented.

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. Indeed, 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).

In this paper we consider, for time series prediction, a simple procedure proposed by Fonseca *et al.* (2010), which gives well-calibrated predictive distributions. The associated prediction intervals have coverage probability equal or close to the target nominal value. Although the exact computation of the proposed distribution 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 to AR, ARCH and MA models are presented.

2 CALIBRATED PREDICTIVE DISTRIBUTION FOR TIME SERIES

Suppose that $\{Y_t\}_{t\geq 1}$ is a discrete-time stochastic process with probability distribution specified by the unknown d-dimensional parameter $\theta \in \Theta \subseteq \mathbf{R}^d$, $d \geq 1$; $Y = (Y_1, \cdots, Y_n)$, $n > 1$, is observable, while $Z = Y_{n+1}$ is a future or not yet available observation. We assume that (Y, Z) is a continuous random vector with joint density $g(y, z; \theta)$. In some cases, as for autoregressive processes, there exists a transitive statistic (Barndorff-Nielsen and Cox, 1996) $U = U(Y)$, with a fixed dimension independent of the sample size n, so that Y and Z are conditionally independent given U. We shall indicate with $g(z|y; \theta)$ and $G(z|y; \theta)$ the conditional density and distribution function of Z given $Y = y$, respectively; in the presence of a transitive statistic U, y is substituted by the observed value u of U.

Given the observed sample $y = (y_1, \ldots, y_n)$, an α-prediction limit for Z is a function $c_\alpha(y)$ such that, exactly or approximately,

$$P_{Y,Z}\{Z \leq c_\alpha(Y); \theta\} = \alpha, \tag{1}$$

for every $\theta \in \Theta$ and for any fixed $\alpha \in (0, 1)$. The above probability is called coverage probability and it is calculated with respect to the joint distribution of (Z, Y). When there exists a transitive statistics, it is natural (see, for example, Kabaila and Syuhada, 2007) to require that the conditional coverage probability is such that, exactly or approximately,

$$P_{Y,Z|U}\{Z \leq c_\alpha(Y)|U = u; \theta\} = \alpha, \tag{2}$$

where the probability is calculated with respect to the conditional distribution of (Z, Y) given $U = u$. Obviously, conditional solutions satisfying (2) also satisfy condition (1), and they are in some settings much easier to find. On the other hand, when we can not find a transitive statistic, the conditional approach is meaningless.

The calibrating approach proposed by Fonseca *et al.* (2010) extends that one suggested by Beran (1990) and it provides a predictive distribution function which gives, as quantiles, prediction limits with well-calibrated (conditional) coverage probability. This proposal is here applied to the case of time series models.

Consider the maximum likelihood estimator $\hat{\theta} = \hat{\theta}(Y)$ for θ, or an asymptotically equivalent alternative, and the estimative prediction limit $z_\alpha(y; \hat{\theta})$, which is obtained by substituting θ with $\hat{\theta}$ in the α-quantile $z_\alpha(y; \theta) = G^{-1}(\alpha|y; \theta)$, where $G^{-1}(\cdot|y; \theta)$ is the inverse of the distribution function $G(\cdot|y; \theta)$. The associated coverage probability is

$$P_{Y,Z}\{Z \leq z_\alpha(Y; \hat{\theta}); \theta\} = E_Y[G\{z_\alpha(Y; \hat{\theta})|Y; \theta\}; \theta] = C(\alpha, \theta)$$

and, although its explicit expression is rarely available, it is well-known that it does not match the target value α even if, asymptotically, $C(\alpha, \theta) = \alpha + O(n^{-1})$, as $n \to +\infty$. As proved in Fonseca *et al.* (2010), function

$$G_c(z|y; \hat{\theta}, \theta) = C\{G(z|y; \hat{\theta}), \theta\}, \tag{3}$$

which is obtained by substituting α with $G(z|y; \hat{\theta})$ in $C(\alpha, \theta)$, is a proper predictive distribution function, provided that $C(\cdot, \theta)$ is a sufficiently smooth function. Furthermore, it gives, as quantiles, prediction limits $z_\alpha^c(y; \hat{\theta}, \theta)$ which coverage probability equals the target nominal value α, for all $\alpha \in (0, 1)$.

The calibrated predictive distribution (3) is not useful in practice, since it depends on the unknown θ and a closed form expression for $C(\alpha, \theta)$ is rarely available. However, a suitable parametric bootstrap estimator for $G_c(z|y; \hat{\theta}, \theta)$ may be readily defined. Let y_*^j, $j = 1, \ldots, B$, be parametric bootstrap samples generated from the estimative distribution of the data and let $\hat{\theta}_*^j$, $j = 1, \ldots, B$, be the corresponding maximum likelihood estimates. Since $C(\alpha, \theta) = E_Y[G\{z_\alpha(Y; \hat{\theta})|Y; \theta\}; \theta]$, we define the bootstrap-calibrated predictive distribution as

$$G_c^b(z|y; \hat{\theta}) = \frac{1}{B} \sum_{j=1}^{B} G\{z_\alpha(y_*^j; \hat{\theta}_j^*)|y_*^j; \hat{\theta}\}\big|_{\alpha = G(z|y; \hat{\theta})}. \tag{4}$$

The corresponding α-quantile defines, for each $\alpha \in (0, 1)$, a prediction limit having coverage probability equal to the target α, with an error term which depends on the efficiency of the bootstrap simulation procedure.

In the presence of a transitive statistic U, a similar procedure may be considered. In this case, y is substituted by u and the conditional coverage probability of the estimative prediction limit $z_\alpha(u; \hat{\theta})$ is

$$P_{Y,Z|U}\{Z \leq z_\alpha(U; \hat{\theta})|U = u; \theta\} = E_{Y|U}[G\{z_\alpha(U; \hat{\theta})|U; \theta\}|U = u; \theta] = C_u(\alpha, \theta).$$

Indeed, the calibrated predictive distribution, giving prediction limits with conditional and unconditional coverage probability equal to α, corresponds to (3) with $C_u(\cdot, \theta)$ instead of $C(\cdot, \theta)$. Finally, the associated parametric bootstrap estimator is similar to (4), but it is based on simulated samples from the conditional distribution of Y given $U = u$, assuming $\theta = \hat{\theta}$. In this context, it could be convenient to apply the simulation technique proposed by Kabaila (1999) for estimating conditional expectations or, whenever possible, to use the backward representation of stationary autoregressive processes.

3 AUTOREGRESSIVE MODELS

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

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

where μ and ρ are unknown parameters and $\{\varepsilon_t\}_{t \geq 1}$ 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. The unknown parameter is $\theta = (\mu, \rho, \sigma^2)$ and likelihood inference is conditioned on $Y_0 = y_0$, with y_0 known. The observable random vector is $Y = (Y_1, \ldots, Y_n)$ and the next realization of the process is $Z = Y_{n+1}$. The conditional distribution of Z given $Y = y$ is Gaussian with mean $\mu_{n+1} = \mu + \rho(y_n - \mu)$ and variance $\sigma^2_{n+1} = \sigma^2$. Indeed, Y_n is a transitive statistic and, as explained in the previous section, we evaluate the performance of a prediction limit by means of its coverage probability conditioned on the observed value y_n of Y_n.

An approximated solution to this problem has already been considered by Vidoni (2004), using complicated asymptotic calculations. An alternative easier solution is given by the parametric bootstrap estimator (4), where the bootstrap samples are generated keeping the observed value of the transitive statistic fixed. A simulation study shows the performance of this predictive solution. Conditional coverage probabilities for estimative, asymptotically approximated and bootstrap calibrated prediction limits of level $\alpha = 0.9, 0.95$ are calculated by means of the simulation technique presented in Kabaila (1999), keeping the last observed value fixed to $y_n = 1, 0, -1$ and $y_0 = 0$. The results are collected in Table 1 and show that both the approximated asymptotic solution and, in particular, the bootstrap solution remarkably improve on the estimative one. Thus, the bootstrap estimator can be fruitfully considered as a valid and simpler alternative to asymptotic methods.

α	n	y_n	Estimative	Approximated	Bootstrap
0.9	25	-1	0.868	0.889	0.893
		0	0.871	0.884	0.900
		1	0.860	0.883	0.898
	50	-1	0.884	0.895	0.897
		0	0.887	0.894	0.900
		1	0.882	0.893	0.900
0.95	25	-1	0.924	0.944	0.946
		0	0.925	0.940	0.950
		1	0.918	0.939	0.949
	50	-1	0.937	0.947	0.949
		0	0.939	0.946	0.950
		1	0.936	0.946	0.950

Table 1. AR(1) Gaussian model. Conditional coverage probabilities for estimative, approximated and bootstrap calibrated prediction limits of level $\alpha = 0.9, 0.95$, conditioned on $y_n = -1, 0, 1$; $\rho = 0.5$, $\mu = 0$, $\sigma^2 = 1$, $y_0 = 0$ and $n = 25, 50$. Estimation is based on 5,000 Monte Carlo replications. Bootstrap procedure is based on 2,000 bootstrap samples. Estimated standard errors are always smaller than 0.005.

4 ARCH MODELS

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

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

where β and γ are unknown parameters and $\{\varepsilon_t\}_{t\geq 1}$ 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 $Y = (Y_1, \ldots, Y_n)$ and the next realization of the process is $Z = Y_{n+1}$. The conditional distribution of Z given $Y = y$ is Gaussian with zero mean and variance $\beta + \gamma y_n^2$. Indeed, Y_n is a transitive statistic and we evaluate a prediction limit by means of its coverage probability conditioned on the observed value y_n of Y_n.

As for autoregressive models, an approximated solution to this problem, based on asymptotic calculations, has already been considered by Vidoni (2004). However, it is possible to obtain an alternative simpler solution by means of the parametric bootstrap estimator (4). A simulation study shows the performance of this predictive solution. Here, the bootstrap samples are generated keeping the observed value of the transitive statistic fixed. Conditional coverage probabilities for estimative, asymptotically approximated and bootstrap calibrated prediction limits of different levels are calculated by means of the simulation technique presented in Kabaila (1999). The results are collected in Table 2 and confirm the superiority of the bootstrap calibrated prediction limits over the estimative and the approximated ones.

α	n	y_n	Estimative	Approximated	Bootstrap
0.95	25	0	0.934	0.965	0.950
		1	0.929	0.943	0.944
		2	0.926	0.932	0.941
	50	0	0.943	0.955	0.950
		1	0.940	0.947	0.949
		2	0.937	0.941	0.949
0.99	25	0	0.976	0.986	0.991
		1	0.977	0.980	0.985
		2	0.972	0.963	0.981
	50	0	0.984	0.988	0.990
		1	0.985	0.986	0.990
		2	0.981	0.978	0.990

Table 2. ARCH(1) Gaussian model. Conditional coverage probabilities for estimative, approximated and bootstrap calibrated prediction limits of level $\alpha = 0.95, 0.99$, conditioned on $y_n = 0, 1, 2$; $\beta = 0.5$, $\gamma = 1$, $y_0 = 0$ and $n = 25, 50$. Estimation is based on 5,000 Monte Carlo replications. Bootstrap procedure is based on 2,000 bootstrap samples. Estimated standard errors are always smaller than 0.005.

5 MOVING AVERAGE MODELS

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

$$Y_t = \mu + \varepsilon_t + \rho\varepsilon_{t-1}, \qquad t \geq 1,$$

with $\varepsilon_t \sim N(0, \sigma^2)$, $t \geq 0$, independent Gaussian distributed random variables. We want to predict $Z = Y_{n+1}$ on the basis of an observed sample $y = (y_1, \ldots, y_n)$ from $Y = (Y_1, \ldots, Y_n)$.

Here the unknown parameter is $\theta = (\mu, \rho, \sigma^2)$, with $\sigma \in \mathbf{R}^+$ and $|\rho| < 1$. The conditional distribution of Z given Y is $Z|Y \sim N(\mu_{n+1}, \sigma_{n+1}^2)$, where $\mu_{n+1} = \mu + \sum_{i=1}^{n} (-1)^{i+1} \rho^i (Y_{n+1-i} - \mu)$, and $\sigma_{n+1}^2 = \sigma^2$. Since we work conditionally on $\varepsilon_0 = 0$ and $|\rho| < 1$, the effect of conditioning vanishes as the sample size n increases.

An asymptotic solution to this problem, involving tedious calculations, has already been considered by Giummolè and Vidoni (2010). Also in this case, it is possible to obtain an alternative simpler solution by means of the parametric bootstrap estimator (4). Table 3 shows the results of a simulation study for comparing coverage probabilities associated to estimative, asymptotically approximated and bootstrap calibrated prediction limits of different levels. In this case, the approximated and the bootstrap calibrated solutions behave similarly and both improve the estimative one.

α	n	Estimative	Approximated	Bootstrap
0.9	25	0.859	0.870	0.886
	50	0.889	0.898	0.900
0.95	25	0.915	0.933	0.938
	50	0.940	0.948	0.949
0.99	25	0.971	0.984	0.983
	50	0.984	0.989	0.988

Table 3. MA(1) Gaussian model. Coverage probabilities for estimative, approximated and bootstrap calibrated prediction limits of level $\alpha = 0.9, 0.95, 0.99$; $\mu = 1$, $\sigma^2 = 1$, $\rho = 0.5$, $\varepsilon_0 = 0$ and $n = 25, 50$. Estimation is based on 5,000 Monte Carlo replications. Bootstrap procedure is based on 2,000 bootstrap samples. Estimated standard errors are always smaller than 0.005.

REFERENCES

BARNDORFF-NIELSEN, O.E., COX, D.R. (1996): Prediction and asymptotics. *Bernoulli*, *2*, 319–340.
BERAN, R. (1990): Calibrating prediction regions. *Journal of the American Statistical Association*, *85*, 715–723.
CLEMENTS, M.P., KIM, J.H. (2007): Bootstrap prediction intervals for autoregressive time series. *Computational Statistics & Data Analysis*, *51*, 3580–3594.
FONSECA, G., GIUMMOLÈ, F., VIDONI, P. (2010): Calibrating predictive distributions. *Redazioni Provvisorie 2/2010, Dept. of Statistics, Ca' Foscari University of Venice*.
GIUMMOLÈ, F, VIDONI, P. (2010): Improved prediction limits for a general class of Gaussian models. *Journal of Time Series Analysis*, *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. *Australian & New Zeland Journal of Statistics*, *41*, 331–336.
KABAILA, P., SYUHADA, K. (2007): Improved prediction limits for AR(p) and ARCH(p) processes. *Journal of Time Series Analysis*, *29*, 213–223.
VIDONI, P. (2004): Improved prediction intervals for stochastic process models. *Journal of Time Series Analysis*, *25*, 137–154.