ARMA models with time dependent coefficients

Keywords: nonstationary time series, time-varying models, official statistics, production.

1. Introduction

Two decades ago, effective methods for dealing with time series models that vary with time have appeared in the statistical literature. Except in a case of marginal heteroscedasticity [1], they have never been used for official statistics. In this paper, we consider autoregressive integrated moving average (ARIMA) models with time-dependent coefficients applied to very long U.S. industrial production series. There was well an attempt to handle time-dependent integrated autoregressive (AR) models [2] but the case study was small. Here, we investigate the case of ARIMA models on the basis of [3, 4, 5]. As an illustration, we consider a big dataset of U.S. industrial production time series already used in [6]. We employ the software package Tramo in [7] to obtain linearized series and we built both ARIMA with constant coefficients (cARIMA) and ARIMA models with time-dependent coefficients (tdARIMA). In these tdARIMA models we use the simplest specification: a regression with respect to time. Surprisingly, for a large part of the series, there are statistically significant slopes, indicating that the tdARIMA models fit better the series than the cARIMA models.

2. Methods

We consider the well-known class of seasonal ARIMA models, see e.g. [7, 8]. Models with time-dependent coefficients are often used in econometrics but not for ARIMA models. For very long time series, there is no reason that the coefficients would stay constant. They can be supposed to vary slowly with time although breaks can also be considered. Time series models with time-varying coefficients have nevertheless been studied, mainly from a theoretical point of view. Several papers [3, 4, 5] provide conditions for asymptotic properties, hence justification for statistical inference.

2.1. The model

To illustrate a simple ARIMA model with time dependent coefficient, we can consider the ARMA(1,1) model. Let the series be denoted by $y = (y_1, y_2, \ldots, y_n)$. Then a tdARMA(1,1) model is described by the equation

$$y_t = \phi_t^{(n)} y_{t-1} + e_t - \theta_t^{(n)} e_{t-1}, \quad (2.1)$$

where the e_t are independent random variables with mean zero and with standard deviation σ, and the time dependent coefficients $\phi_t^{(n)}$ and $\theta_t^{(n)}$ depend on time t, also on n, the length of the series but also on a small number of parameters. The simplest specification for $\phi_t^{(n)}$, for example, is

$$\phi_t^{(n)}(\beta) = \phi + \frac{1}{n-1}\left(t - \frac{n+1}{2}\right)\phi', \qquad (2.2)$$

and a similar expression for $\theta_t^{(n)}$ using two other parameters θ and θ'. The vector β contains all parameters to be estimated, those in $\phi_t^{(n)}$ (like ϕ and ϕ', here) and $\theta_t^{(n)}$ (θ and θ'), but not the scale factor σ^2 which is estimated separately. For the corresponding cARIMA model, we have of course $\phi' = \theta' = 0$. For a lag k instead of 1, we add a subscript k to the parameter symbols.

2.2. The estimation method

For any tdARIMA model, we can estimate the parameters by maximizing the logarithm of the Gaussian likelihood. Under some very general conditions, it can be shown that the quasi maximum likelihood estimator $\hat{\beta}$ converges to the true value of β, and is asymptotically normal. Moreover, the asymptotic covariance matrix of $\hat{\beta}$ can be evaluated as a by-product of estimation. An improved Marquardt algorithm is used for that purpose. The Student *t* statistics shown in the next section are based on the standard errors deduced from the evaluation of that asymptotic covariance matrix.

2.3. The dataset

We use a big dataset of U.S. industrial production series already used by in [6]. See http://www.federalreserve.gov/releases/g17/ipdisk/ip_nsa.txt. These are 321 time series from January 1986 to December 2017. The models were fitted until December 2013 leaving the remaining months to compare the data to the ex-post forecasts, using either forecasts with a fixed forecast origin or rolling forecasts.

We employ the software package Tramo described in [7] to obtain partially linearized series by removing outliers. Indeed, presence of outliers can distort the analysis. Finding the cARIMA models in an automated way is also done using Tramo. Then we replace the constant coefficients by linear functions of t for order $k \leq 13$, giving tdARIMA models, using the simple linear regression approach illustrated in (2.2) for each lag k coefficient in the model. At this stage, we do not omit non-significant parameters. The cARIMA and tdARIMA models are fitted using the same specialised software package. In an earlier attempt of this project, the data were limited to 2016 and were not corrected for outliers, and only fixed forecasts were considered.

3. RESULTS

We compare the results of tdARIMA versus cARIMA models using the following criteria:

• is the highest Student *t* statistic of the slopes, the td parameters, in absolute value, larger than 1.96 ?

• is tdARIMA SBIC smaller than the corresponding cARIMA SBIC ?

• is tdARIMA residual standard deviation smaller than the corresponding cARIMA one ?

• is tdARIMA P-value of the Ljung-Box (LB) statistic for residual autocorrelation larger (with 48 lags) than the corresponding cARIMA one ?

• is tdARIMA mean absolute percentage error (MAPE) in percent for 2014, for all horizons 1 to 12, smaller than the corresponding cARIMA one ?

• are tdARIMA mean absolute percentage error (MAPE) in percent for rolling forecasts, at horizons 1, 3, 6 and 12, smaller than the corresponding cARIMA one ?

We count the percentages for each criterion over the 321 series. The results are shown in Table 1. They are very interesting since about 40% of the series show at least one statistically significant slope parameter at the 5% level. A majority of the series have smaller residual standard deviation, and less residual autocorrelation. That SBIC is worse for most of the series which can be partly explained by the fact that non-significant slope parameters were left in the model. The only unsatisfactory aspect of tdARIMA models is that they fail to improve the forecasts for a majority of series. To provide a better view of these results, we have added percentages conditional to significant time-dependency in the last column of Table 1. Except a smaller residual standard deviation, they confirm that only one third of the "time dependent series", i.e. those series which have at least one statistically significant slope parameter, provide better forecasts with a tdARIMA model than with a cARIMA model. For rolling forecasts, the results are still slightly worse. This is surprising although we know that a better fit is not a guarantee for better forecasts.

Table 1. For each criterion, we give the percentages of improvement over the 321 U.S. series from cARIMA models to tdARIMA models. The last column contains percentages conditional to the existence of at least one statistically significant slope parameter ϕ_k' or θ_k'.

Criteria	Percentage	Percentage if slope significant
Highest $\|t\|$ statistic of slopes > 1.96	42.06	100.00
tdARIMA SBIC < cARIMA SBIC	04.67	09.63
tdARIMA residual std dev < cARIMA	52.02	72.59
tdARIMA LB P-value > cARIMA	57.63	60.00
tdARIMA 2014 forecasting MAPE < cARIMA	33.96	33.33
tdARIMA horizon 1 rolling forecasts MAPE < cARIMA	26.48	20.74
tdARIMA horizon 3 rolling forecasts MAPE < cARIMA	25.55	18.52
tdARIMA horizon 6 rolling forecasts MAPE < cARIMA	27.73	21.48
tdARIMA horizon 12 rolling forecasts MAPE < cARIMA	32.71	29.63

4. CONCLUSIONS

These results seem to confirm the first results obtained with that dataset of very long official statistics time series. Indeed, a large proportion of the tdARIMA models contain at least one statistically significant slope. It was a majority with slightly shorter series where the outliers were not removed. So it is not the presence of outliers that could lead to better fits by tdARIMA models. A common feature is nevertheless that the forecasts are not improved by replacing the cARIMA models by tdARIMA models. Consequently, we have not solved the paradox yet and more research is deserved. On the other side, it would be also interesting to conclude that traditional cARIMA models are enough to forecast very long time series and that no substantial gain can be obtained by considering tdARIMA models.

Besides the results presented above, it is intended to work with slightly longer series (up to December 2018) that are completely linearized (where outliers and calendar effects are removed), not only corrected for outliers. We will also delete one by one the insignificant slope parameters in order to have tdARIMA models that are more parsimonious. For the statistical tests on the coefficients, we will use global tests instead of individual tests on each parameter involved in time-dependency. Finally, that kind of analysis can be repeated with other datasets, like those maintained by Eurostat.

5. ACKNOWLEDGEMENTS

We thank Agustín Maravall for his help in order to produce linearized time series in Tramo, and Ahmed Ben Amara for having contributed to a very first version of a part of the program chain (which includes Tramo-Seats and Microsoft Excel Visual Basic modules), in addition to the specialised code for estimating tdARIMA models.

REFERENCES

[1] S. Van Bellegem and R. von Sachs, Forecasting economic time series with unconditional time-varying variance, International Journal of Forecasting 20 (2004), 611-627.

[2] R. Azrak and G. Mélard, AR models with time-dependent coefficients with a macro-economic example, NTTS 2017 Conference on New Techniques and Technologies for Statistics, Eurostat, Brussels, 14-16 March, 2017 (2017).

[3] R. Azrak and G. Mélard, Asymptotic properties of quasi-maximum likelihood estimators for ARMA models with time-dependent coefficients, Statistical Inference for Stochastic Processes 9 (2006), 279-330.

[4] R. Dahlhaus, Maximum likelihood estimation and model selection for locally stationary processes, Journal of Nonparametric Statistics 6 (1996), 171-191.

[5] R. Dahlhaus, Fitting time series models to nonstationary processes, Annals of Statististics 25 (1997), 1-37.

[6] T. Prioetti and H. Lüktepohl, Does the Box-Cox transformation help in forecasting macroeconomic time series?, International Journal of Forecasting 29 (2013), 88-99.

[7] V. Gómez and A. Maravall, Automatic modelling methods for univariate series, In A Course in Time Series Analysis, D. Peña, G.C. Tiao, and R.S. Tsay (Eds), New York: Wiley, Chapter 7 (2001), 171-201.

[8] G. E. P. Box, G. M. Jenkins, G. C. Reinsel and G. S. Ljung, Time Series Analysis, Forecasting and Control, Fifth Edition, New Jersey: Prentice-Hall (2016).