Forecasting the Rand

MG Ferreira

November 2002

Abstract

This note looks at the performance of various time series models in predicting the near term future of the rand/dollar exchange rate. Each technique's prediction power is evaluated over a period of 24 and 36 months respectively by evaluating the one-month forecast's mean absolute percentage error. The theta decomposition model is the clear winner in both the 24 and 36 month 'competitions'. Thereafter the various models' forecasts are combined optimally in such a way that the one-month forecast mean absolute percentage error of the weighted average of forecasts is minimised. This yields a further marginal improvement in the mean absolute percentage error which suggests that a combination of models should outperform the best of the suite of models.

1 Forecasting in the very short term

Typical financial time series such as exchange rates are driven by a combination of *internal* and *external* factors. Here internal refers to those factors inherent in the time series itself such as trend and seasonal fluctuations. These are typically the subject of univariate time series models and in the short term play the dominant role in determining the trajectory of the series. External factors, on the other hand, refer to 'fundamental' factors as dicated by economic theory, such as interest rate and price differentials. In the short term these factors have almost no influence but in the long term they play a dominant role in determining the series' trajectory. Models based on external factors are typically multivariate regression models. Models that try to capture the middle ground between time series and fundamental models are vector autoregression and neural network type models.

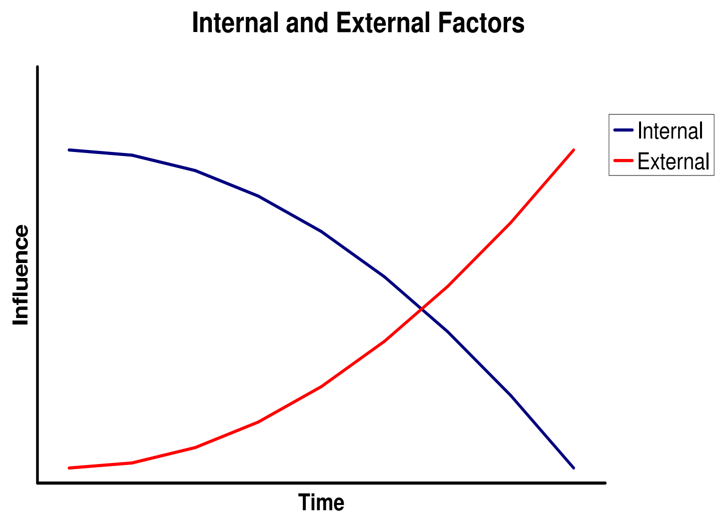

Here the forecast performance of a variety of models are investigated in forecasting the short term, one month ahead future of the rand/dollar exchange rate. Thus the emphasis is on time series models. A variety of ARIMA type models as well as Holt-Winters exponential smoothing techniques are investigated. The variety of ARIMA models stems from a visual inspection of the correlogram and partial correlogram of log-changes in the rand and a selection made from this visual inspection of potential orders for ARIMA models. These models, ARIMA and exponential smoothing, can be viewed as the classical approach to short term forecasting.

Two other types of models are also examined - these are the STL (seasonal-trend-Loess) and Theta models. Both these models are decomposition models and have received considerable interest recently, especially due to the excellent performance of the Theta model in the most recent M competition as published by the International Journal of Forecasting.

2 Methodology

A fairly large sample of month-end exchange rates has been used in this study. The series starts in 1986 and spans several years of the rand's fluctuating and volatile history with the last point being the month-end figure of October 2002. From the full data set one observation was removed and all the remaining points used to estimate the various models' parameters. Each model was then required to predict one month into the future and this prediction was compared to the observation in the hold-out period. Then the sample was again reduced by removing the last observation (this time the second-to-last observation in the full sample) and the various models' parameters was estimated. Again a one-month prediction was obtained from each model and this was compared to the actual outcome in the hold-out sample. This process was repeated until at first 24 and later on 36 actual observations vs one month predictions were obtained. As in the M competition the MAPE (mean absolute percentage error) was used as a yardstick to benchmark model performance.

For each data set thus obtained, the 24 and 36 point sets, the forecasts from the various models were combined using a simple weighted average. The weights were determined optimally using the constrained limited memory modification of the quasi-Newton BFGS (Broyden, Fletcher, Goldfarb and Shanno, 1970) algorithm. Weights were constrained to be positive and less than unity. The goal function that was minimsed was the MAPE of the weighted average forecast. Note that the optimisation process, other than the individual models, does look into the future in that it optimises the known forecast performance of the models.

3 Model performance

The following figures present the various models' performance for the 24 and the 36 point forecast samples.

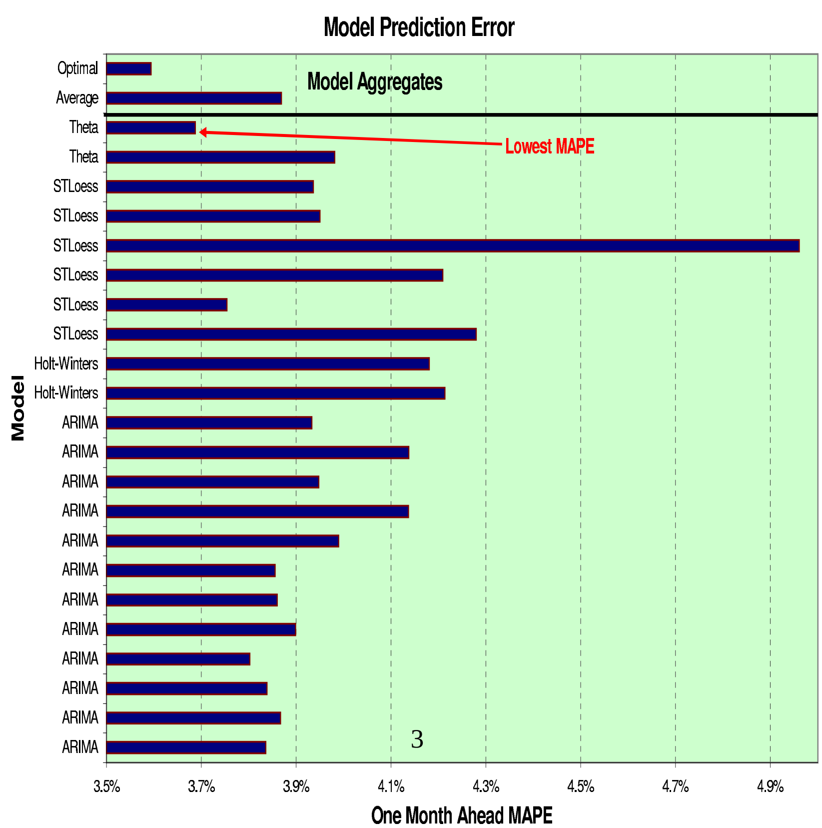

In the 24 month sample note the outperformance of the Theta model as well as the apparent barrier at around 3.6%. Also note that if a simple (equally weighted) average of all the models is used quite a mediocre prediction is obtained while an average that optimally combine predictions see a fairly large improvement in the MAPE. It is interesting to note that the optimisation process assigned zero weights to most models and roughly equal weights to the 'best' model from each of the ARIMA, STLoess and Theta categories. This is encouraging as it is quite similar to the way in which one would combine forecasts without the advantage of an optimiser - pick the best model from each category. Again, it is encouraging to note that this optimal combination does improve the MAPE. Other financial series, especially commodities, yield an even more significant improvement in MAPE if different forecasts are optimally combined.

If a 36 point sample is used the MAPE 'barrier' lowers dramatically to around 3.1%. Once again the Theta model is superior while a simple average of all models does not appear to yield an exceptional forecast. Again the optimal combination of weights succeeds in a notable MAPE improvement.