Available online at www.sciencedirect.com

ScienceDirect

International Journal of Forecasting 26 (2010) 885–907

international journal of forecasting

www.elsevier.com/locate/ijforecast

General-to-specific modelling of exchange rate volatility: A forecast evaluation

Luc Bauwens[a], Genaro Sucarrat[b,*]

[a] CORE and Department of Economics, Université catolique de Louvain, Belgium
[b] Norwegian School of Management, Norway

Abstract

The general-to-specific (GETS) methodology is widely employed in the modelling of economic series, but less so in financial volatility modelling, due to its computational complexity when many explanatory variables are involved. This study proposes a simple way of avoiding this problem when the conditional mean can appropriately be restricted to zero, and undertakes an out-of-sample forecast evaluation of the methodology applied to the modelling of the weekly exchange rate volatility. Our findings suggest that GETS specifications perform comparatively well in both *ex post* and *ex ante* forecasting as long as sufficient care is taken with respect to the functional form and the way in which the conditioning information is used. Also, our forecast comparison provides an example of a discrete time explanatory model being more accurate than the realised volatility *ex post* in 1-step-ahead forecasting.

Keywords: Exchange rate volatility; General-to-specific; Forecasting

1. Introduction

Exchange rate variability is an issue of great importance to both businesses and policymakers. Businesses use volatility models as tools in their risk management and as inputs in derivative pricing, whereas policymakers use them to acquire knowledge about the impact of economic factors on exchange rate variability for informed policymaking. Most volatility models are highly non-linear, and thus require complex optimisation algorithms for their empirical application. For models with few parameters and few explanatory variables, this may not pose unsurmountable problems. However, as the number of parameters and explanatory variables increases, the resources required for reliable estimation and model validation multiply. Indeed, this may even become an obstacle to the application of certain econometric modelling strategies, as was argued by Granger and Timmermann (1999) and McAleer (2005), for example, regarding automated general-to-specific (GETS)

* Corresponding author.
E-mail addresses: bauwens@core.ucl.ac.be (L. Bauwens), genaro.sucarrat@bi.no (G. Sucarrat).

0169-2070/$ - see front matter
doi:10.1016/j.ijforecast.2010.07.001

modelling of financial volatility.[1] GETS modelling is particularly well suited to explanatory econometric modelling, since it provides a systematic framework for statistical economic hypothesis testing, model development and model (re-)evaluation, and the methodology is relatively popular among large scale econometric model developers and proprietors. However, since the initial model formulation typically requires many explanatory variables, this poses challenges for computationally complex models at the outset.

The recent developments by Doornik (2009) and Hendry, Johansen, and Santos (2008) might be a step towards overcoming some of the computational challenges associated with the maximum likelihood estimation of financial models when many variables are included in the variance specification. However, this is still to be investigated, since their work is on the conditional mean using ordinary least squares estimation. Meanwhile, in this study we overcome the computational challenges traditionally associated with the application of the GETS methodology in the modelling of financial volatility by modelling volatility within an exponential model of variability (EMOV), where the variability is defined as the squared returns. The parameters of interest can therefore be estimated consistently with ordinary least squares (OLS) under rather weak assumptions. This setup implies that the conditional mean is restricted to zero, but enables us in return to apply GETS to a general specification, with, in our case, a constant and 24 regressors, including lags of the log of squared returns, an asymmetry term, a skewness term, seasonality variables, and economic covariates. Compared with models belonging to the autoregressive conditional heteroscedasticity (ARCH) and stochastic volatility (SV) classes, we estimate and simplify our specification with little effort, and obtain a parsimonious encompassing specification with uncorrelated homoscedastic residuals and relatively stable parameters. Moreover, our out-of-sample forecast evaluation suggests that GETS specifications can be particularly valuable in conditional forecasting—as long as sufficient care is taken as to where and how the conditioning information enters, since the *ex post* EMOV specification performs particularly well.

Another contribution of this study is a qualificatory note on the evaluation of explanatory economic models of financial volatility against estimates based on continuous time theory. Highly simplified, the return volatility forecasting literature can be divided in two parts: before and after the highly influential publication of Andersen and Bollerslev (1998). Although in-sample estimates suggest the widespread presence of ARCH, asymmetry effects, jumps, volume effects, and so on in financial returns volatility, models that include these effects tend to explain a very small proportion of the return variability out-of-sample (see Poon & Granger, 2003, for a review of the literature). Andersen and Bollerslev (1998) argued that this is because the standard estimates of volatility are very noisy, and suggested instead that forecasts of volatility should be evaluated against high frequency *ex post* estimates, for example the realised volatility (sums of intra-period squared returns). Andersen and Bollerslev (1998) were not the first to put forward this explanation and solution; nevertheless, they had the greatest impact. Subsequently, the general view that has emerged is that discrete time models of financial volatility should be evaluated against estimates derived from continuous time theory, not against the return variability (for example squared returns); see *inter alia* Andersen, Bollerslev, and Lange (1999), Andersen, Bollerslev, Diebold, and Labys (2003), Andersen, Bollerslev, and Meddahi (2005), Andersen, Bollerslev, Christoffersen, and Diebold (2006), and Hansen and Lunde (2005, 2006). As a consequence, little if any role is left for the residuals to play—directly or indirectly—in the forecast evaluation. This is the direct opposite of the GETS methodology, where the analysis of the residuals plays a key role in model evaluation and comparison, since any empirical model is a highly simplified representation

[1] Making a distinction between the general-to-specific specification search on the one hand and the GETS methodology on the other is useful at this point. General-to-specific specification search plays a central role in the GETS methodology, but the methodology also embodies a particular view of the relationship between reality and the empirical models, which gives rise to a certain set of model evaluation criteria and modelling objectives. The GETS methodology is also referred to on occasion as the "LSE methodology", after the institution where the methodology to a large extent originated, the "Hendry methodology", after the most influential and arguably the most important contributor to the development of the methodology, and sometimes even "British econometrics", see Campos, Ericsson, and Hendry (2005), Gilbert (1989, 1990), Hendry (2003) and Mizon (1995).