Available online at www.sciencedirect.com

SciVerse ScienceDirect

International Journal of Forecasting 27 (2011) 1089–1107

international journal of forecasting

www.elsevier.com/locate/ijforecast

Forecasting exchange rate volatility using high-frequency data: Is the euro different?

Georgios Chortareas[a,1], Ying Jiang[b,*], John. C. Nankervis[c,2]

[a] *Department of Economics, University of Athens, Athens, 10559, Greece*
[b] *Nottingham University Business School China, University of Nottingham Ningbo, 199 Taikang East Road, Ningbo 315100, China*
[c] *Essex Business School, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, UK*

Abstract

We assess the performances of alternative procedures for forecasting the daily volatility of the euro's bilateral exchange rates using 15 min data. We use realized volatility and traditional time series volatility models. Our results indicate that using high-frequency data and considering their long memory dimension enhances the performance of volatility forecasts significantly. We find that the intraday FIGARCH model and the ARFIMA model outperform other traditional models for all exchange rate series.

Keywords: Euro exchange rates; Volatility forecasting; High-frequency data; GARCH model; Long memory time series; Forecast evaluation

1. Introduction

Volatility forecasting of asset prices in general, and exchange rates in particular, has been the focus of research in areas such as investment analysis, derivative securities pricing and risk management. Moreover, since the volatility of financial markets has a direct influence on policymaking, volatility forecasts can play the role of a 'barometer for the vulnerability of financial markets and the economy' (Poon & Granger, 2003). Poon and Granger (2003) review 93 papers in the volatility forecasting field and show that different models for forecasting the exchange rate volatility perform differently for different currencies. In this paper we evaluate the daily volatility forecasting performances of alternative models for euro exchange rates using high-frequency data.

Until quite recently, the literature typically focused on daily returns for forecasting the daily volatility, and used the daily squared returns as a measure of the 'true volatility'. However, daily squared returns

* Corresponding author. Tel.: +86 0 574 88180200; fax: +86 0 574 88180125.
E-mail addresses: gchortar@econ.uoa.gr (G. Chortareas), ying.jiang@nottingham.edu.cn (Y. Jiang), jcnank@essex.ac.uk (J.C. Nankervis).
[1] Tel.: +30 210 3689805; fax: +30 210 3689810.
[2] Tel.: +44 0 1206 873973; fax: +44 0 1206 873429.

0169-2070/$ - see front matter
doi:10.1016/j.ijforecast.2010.07.003

are not an accurate measure of the true volatility, since they are calculated from closing prices and therefore cannot capture price fluctuations during the day (see Andersen & Bollerslev, 1998). In response to these limitations, Andersen and Bollerslev (1998) propose the realized volatility (constructed from intraday returns) as a measure of the true volatility, and this measure has since become very popular. High-frequency data carry more information on daily transactions, and are useful not only in measuring volatility, but also in direct model estimation and forecast evaluation. Many recent methodological advances focus on high-frequency data,[3] while a number of studies build on this literature to evaluate the performance of alternative models for volatility forecasting.[4]

While there exist a number of studies on foreign exchange volatility forecasting,[5] as is discussed in Section 2, to the best of our knowledge, limited work has been done on forecasting the volatility of euro exchange rates. Since its introduction in 1999, the euro has become a major international currency, quickly establishing itself as the second most widely used international currency after the US dollar.[6] Nevertheless, the literature on exchange rate volatility forecasting focuses on USD exchange rates alone.

Our study addresses this gap in the literature by providing a characterization of the euro's exchange rate volatility at both the daily and intraday frequencies, and considers questions such as: Are the same models appropriate for the euro exchange rate as for the USD exchange rate? Do high-frequency euro exchange rates have properties similar to those of other high-frequency data? Can a long memory factor improve the performance of exchange rate volatility forecasting?

To answer these questions we compare the out-of-sample daily volatility forecast performances of traditional time series volatility models with that of a realized volatility model at high frequencies. The traditional time series volatility models considered include the GARCH model, the stochastic volatility (SV) model, the stochastic volatility with exogenous variables (SVX) model, and finally, the fractionally integrated GARCH (FIGARCH) model. The realized volatility model is an ARFIMA model.[7,8] We compare the performances of the two types of long memory models (FIGARCH and ARFIMA) using high-frequency data. We also compare the properties of the intraday GARCH and FIGARCH models with those of ARFIMA models which use the daily realized volatility. Finally, we compare the intraday GARCH model with the intraday FIGARCH model to provide evidence on whether modelling the long memory property in a high-frequency volatility process can improve the daily forecast performance.

For the intraday GARCH and FIGARCH models we use deseasonalized 15 min data on returns for a period covering almost four years. We thus obtain a very large number of observations relative to other studies that apply standard volatility models to intraday returns (e.g., Beltratti & Morana, 1999; Marlik, 2005; Martens, 2001; Rahman & Ang, 2002). Marlik (2005), for example, uses 30 min data covering a period of four months.

We employ a battery of tests to evaluate the out-of-sample forecast performances of the models considered. In addition to the regression test and the accuracy test, we also use the superior predictive ability test (Hansen, 2005) and an equal accuracy test, namely the adjusted Diebold-Mariano (1995) test. The results of these tests show that the intraday FIGARCH model always outperforms other traditional models, and produces results that are not significantly different to those from the realized volatility (ARFIMA) model. This is not atypical of the

[3] Examples of this type of analysis include the use of long memory ARFIMA (Autoregressive Fractional Integration Moving Average) models for forecasting the realized volatility (e.g., Pong, Shackleton, Taylor, & Xu, 2004), extending the daily model to include intraday information (e.g., Koopman, Jungbacker, & Hol, 2005) and the direct modelling of intraday returns using standard volatility models (e.g., Marlik, 2005).

[4] For example, see Andersen, Bollerslev, Diebold, and Labys (2003, ABDL hereafter), Hol and Koopman (2002), Martens (2001) and Martens and Zein (2004).

[5] Examples include Andersen and Bollerslev (1998), Andersen, Bollerslev, Diebold, and Labys (1999, 2000, 2003), Martens (2001), Vilasuso (2002) and West and Cho (1995).

[6] European Central Bank (1999).

[7] The ARFIMA model is fitted to the daily realized volatility. The GARCH and FIGARCH models are both fitted to deseasonalized 15 min returns. The GARCH, SV and SVX models are also fitted to daily returns.

[8] This paper focuses on the ARFIMA model alone. Other newly developed realized volatility models include the Heterogeneous Autoregressive (HAR) model (Corsi, 2009) and the Mixed Data Sampling (MIDAS) model (Ghysels, Sinko, & Valkanov, 2007).