International Review of Financial Analysis 27 (2013) 21–33

Contents lists available at SciVerse ScienceDirect

International Review of Financial Analysis

Forecasting value-at-risk and expected shortfall using fractionally integrated models of conditional volatility: International evidence

Stavros Degiannakis [a,b,*], Christos Floros [a], Pamela Dent [a]

[a] Department of Economics, Portsmouth Business School, University of Portsmouth, Richmond Building, Portland Street, Portsmouth, PO1 3DE, UK
[b] Postgraduate Department of Business Administration, Hellenic Open University, Aristotelous 18, 26 335, Greece

ARTICLE INFO

Article history:
Received 6 May 2011
Received in revised form 16 May 2012
Accepted 1 June 2012
Available online 12 June 2012

JEL classification:
G17
G15
C15
C32
C53

Keywords:
Expected shortfall
Long memory
Multi-period forecasting
Value-at-risk
Volatility forecasting

ABSTRACT

The present study compares the performance of the long memory FIGARCH model, with that of the short memory GARCH specification, in the forecasting of multi-period value-at-risk (VaR) and expected shortfall (ES) across 20 stock indices worldwide. The dataset is composed of daily data covering the period from 1989 to 2009. The research addresses the question of whether or not accounting for long memory in the conditional variance specification improves the accuracy of the VaR and ES forecasts produced, particularly for longer time horizons. Accounting for fractional integration in the conditional variance model does not appear to improve the accuracy of the VaR forecasts for the 1-day-ahead, 10-day-ahead and 20-day-ahead forecasting horizons relative to the short memory GARCH specification. Additionally, the results suggest that underestimation of the true VaR figure becomes less prevalent as the forecasting horizon increases. Furthermore, the GARCH model has a lower quadratic loss between actual returns and ES forecasts, for the majority of the indices considered for the 10-day and 20-day forecasting horizons. Therefore, a long memory volatility model compared to a short memory GARCH model does not appear to improve the VaR and ES forecasting accuracy, even for longer forecasting horizons. Finally, the rolling-sampled estimated FIGARCH parameters change less smoothly over time compared to the GARCH models. Hence, the parameters' time-variant characteristic cannot be entirely due to the news information arrival process of the market; a portion must be due to the FIGARCH modelling process itself.

1. Introduction — motivation and review of literature

The recent financial crisis has emphasised the importance for financial institutions of producing reliable value-at-risk (VaR) and expected shortfall (ES) forecasts. VaR quantifies the maximum amount of loss for a portfolio of assets, under normal market conditions over a given period of time and at a certain confidence level. ES quantifies the expected value of the loss, given that a VaR violation has occurred.

Following the recommendations of the Basel Committee on Banking Supervision (1996, 2006), many financial institutions have flexibility over their choice of model for estimating VaR. The guidelines prescribe, however, that financial institutions should use up to one year of data to calculate the VaR of their portfolios for a ten-day holding period.[1] The Basel Committee recommend producing multi-step VaR forecasts by scaling up the daily VaR figure using the square root of time rule.[2] However, this method is criticised in the literature, with Engle (2004) noting that it makes the invalid assumption that volatilities over time are constant. Further, Rossignolo, Fethi, and Shaban (in press) give emphasis to both the current (Basel II[3]) and proposed regulations (Basel III[4]) with regard to VaR estimation. Focusing on 1-trading-day VaR, they compare results from current and proposed regulations and suggest that heavy-tailed distributions are the most accurate technique to model market risks.

The majority of existing models for forecasting VaR and ES are focused on producing accurate forecasts for 1-trading-day. An enormous variety of VaR models have been tested in the literature,

* Corresponding author at: Department of Economics, Portsmouth Business School, University of Portsmouth, Richmond Building, Portland Street, Portsmouth, PO1 3DE, UK.
E-mail address: stavros.degiannakis@gmail.com (S. Degiannakis).

[1] Following the financial crash, amendments to the regulations were announced, necessitating financial institutions to calculate a 'stressed value-at-risk' measure, using data covering a year of trading in which the financial institution incurred significant losses (Basel Committee on Banking Supervision, 2009).

[2] To account for the non-linear price characteristics of option contracts, financial institutions are expected to move towards calculating a full 10-day VaR for positions involving such contracts.

[3] Basel II VaR quantitative requirements include: (a) daily-basis estimation; (b) confidence level set at 99%; (c) one-year minimum sample extension with quarterly or more frequent updates; (d) no specific models prescribed: banks are free to adopt their own schemes; (e) regular backtesting and stress testing programme for validation purposes, see Rossignolo et al. (in press).

[4] Basel III captures fat-tail risks (that most VaR models are not able to do under Basel II) by introducing a stressed VaR (sVaR) metric to increase the minimum capital requirements (MCR), see Rossignolo et al. (in press).

1057-5219/$ – see front matter © 2012 Elsevier Inc. All rights reserved.
doi:10.1016/j.irfa.2012.06.001

including both parametric and non-parametric models. The results have not been entirely consistent, often suggesting that the optimum choice of model, as well as the distributional assumptions, may depend upon a number of factors including the market for which the model is being estimated, the length and frequency of the data series, and whether or not the VaR relates to long or short trading positions (Angelidis, Benos, & Degiannakis, 2004; Shao, Lian, & Yin, 2009).

The Generalised Autoregressive Conditionally Heteroskedastic (GARCH) model has been shown in the literature to produce reasonable low and high frequency VaR forecasts across a variety of markets and under different distributional assumptions. For example Srinanthakumar and Silvapulle (2003) estimate the VaR for daily returns and select the simple GARCH(1,1) model with Student-*t* errors as the preferred model. Some studies have concluded that the use of a skewed, rather than a symmetrical, distribution for the standardised residuals produces superior VaR forecasts. For example, Giot and Laurent (2003, 2004) find the skewed Student-*t* APARCH model to be superior to other specifications for estimating both in-sample and out-of-sample VaR. On the other hand, Angelidis and Degiannakis (2007) conclude that the Student-*t* and skewed Student-*t* overestimate the true VaR, and consequently other distributions such as the normal may be more appropriate for the standardised residuals. There is some debate over the relative merits of conditional volatility models compared to other specifications. Whilst, Danielsson and Morimoto (2000) find that conditional volatility models produce more volatile VaR predictions, Kuester, Mittnik, and Paolella (2006) conclude that the VaR violations arising from unconditional VaR models do not occur independently throughout the estimation period, but may be clustered together.

Accounting for long memory and asymmetries in the conditional volatility process has been shown to improve VaR and ES forecasting accuracy for short (1-day and 5-day) forecasting horizons (Angelidis & Degiannakis, 2007; Härdle & Mungo, 2008).

Recently, Halbleib and Pohlmeier (2012) propose a methodology of computing VaR based on the principle of optimal combination that accurately predicts losses during periods of high financial risk. They develop data-driven VaR approaches that provide robust VaR forecasts; the examined methods include the ARMA-GARCH, RiskMetrics™ and ARMA-FIGARCH. They argue that popular VaR methods perform very differently from calm to crisis periods. Further, they show that, in the case of 1-day VaR forecasts, proper distributional assumptions (Student-*t* with estimated degrees of freedom, skewed Student-*t* and extreme value theory), deliver better quantile estimates and VaR forecasts.

Rossignolo et al. (in press) give a detailed theoretical description of the regulatory framework (Basel II and III Capital Accord) as well as a synopsis of VaR models. Using data from 10 stock market blue-chip indices of six emerging markets (Brazil, Hungary, India, Czech Republic, Indonesia and Malaysia) and four frontier markets (Argentina, Lithuania, Tunisia and Croatia), they argue that "No improvement is virtually recorded employing a heavy-tailed *t* distribution instead of the normal one as the underlying risk measure is inherently flawed". Further, they show that the EGARCH technique brings no significant advantage over the GARCH method for daily time horizon.

Finally, Chen and Lu (2010) review the robustness and accuracy of several VaR estimation methods, under normal, Student-*t* and normal inverse Gaussian (NIG) distributional assumptions, and further test both the unconditional and conditional coverage properties of all the models using the Christoffersen's test, the Ljung–Box test and the dynamic quantile test. Using data from Dow Jones Industrial, DAX 30 and Singapore STI, they argue that conditional autoregressive VaR (CAViaR) and the NIG-based estimation are robust and deliver accurate VaR estimation for the 1-day forecasting interval, whilst the filtered historical simulation (FHS) and filtered EVT perform well for the 5-day forecasting interval.[5]

The aim of this paper is to test empirically whether the short memory GARCH model is outperformed for forecasting multi-period VaR for longer time horizons (10-day and 20-day) by the long memory FIGARCH model, which accounts for the persistence of financial volatility (Baillie, Bollerslev, & Mikkelsen, 1996; Bollerslev & Mikkelsen, 1996; Nagayasu, 2008).[6]

The FIGARCH specification has been shown in some empirical studies to produce superior VaR forecasts (Caporin, 2008; Tang & Shieh, 2006). However, these contrast with the findings of McMillan and Kambouroudis (2009) who conclude that the FIGARCH (as well as the RiskMetrics™ and HYGARCH) specifications are adequate to forecast the volatility of smaller emerging markets at a 5% significance, but that the APARCH model is superior for modelling a 99% VaR.

Recently, attention has turned towards extending the existing literature on the accuracy of various modelling specifications to produce one-step-ahead VaR forecasts, and to formulate reliable modelling techniques for multi-step-ahead VaR forecasts. For example, historical simulation using past data on the sensitivity of the assets within a portfolio to macroeconomic factors has been used to estimate 1-day and 10-day VaRs (Semenov, 2009). Furthermore, a Monte Carlo simulation has been shown to produce useful estimates of intra-day VaR using tick-by-tick data (Brooks & Persand, 2003; Dionne, Duchesne, & Pacurar, 2009).

The empirical analysis in this paper makes use of an adaptation of the Monte Carlo simulation technique of Christoffersen (2003) for estimating multiple-step-ahead VaR and ES forecasts to the FIGARCH model. This enables comparisons to be made between the forecasting performances of the GARCH and FIGARCH models for i) 1-step-ahead, ii) 10-step-ahead and iii) 20-step-ahead VaR and ES predictions. The 95% VaR and 95% ES forecasting performances of the GARCH and FIGARCH models are tested on daily data across 20 leading stock indices worldwide.

This study further provides evidence for the time-variant characteristic of the estimated parameters.[7] In particular, this paper contributes to the debate on the out-of-sample forecast performance of fractionally integrated models (see Ellis & Wilson, 2004). The out-of-sample forecast performance of the GARCH and FIGARCH models is investigated in order to examine (i) whether the FIGARCH model provides superior multi-period VaR and ES forecasts and (ii) in what extent do the rolling-sampled estimated parameters confirm a time-variant characteristic (see Degiannakis, Livada, & Panas, 2008).

We show that i) the long memory FIGARCH model, as compared to the short memory GARCH model, does not appear to improve the VaR and ES forecasting accuracy and ii) the estimated parameters of the models present a time-varying characteristic, which can be linked to market dynamics in response to the unexpected news. However, the estimated parameters of the FIGARCH model exhibit relatively a more time-varying characteristic than those of the GARCH model, inferring evidence that not all of the time-varying characteristics can be due to the news information arrival process of the market. These findings are similar to those of Ellis and Wilson (2004) who argue that fractionally integrated models for forecasting the conditional mean of financial asset returns (i.e. ARFIMA model) fail to outperform forecasts derived from short memory models.

Furthermore, we conclude that the models should be constructed carefully, either by risk managers or by market regulators. The ES

[5] Chen and Lu (2010) show that NIG works well if the market is normal, whereas the method provides low accurate VaR values within a financial crisis period.

[6] It should be recognised that some authors suggest that accounting for structural breaks in volatility (Granger & Hyung, 2004), or allowing the unconditional variance to change over time (McMillan & Ruiz, 2009) can reduce the strength of the evidence in favour of the persistence of financial volatility.

[7] To this end, we allow the standardised residuals of the model to follow the relatively parsimonious normal distribution, since we are only interested in comparing the effects of modelling for short memory and long memory on the VaR and ES forecasting accuracy. The normal model has been shown by Angelidis and Degiannakis (2007) to be preferable to more parameterised distributions for the standardised residuals in some cases.