ARTICLE IN PRESS

International Journal of Forecasting ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Selecting volatility forecasting models for portfolio allocation purposes

R. Becker[b], A.E. Clements[a], M.B. Doolan[a,*], A.S. Hurn[a]

[a] School of Economics and Finance, Queensland University of Technology, Australia
[b] Economics, School of Social Sciences, University of Manchester, United Kingdom

ARTICLE INFO

Keywords:
Multivariate time series
Loss functions
Evaluating forecasts
Covariance matrix
GARCH models
Model confidence set

ABSTRACT

Techniques for evaluating and selecting multivariate volatility forecasts are not yet understood as well as their univariate counterparts. This paper considers the ability of different loss functions to discriminate between a set of competing forecasting models which are subsequently applied in a portfolio allocation context. It is found that a likelihood-based loss function outperforms its competitors, including those based on the given portfolio application. This result indicates that considering the particular application of forecasts is not necessarily the most effective basis on which to select models.

1. Introduction

A recent survey by Amenc, Goltz, Tang, and Vaidyanathan (2012) of 139 North American Investment Managers representing $12 trillion worth of assets under management reports that the majority of fund managers use volatility and correlation forecasts to construct equity portfolios. Given the range of models which are capable of forecasting multivariate volatility, it can be inferred that these managers must all apply some discriminatory procedure when selecting a preferred forecasting model. It follows, therefore, that the process of evaluating volatility forecasts for the purpose of model selection is of enormous practical importance.

The literature on multivariate volatility modeling is extensive. A comprehensive survey of volatility modeling is provided by Andersen, Bollerslev, Christoffersen, and Diebold (2006), while Bauwens, Laurent, and Rombouts (2006), and Silvennoinen and Teräsvirta (2009) survey the multivariate versions of the popular generalised autoregressive conditional heteroskedastic (GARCH) models. Broadly speaking, the papers surveyed introduce new models, specify variations to an existing model or outline different estimation procedures, or provide some combination of these facets. It is fair to say that these studies have developed models that capture the stylized facts of volatility and produce adequate forecasts in relatively large-dimensional problems. Nevertheless, despite these advances, no model or group of models has yet emerged as the obvious method of choice. The fact remains that the investment manager must still select a forecasting model from a vast array of competing models.

Model selection generally involves the evaluation of forecasts of volatility within loss functions, which are classified as either direct or indirect by Patton and Sheppard (2009). Direct loss functions are measures of the forecast accuracy based on traditional statistical measures of precision. Although it would appear that direct loss functions should be easy to implement and interpret, the fact that volatility is unobservable, thereby necessitating the use of an observable proxy for volatility, confounds the issue. Indeed, Hansen and Lunde (2006) and Patton (2011) demonstrate that noise in the volatility proxy renders certain direct loss functions incapable of ranking forecasts consistently in the univariate setting. Subsequent studies by Laurent, Rombouts, and Violante (2013) and Patton and Sheppard (2009) have reported equivalent results in the multivariate setting.

* Correspondence to: School of Economics and Finance, Queensland University of Technology, Brisbane, 4001 Queensland, Australia.
E-mail address: m.doolan@qut.edu.au (M.B. Doolan).

0169-2070/$ – see front matter © 2014 International Institute of Forecasters. Published by Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ijforecast.2013.11.007

On the other hand, indirect measures of the volatility forecasting performance evaluate forecast efficacy in the context of the application for which the forecast is required, for example the portfolio allocation problem. One appealing attribute of this type of evaluation is that it is specifically related to the economic decision from which the forecast derives its value (Elliott & Timmermann, 2008). Danielsson (2011, p. 44) argues that forecasts should be evaluated and selected on the basis of their intended application. Many studies have used indirect measures to evaluate volatility forecasting models. For example, Engle and Colacito (2006) evaluate the forecasting performance in terms of portfolio return variance, while Fleming, Kirby, and Ostdiek (2001, 2003) apply a quadratic utility function that values one forecast relative to another. Despite the strong economic appeal of measures that combine risk and return, especially those that report a measure of relative economic value, it is easy to show that these measures can favour incorrect forecasts of volatility. One notable exception is the portfolio variance, which does not display this problem. Engle and Colacito (2006) and Patton and Sheppard (2009) have demonstrated that the portfolio variance is minimised when the correct forecast is applied; a result that links the portfolio variance with robust statistical loss functions.

This paper extends the previous literature by considering the role played by loss functions in ex-ante multivariate volatility model selection, where forecasts from these models will subsequently be used in mean–variance portfolio optimisation. In doing so, it will assess the ability of a range of loss functions to discriminate between volatility forecasting models where the intended use of the forecasts is a portfolio optimisation problem. While this paper focuses on mean–variance portfolio optimisation, there is nothing to prevent the consideration of higher moments of returns, estimation error or portfolio constraints. However, the main focus here is not the final application itself, but rather the way in which the loss functions perform in terms of model selection with a given application in mind. This is achieved in part by a simulation study that considers the relative powers of a range of statistical loss functions and portfolio variances. Power is important in the context of this problem because it reflects the ability of loss functions to discriminate between forecasts. A subsequent empirical study then assesses the consistency between the various loss functions and the final portfolio application. It will gauge whether the best models selected from the evaluation period continue to be the best performers in the application period, the optimal outcome in terms of model selection. This differs from the traditional forecast evaluation literature in that it considers the use of statistical measures to discriminate between models, the performances of which are then measured based on an economic criterion in a subsequent period.

Very briefly, two important general results emerge from the research reported in this paper. The first is that a likelihood-based loss function is preferred when selecting models whose forecasts will subsequently be applied in a portfolio allocation context. Using portfolio-variance-based evaluation to select models does not appear to provide a strong enough discrimination between the competing models. This result suggests that selecting models solely on the basis of their intended use may not be an optimal strategy. The second interesting result is that more precise volatility proxies do not necessarily improve model selection outcomes. While a high-frequency data proxy for volatility (realised volatility) improves the ability to discriminate between models, it does not improve model selection outcomes in the context of the portfolio application. That is, while fewer models may be selected with the more precise volatility proxy, there are instances where these models do not perform particularly well when evaluated in the context of the portfolio application. In those instances where the more precise volatility proxy selects superior models in the portfolio application, the outcome is the same as that from the less precise proxy. Overall, these findings should assist practitioners who are involved in forecasting multivariate volatility for the purposes of portfolio construction.

The paper proceeds as follows. Section 2 outlines the portfolio allocation problem and the loss functions used to evaluate forecasts. The econometric methodology used to distinguish between forecasts is described in Section 3. Section 4 reports simulation evidence relating to the ability of a number of loss functions to distinguish between competing forecasts. Section 5 provides an empirical illustration of the model selection problem. Here, models will be selected from an evaluation period using a range of loss functions, and their performances compared in the subsequent portfolio application. The outcome of this exercise identifies the optimal loss functions to use for model selection. Section 6 is a brief conclusion.

2. Forecast application and evaluation

This section describes the intended portfolio application and the loss functions employed to evaluate the volatility forecasts. Begin by considering a system of N asset excess returns

$$r_t = \mu_t + \varepsilon_t, \quad \varepsilon_t \sim F(0, \Sigma_t), \tag{1}$$

where r_t is an $N \times 1$ vector, μ_t is an $N \times 1$ vector of expected excess returns and ε_t is an $N \times 1$ vector of disturbances following the multivariate distribution F.

In this context, the optimization problem of an investor who seeks to minimise the variance of a portfolio of N risky assets and a risk-free asset is

$$\min_{w_t} \; w_t' \Sigma_t w_t \quad \text{s.t. } w_t' \mu_t = \mu_0, \tag{2}$$

where w_t is an $N \times 1$ vector of portfolio weights and μ_0 is the target excess return for the portfolio. The unconstrained solution to the problem posed in Eq. (2) is

$$w_t = \frac{\Sigma_t^{-1}\mu_t}{\mu_t' \Sigma_t^{-1}\mu_t}\, \mu_0. \tag{3}$$

In this setting, $1 - w_t' \iota$ represents the proportion of wealth invested in the risk-free asset and ι is an $N \times 1$ vector of ones. One may avoid making any assumptions regarding the vector of expected excess returns, μ_t, by considering the global minimum variance portfolio for risky assets