Journal of Econometrics 160 (2011) 93–101

Contents lists available at ScienceDirect

Journal of Econometrics

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

Forecasting multivariate realized stock market volatility

Gregory H. Bauer[a,*], Keith Vorkink[b]

[a] Financial Markets Department 4E, Bank of Canada, 234 Wellington, Ottawa, Ontario, Canada K1A 0G9

[b] 667 TNRB, Provo, UT 84602, USA

ARTICLE INFO

Article history:
Available online 6 March 2010

JEL classification:
G14
C53
C32

Keywords:
HAR-RV model
Realized volatility
Covariance matrix
Factor model

ABSTRACT

We present a new matrix-logarithm model of the realized covariance matrix of stock returns. The model uses latent factors which are functions of lagged volatility, lagged returns and other forecasting variables. The model has several advantages: it is parsimonious; it does not require imposing parameter restrictions; and, it results in a positive-definite estimated covariance matrix. We apply the model to the covariance matrix of size-sorted stock returns and find that two factors are sufficient to capture most of the dynamics.

1. Introduction

The variances and covariances of stock returns vary over time (e.g. Andersen et al., 2005). As a result, many important financial applications require a model of the conditional covariance matrix. Three distinct categories of methods for estimating a latent conditional covariance matrix have evolved in the literature. In the first category are the various forms of the multivariate GARCH model where forecasts of future volatility depend on past volatility and shocks (e.g. Bauwens et al., 2006). In the second category, authors have modeled asset return variances and covariances as functions of a number of predetermined variables (e.g. Ferson, 1995). The third category includes multivariate stochastic volatility models (e.g. Asai et al., 2006).

In this paper, we introduce a new model of the *realized* covariance matrix.[1] We use high-frequency data to construct estimates of the daily realized variances and covariances of five size-sorted stock portfolios. By using high-frequency data we obtain an estimate of the matrix of 'quadratic variations and covariations' that differs from the true conditional covariance matrix by mean zero errors (e.g. Andersen et al. (2003) and Barndorff-Nielsen and Shephard (2004a)). This provides greater power in determining the effects of alternative forecasting variables on equity market volatility when compared to efforts based on latent volatility models.

We transform the realized covariance matrix using the matrix logarithm function to yield a series of transformed volatilities which we term the *log-space volatilities*. The matrix logarithm is a non-linear function of all of the elements of the covariance matrix and thus the log-space volatilities do not correspond one-to-one with their counterparts in the realized covariance matrix.[2] However, modeling the time variation of the log-space volatilities is straightforward and avoids the problems that plague existing estimators of the latent volatility matrix. In particular, we do not have to impose any constraints on our estimates of the log-space volatilities.

We then model the dynamics of the log-space volatility matrix using a latent factor model. The factors consist of *both* past volatilities and other variables that can help forecast future volatility. We thus are able to model the conditional covariance matrix by combining a large number of forecasting variables into a relatively small number of factors. Indeed we show that two factors can capture the volatility dynamics of the size-sorted stock portfolios.

The factor model is estimated by GMM yielding a series of filtered estimates. We then transform these fitted values, using the matrix exponential function, back into forecasts of the realized covariance matrix. Our estimated matrix is positive

* Corresponding author.
E-mail addresses: gbauer@bankofcanada.ca (G.H. Bauer), keith_vorkink@byu.edu (K. Vorkink).

1 Andersen et al. (2001) and Barndorff-Nielsen and Shephard (2002) formalized the notion of *realized volatility*.

2 The matrix logarithm has been used for estimators of latent volatility by Chiu et al. (1996) and Kawakatsu (2006) and was also suggested in Asai et al. (2006).

0304-4076/$ – see front matter
doi:10.1016/j.jeconom.2010.03.021