Journal of Econometrics 183 (2014) 181–192

Contents lists available at ScienceDirect

Journal of Econometrics

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

The VIX, the variance premium and stock market volatility[☆]

Geert Bekaert[a], Marie Hoerova[b,*]

[a] Columbia Business School and NBER, United States
[b] European Central Bank, Germany

ARTICLE INFO

Article history:
Available online 20 May 2014

JEL classification:
C22
C52
G12
E32

Keywords:
Option implied volatility
Realized volatility
VIX
Variance risk premium
Risk aversion
Stock return predictability
Risk–return trade-off
Economic uncertainty
Financial instability

ABSTRACT

We decompose the squared VIX index, derived from US S&P500 options prices, into the conditional variance of stock returns and the equity variance premium. We evaluate a plethora of state-of-the-art volatility forecasting models to produce an accurate measure of the conditional variance. We then examine the predictive power of the VIX and its two components for stock market returns, economic activity and financial instability. The variance premium predicts stock returns while the conditional stock market variance predicts economic activity and has a relatively higher predictive power for financial instability than does the variance premium.

1. Introduction

The 2007–2009 crisis has intensified the need for indicators of the risk aversion of market participants. It has also become increasingly commonplace to assume that changes in risk appetites are an important determinant of asset prices. Not surprisingly, the behavioral finance literature (see e.g. Baker and Wurgler, 2007) has developed "sentiment indices", and financial institutions have created a wide variety of "risk aversion" indicators (see Coudert and Gex, 2008, for a survey).

One simple candidate indicator is the equity variance premium, the difference between the squared VIX index and an estimate of the conditional variance of the stock market. The VIX index is the "risk-neutral" expected stock market variance for the US S&P500 contract and is computed from a panel of options prices. Well-known as a "fear index" (Whaley, 2000) for asset markets, it reflects both stock market uncertainty (the "physical" expected volatility), and a variance risk premium, which is also the expected premium from selling stock market variance in a swap contract. Bollerslev et al. (2009) show that an estimate of this variance premium predicts stock returns; Bekaert et al. (2013) show that there are strong interactions between monetary policy and the variance premium, suggesting that monetary policy may actually affect risk aversion in the market place. The variance premium uses objective financial market information and naturally "cleanses" option-implied volatility from the effect of physical volatility dynamics and uncertainty, leaving a measure correlated with risk aversion.

How to measure the variance premium is not without controversy, however, because it relies on an estimate of the conditional variance of stock returns. For example, the measure proposed in Bollerslev et al. (2009), BTZ, henceforth, assumes that the conditional variance of stock market returns is a martingale, an assumption which is not supported by the data, leading to potentially biased variance premiums. In this paper, we tackle several measurement issues for the variance premium, assessing a plethora of state-of-the-art volatility models and making full use of overlapping daily data, rather than sparse end-of-month data, which is standard.

☆ We thank the Editor, Alok Bhargava, and three anonymous reviewers for useful comments and suggestions that improved the paper. Carlos Garcia de Andoain provided excellent research assistance. Geert Bekaert acknowledges financial support from Netspar. The views expressed do not necessarily reflect those of the European Central Bank or the Eurosystem.

* Correspondence to: European Central Bank, Kaiserstrasse 29, 60311 Frankfurt am Main, Germany. Tel.: +49 69 1344 8710; fax: +49 69 1344 8552.
E-mail address: marie.hoerova@ecb.int (M. Hoerova).

http://dx.doi.org/10.1016/j.jeconom.2014.05.008
0304-4076/

The conditional variance measure is of interest in its own right. First, there is a long literature on the trade-off between risk, as measured by the conditional variance of stock market returns, and the aggregate risk premium on the market (see e.g. French et al., 1987, for a seminal contribution). This long line of research has mostly failed to uncover a strong positive relationship between risk and return (see Bali, 2008, for a summary). Second, stock market volatility can also be viewed as a market-based measure of economic uncertainty. For example, Bloom (2009) shows that heightened "economic uncertainty" decreases employment and output. Interestingly, he uses the VIX index to measure uncertainty, so that his results may actually be driven by the variance premium rather than uncertainty per se.

Using more plausible estimates of the variance premium and stock market volatility, we then assess whether they predict stock returns, economic activity, as well as financial instability, an economic outcome whose monitoring is of considerable policy interest. We find that the well-known results in BTZ exaggerate the predictive power of the variance premium for stock returns. However, the equity variance risk premium remains a reliable predictor of stock returns. Stock market volatility does not predict the stock market, but it is a much better predictor of economic activity than is the equity variance premium. It also predicts financial instability more strongly than does the variance premium, especially at longer horizons.

The remainder of the paper is organized as follows. Section 2 discusses the econometric framework that we use to forecast volatility, and lays out our model selection procedure. Section 3 reports the results of our specification analysis and forecasting performance comparison. Section 4 uses the preferred estimates of the variance premium and stock market volatility to predict stock returns, economic activity and financial instability. Section 5 concludes.

2. Econometric framework

We define the variance risk premium as:

$$VP_t = VIX_t^2 - E_t\left[RV_{t+1}^{(22)}\right]. \tag{1}$$

Here the *VIX* is the implied option volatility of the S&P500 index for contracts with a maturity of one month, and $RV_{t+1}^{(22)}$ is the S&P500 realized variance measured over the next month (22 trading days) using five-minute returns. Note that $RV_{t+1}^{(22)} - VIX_t^2$ is the return to buying variance in a variance swap contract. Therefore, technically speaking, the variance risk premium refers to the negative of VP (see Carr and Wu, 2009). Since that number is mostly negative, we prefer to define it as we did in Eq. (1).

Economically, the squared *VIX* is the conditional return variance using a "risk-neutral" probability measure, whereas the conditional variance uses the actual "physical" probability measure. The risk-adjusted measure shifts probability mass to states with higher marginal utility (bad states) and this implies that in many realistic economic settings, the variance premium will be increasing in the economy's risk aversion.

The unconditional mean of the variance premium is easy to compute by simply computing the average of $VIX_t^2 - RV_{t+1}^{(22)}$. However, we are interested in the conditional variance premium as described in Eq. (1), which relies on the physical conditional expected value of the future realized variance. The common approach to estimate this uses empirical projections of the realized variance on variables in the information set, and subtracts this estimated expected variance from the VIX^2 to arrive at VP. Hence, the problem is reduced to one of variance forecasting.

Our data start on January 02, 1990 (the start of the model-free VIX series)[1] and covers the period until October 01, 2010. We have a total of 5208 daily, overlapping observations. The recent crisis period presents special challenges as stock market volatilities peaked at unprecedented levels, but at the same time the crisis represents an informative period during which uncertainty and risk aversion may have been particularly pronounced. Nevertheless, if we decompose the sample variance of the implied and realized volatility series in contributions by crisis and non-crisis observations, the crisis observations dominate despite representing a relatively small part of the sample. We deal with the crisis-induced challenges by considering both models that predict the level and the logarithm of realized variances, and by putting much emphasis on parameter stability in our model selection procedure. In addition, we focus on out-of-sample forecasting exercises where we conduct the in-sample estimations mostly on non-crisis observations, so that the influence of the crisis on the parameter estimates and model selection is mitigated.

Variance forecasting

There is an extensive econometric literature on volatility forecasting. It is now generally accepted that models based on high frequency realized variances dominate standard models in the GARCH class (see e.g. Chen and Ghysels, 2012) and we therefore examine the state-of-the-art models in that class. These models stress the importance of persistence (using lagged realized variances as predictors), additional information content in the most recent return variances (Corsi, 2009), asymmetry between positive and negative return shocks (the classic volatility asymmetry, see e.g. Engle and Ng, 1993) and potentially differing predictive information present in jump versus continuous volatility components (Andersen et al., 2007). We accommodate all of these elements in our model.

In the finance literature, it has been pointed out as early as in Christensen and Prabhala (1998) that option prices as reflected in implied volatility should have information about future stock market volatility. This motivates using the *VIX* as a predictive variable. Recent articles using the *VIX* in similar forecasting exercises include Busch et al. (2011) who examine a number of variance forecasting models embedding option-implied volatility for bond, currency and stock markets, and Andersen and Bondarenko (2007) who mostly focus on measurement issues with the officially published *VIX* index. Of course, because the *VIX* also embeds a risk premium, it will not be an unbiased predictor of future realized volatility. Chernov (2007) argues that spot volatility is likely to have additional information about future volatility.

Finally, it is well-known that estimation noise hurts out-of-sample forecasting performance. Simple models such as the martingale model may therefore outperform more complex models. We therefore also consider a number of non-estimated models that are special cases of our general framework.

Our most general forecasting model can be represented as follows:

$$\begin{aligned} RV_t^{(22)} = {} & c + \alpha VIX_{t-22}^2 + \beta^m C_{t-22}^{(22)} + \beta^w C_{t-22}^{(5)} + \beta^d C_{t-22}^{(1)} \\ & + \gamma^m J_{t-22}^{(22)} + \gamma^w J_{t-22}^{(5)} + \gamma^d J_{t-22}^{(1)} + \delta^m r_{t-22}^{(22)-} \\ & + \delta^w r_{t-22}^{(5)-} + \delta^d r_{t-22}^{(1)-} + \varepsilon_t. \end{aligned} \tag{2}$$

[1] The CBOE changed the methodology for calculating the VIX, initially measuring implied volatility for the S&P100 index, to be measured in a model-free manner from a panel of option prices (see Bakshi et al., 2003, for details) only in September 2003. It then backdated the new model-free index to 1990 using historical option prices (see CBOE, 2004, for more details).