Journal of Banking & Finance 43 (2014) 200–211

Contents lists available at ScienceDirect

Journal of Banking & Finance

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

Estimating and using GARCH models with VIX data for option valuation

Juho Kanniainen[a,*], Binghuan Lin[a,b], Hanxue Yang[a]

[a] Department of Industrial Management, Tampere University of Technology, P.O. Box 541, FI-33101 Tampere, Finland
[b] Techila Technologies Ltd, Itsenäisyydenkatu 2, FI-33100 TAMPERE, Finland

ARTICLE INFO

Article history:
Received 2 April 2013
Accepted 23 March 2014
Available online 2 April 2014

Keywords:
Option valuation
VIX
GARCH
Estimation

ABSTRACT

This paper uses information on VIX to improve the empirical performance of GARCH models for pricing options on the S&P 500. In pricing multiple cross-sections of options, the models' performance can clearly be improved by extracting daily spot volatilities from the series of VIX rather than by linking spot volatility with different dates by using the series of the underlying's returns. Moreover, in contrast to traditional returns-based Maximum Likelihood Estimation (MLE), a joint MLE with returns and VIX improves option pricing performance, and for NGARCH, joint MLE can yield empirically almost the same out-of-sample option pricing performance as direct calibration does to in-sample options, but without costly computations. Finally, consistently with the existing research, this paper finds that non-affine models clearly outperform affine models.

1. Introduction

The empirical option pricing performance of the GARCH family models has been well studied in the recent literature (see Christoffersen et al., 2012, and references therein). For option valuation, GARCH model parameters are often estimated by the Maximum Likelihood Estimation (MLE) method using return series, Non-linear Least-Squares (NLS) on (multiple) cross-sections of option data, or both returns and option data. However, the MLE approach with returns only does not necessarily yield good estimates for option pricing; therefore, one might prefer to estimate structural parameters directly using information on option price observations (see, e.g., Christoffersen and Jacobs, 2004; Christoffersen et al., 2012). On the other hand, as pointed out, for example, in (Broadie and Detemple, 2004 and Duan and Yeh, 2010), model calibration on option prices over a long period can result in challenging and costly computations, especially if a closed-form analytical solution is not available. Importantly, in most GARCH models, no (semi-closed) solutions are available for option valuation, and especially with non-affine models, option prices can be computed only through Monte Carlo simulation or by using approximations (Duan et al., 2006).

In this paper, we investigate alternative approaches to estimating parameters and to valuating index options using information on the VIX index, aiming in general (i) to improve the option pricing performance of GARCH models and (ii) to reduce the computational burden. Many VIX-related papers consider VIX derivatives (see, for example, Lin and Chang, 2010, and references therein) or volatility-forecasting (see, for example, Poon and Granger, 2003, and references therein), but here the goal is different: *to use information on the VIX index to estimate GARCH models and to improve their performance for pricing multiple cross-sections of options on the S&P 500.*

A few recent papers on continuous time volatility models have focused on estimation procedures using volatility proxies constructed from volatility indices such as VXO and VIX (see, for example, Jones, 2003; Bakshi et al., 2006; Aït-Sahalia and Kimmel, 2007; Duan and Yeh, 2010; Kanniainen, 2011). Aït-Sahalia and Kimmel (2007) provide a maximum likelihood estimator for three continuous time models using VIX as a volatility proxy, yet their closed-form results can be easily applied also to other popular continuous time-stochastic volatility models. Duan and Yeh (2010) also used information from the VIX index jointly with returns on the S&P 500 and introduced a maximum likelihood estimation method for a class of continuous-time stochastic volatility models in a jump diffusion framework. Moreover, Duan and Yeh (2012) developed an estimation method to capture the VIX term-structure jointly with returns. However, these papers did not investigate empirically, based on observations of option prices on the S&P 500, how inclusion of the VIX index in parameter estimation improved the models' option pricing performance. One welcome exception in continuous-time stochastic volatility is Kaeck and Alexander (2012), who estimated several continuous-time models with the Markov chain Monte Carlo using information

* Corresponding author. Tel.: +358 40 707 4532.
E-mail address: juho.kanniainen@tut.fi (J. Kanniainen).

http://dx.doi.org/10.1016/j.jbankfin.2014.03.035
0378-4266/

on the VIX term-structure and testing parameter estimates with extensive option data samples.

Most importantly, to our knowledge, the VIX index has not been widely used to estimate or calibrate discrete-time GARCH models, perhaps because the spot volatility state is already inherently available when historical asset returns are used for given GARCH structural parameters. A welcome exception are Hao and Zhang, 2013, who recently proposed a joint likelihood estimation with returns and VIX. However, they used no option data to compare and verify the option pricing performance of different estimation methodologies, and in addition and unlike in this paper, assumed no autocorrelation in disturbances.

As we show in this paper, with GARCH models, we must implement MLE with autoregressive disturbances, because VIX errors (differences between observations and model values) are highly autocorrelated. We show that proper inclusion of information on VIX in MLE can substantially decrease the option pricing error over traditional MLE with returns data only. Even more interestingly, so far in the literature models have been calibrated to option prices over multiple days by linking spot volatility with different dates by using return series. However, the empirical evidence in this paper shows that, in fact, spot volatilities should be extracted from VIX rather than from returns. According to our extensive empirical analysis, using VIX (rather than return series) with multiple cross-sections of options can substantially improve the models' option pricing performance.

Our approaches are based on the fact that VIX approximates the 30-day variance swap rate on the S&P 500 index (see, e.g., Carr and Wu, 2006; Bollerslev et al., 2011).[1] The first approach is to improve calibrations and option valuation to multiple cross-sections of options. When information on option contracts is used over multiple days, spot volatility on different dates has traditionally been 'updated' using the time series of the underlying's returns for given structural parameter values. Instead of calculating daily spot volatilities from returns series, we suggest extracting spot volatilities from lagged values of VIX. For at least two reasons, this VIX-based volatility extraction may work better than the traditional returns-based approach.

First, both VIX and options prices contain forward-looking information, whereas asset returns do not. On the one hand, as VIX approximates the 30-day variance swap rate, which can be interpreted to measure the risk-neutral expectation of integrated variance within the month (see, for example, Carr and Wu, 2006), this approach provides forward-looking parameter estimates under the risk-neutral measure. VIX, on the other hand, can be regarded as the value of a portfolio of options (while at the same it approximates the variance swap rate), representing aggregated information on option contracts Carr and Wu (2006). Second, the returns-based volatility extraction approach is applied under the physical measure, which may pose a problem, if the price of the return risk cannot be identified from the option data.

In the second approach, we aim to improve MLE estimations by incorporating information on VIX into the likelihood function of the bivariate system with autoregressive disturbances. In order to calculate VIX disturbances, we must solve VIX for given structural parameters and return-based filtered conditional spot volatilities. Parameter estimates are then found by maximizing the joint likelihood of returns and VIX. Because VIX represents a portfolio of options, VIX-based parameter estimates can potentially yield better option pricing performance than pure return-based maximum likelihood estimates, and one can reasonable expect that VIX-implied option pricing errors are not far from minima. Compared to the traditional approach of calibrating GARCH models on option price observations, this joint VIX-Returns-MLE definitely saves computation time, especially with non-affine models, which have no analytical formulae for option prices. Instead of using Monte-Carlo methods to repeatedly value a large set of option prices to minimize the option pricing error, VIX-based parameter estimates can be obtained without computationally expensive option valuations.

This paper is organized as follows. In Section 2, we discuss the GARCH models used in the paper and in Section 3 the calibration approaches using information on option data. In addition, we show how conditional spot volatilities can be alternatively extracted from VIX. Section 4 introduces the joint MLE approach with autoregressive disturbances. In Section 5, we describe our data sets and estimate various GARCH models by four methods and examine the option pricing performance of the different models and estimation methods. The final section discusses the results and draws conclusions.

2. Models

In this section, we introduce three widely recognized specifications we employed in this study. We chose a set of models for comparing different estimation and volatility extraction approaches. In particular, for diversified analysis, we chose the GARCH specifications that incorporate volatility asymmetry differently and do not nest each other. The models are the non-affine model by Glosten et al. (1993); the non-affine NGARCH-specification of Engle and Ng (1993); and the affine model originally proposed by Heston and Nandi (2000). Hereafter, the models are referred to as GJR, NGARCH, and HN, respectively.

The HN specification gives rise to a quasi-closed-form solution to European options, which expedites option valuation, whereas GJR and NGARCH are applied with Monte Carlo methods.[2] On the other hand, in some previous papers, non-affine models outperformed affine models (see, e.g.,Hsieh and Ritchken, 2005; Christoffersen et al., 2010a); therefore, GJR and NGARCH serve as interesting benchmarks for HN. The potential advantage of GJR is that under the risk-neutral measure, the price of the return risk and the physical leverage parameter can be identified from option data, whereas with HN and NGARCH only their combination can be estimated. This is important especially in extracting (updating) spot volatilities over a multi-day data set with returns series.

GJR

With GJR, total return dynamics are

$$\begin{aligned} R_{t+1} &\equiv \ln\left(\frac{S_{t+1}}{S_t}\right) = r + \lambda\sqrt{h_{t+1}} - \frac{1}{2}h_{t+1} + \sqrt{h_{t+1}}z_{t+1}, \\ h_{t+1} &= \beta_0 + h_t\left[\beta_1 + \beta_2 z_t^2 + \beta_3 \max(0, -z_t)^2\right], \end{aligned} \tag{1}$$

where $\beta_0 > 0,\ \beta_1,\ \beta_2,\ \beta_3 \geqslant 0$ for the positive conditional variance and $\lambda > 0$ for the positive equity risk-premium. Moreover, z_t is the iid standard normal random variable. According to Duan's (Duan, 1995) locally risk-neutral pricing framework, total return dynamics can be expressed under the risk-neutral measure as

[1] On September 22, 2003, the CBOE reformulated its VIX index to use the *model-free* implied volatility approach on the S&P 500 and created a historical record for a changed S&P 500 VIX dating back to 1990. The old index, based on the Black–Scholes model (and hence not model-free), was renamed VXO. Specifically, VXO is an average of Black–Scholes implied volatility quotes on eight near-the-money options at two nearby maturities on the S&P 100. For further information, see CBOE Documentation 2003. Note that this study uses VIX, not VXO.

[2] Heston and Nandi (2000) provide a closed-form solution to the characteristic function for future asset prices, but for European option pricing, we must rely on the inversion of this characteristic function, which involves numerical integration.