International Journal of Forecasting 33 (2017) 1105–1123

Contents lists available at ScienceDirect

International Journal of Forecasting

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

Threshold stochastic volatility: Properties and forecasting

Xiuping Mao [a], Esther Ruiz [b,c,*], Helena Veiga [b,c,d]

[a] School of Finance, Zhongnan University of Economics and Law, China
[b] Department of Statistics, Universidad Carlos III de Madrid, Spain
[c] Instituto Flores de Lemus, Universidad Carlos III de Madrid, Spain
[d] BRU-IUL, Instituto Universitario de Lisboa, Portugal

ARTICLE INFO

Keywords:
Conditional heteroscedasticity
Leverage effect
MCMC estimator
Option pricing
Volatility forecasting

ABSTRACT

We analyze the ability of Threshold Stochastic Volatility (TSV) models to represent and forecast asymmetric volatilities. First, we derive the statistical properties of TSV models. Second, we demonstrate the good finite sample properties of a MCMC estimator, implemented in the software package WinBUGS, when estimating the parameters of a general specification, denoted CTSV, that nests the TSV and asymmetric autoregressive stochastic volatility (A-ARSV) models. The MCMC estimator also discriminates between the two specifications and allows us to obtain volatility forecasts. Third, we analyze daily S&P 500 and FTSE 100 returns and show that the estimated CTSV model implies plug-in moments that are slightly closer to the observed sample moments than those implied by other nested specifications. Furthermore, different asymmetric specifications generate rather different European options prices. Finally, although none of the models clearly emerge as best out-of-sample, it seems that including both threshold variables and correlated errors may be a good compromise.

1. Introduction

Stochastic volatility (SV) models are a popular choice for representing the second-order dynamics of financial returns. These models have been generalized to capture the *leverage effect* that is characterized by the asymmetric responses of volatility to positive and negative past returns of the same magnitude; see Black (1976), who was the first to introduce the term 'leverage effect'. One of these generalizations is the threshold stochastic volatility (TSV) model proposed by Breidt (1996) and So, Li, and Lam (2002), which incorporates the leverage effect by allowing the parameters of the log-volatility equation to differ depending on the sign of the lagged returns. TSV models are quite a popular choice for representing the volatility of financial returns; see, among others, Asai and McAleer (2004, 2005, 2011), Chen, Liu, and So (2008, 2013), Chen, So, and Liu (2011), Elliott, Liew, and Siu (2011), Fan and Wang (2013), Ghosh, Gurung, and Source (2015), Liu, Wong, An, and Zhang (2014), Montero-Lorenzo, Fernández-Avilés, and García-centeno (2010), Muñoz, Marquez, and Acosta (2007), Smith (2009), So and Choi (2008, 2009), Tsai-Hung and Wang (2013), Wu and Zhou (2015), Wirjanto, Kolkiewicz, and Men (2016), and Xu (2012). However, to the best of our knowledge, the statistical properties of TSV models are unknown, which makes it difficult to assess their advantages and limitations relative to alternative specifications of the leverage. In particular, it would be interesting to compare TSV models with the popular asymmetric autoregressive SV (A-ARSV) model of Harvey and Shephard (1996) and Taylor (1994), which captures the leverage effect through the correlation between the level and log-volatility disturbances. Given that the properties of TSV models are unknown, only empirical comparisons between these models and alternatives have been carried out. For example, Asai and McAleer (2005) fit TSV and A-ARSV models to four data sets of financial returns and

* Correspondence to: C/ Madrid, 126, 28903, Getafe, Madrid, Spain.
E-mail address: ortega@est-econ.uc3m.es (E. Ruiz).

http://dx.doi.org/10.1016/j.ijforecast.2017.07.001
0169-2070/

conclude that the latter model is superior to the former according to both the AIC and the BIC. Similarly, Smith (2009) compares the two models empirically using the same two criteria plus the Vuong (1989) test, and rejects the TSV model in favour of the A-ARSV model for continuously compounded returns on the value-weighted CRSP portfolio. Wang (2012) also compares the TSV model empirically with the asymmetric SV model proposed by Asai and McAleer (2005), in order to consider different distributions for standardized returns, and concludes that the A-ARSV model has a better fit. Finally, Wu and Zhou (2015) find very weak evidence in favour of a TSV specification.

Results on the forecasting performances of TSV models are also scarce. It is important to analyze the similarities and differences between the volatilities predicted using TSV and A-ARSV models, given the implications for practitioners of having an inadequate specification of the asymmetric response of the volatility. In particular, Christie (1982) shows that equity variances have a strong positive association with both financial leverage and interest rates. Furthermore, different specifications of the leverage lead to different formulae for option pricing, meaning that a misspecification of the leverage might result in incorrect option prices; see Yu, Yang, and Zhang (2006). The specification of the leverage effect may also have consequences for financial prices. For example, Brooks, Henry, and Persand (2002) and Lien (2005) examine its effect on optimal hedge ratios, while Brooks and Persand (2003) consider its impact on the Value-at-Risk (VaR), and Pardo and Torro (2007) explore its potential for profitable holding strategies.

This paper contributes to the literature on asymmetric SV models in three ways. First, we analyze the ability of TSV models to explain the empirical properties that are usually observed with time series of real financial returns, namely excess kurtosis, positive autocorrelations of squares, and negative cross-correlations between returns and future squared returns. We derive closed-form expressions of these moments when the errors have a generalized error distribution (GED), and the constant and variance of the log-volatility noise change depending on whether past returns are smaller or larger than a given threshold. When the autoregressive parameter changes with past returns, we obtain the statistical properties by simulation. We show that the TSV model captures asymmetric conditional heteroscedasticity if the constant of the log-volatility equation changes. If the persistence parameter also changes, the TSV model can generate moments that are close to those that we usually observe when dealing with high frequency financial returns. However, changes in the variance of the log-volatility noise generate series without leverage. We compare the properties of the TSV and A-ARSV models and show that the former generate slightly less leverage than the latter for particular combinations of the parameters.

The second contribution of this paper is to analyze the finite sample performances of Markov chain Monte Carlo (MCMC) estimators of the parameters of TSV and A-ARSV models with GED errors. Although more efficient MCMC estimators exist, we consider the Bayesian software package WinBUGS based on the single-update Gibbs sampler, as described by Meyer and Yu (2000), due to its ease of implementation. We show that it has a good finite sample performance and allows us to discriminate among alternative asymmetric specifications. Furthermore, the MCMC estimator also permits the computation of one-step-ahead volatility forecasts.

Our third contribution is an empirical comparison of the TSV and A-ARSV models fitted to daily S&P 500 and FTSE 100 returns. We show that the estimated SV models with both threshold variables and correlated errors imply plug-in moments that are slightly closer to the observed sample moments than those of models which incorporate leverage through either threshold or correlation. The in-sample volatilities estimated by TSV and A-ARSV models generate rather different European option price, meaning that it is important to fit appropriate specifications of the leverage. Finally, although none of the models emerges as clearly the best in the out-of-sample period, it seems that models that include both threshold and correlation may be a good compromise.

The rest of this paper is organized as follows. Section 2 describes the TSV model and derives its analytical properties when the persistence parameter is fixed. Simulations are carried out in order to analyze its properties when the persistence parameter changes. We also compare the statistical properties of the TSV model with those of the A-ARSV model. Section 3 carries out Monte Carlo experiments analyzing the finite sample properties of the MCMC estimator. Section 4 compares the empirical differences among the models in terms of implied plug-in moments, the pricing of European options and the forecasting of volatility in the context of daily S&P 500 and FTSE 100 returns. Finally, Section 5 concludes the paper.

2. Moments of the threshold SV model

This section describes the TSV model and derives its statistical properties.

2.1. The TSV model

Consider the following TSV model:

$$y_t = \exp(h_t/2)\epsilon_t, \tag{1}$$

$$h_t = \begin{cases} \alpha + \phi h_{t-1} + \sigma_\eta \eta_{t-1}, & \epsilon_{t-1} \geq \delta, \\ \alpha + \alpha_0 + (\phi + \phi_0)h_{t-1} + (\sigma_\eta + \sigma_{\eta 0})\eta_{t-1}, & \epsilon_{t-1} < \delta, \end{cases} \tag{2}$$

where y_t is the return at time t, $\sigma_t \equiv \exp(h_t/2)$ is its volatility, η_t is a standardized Gaussian white noise process, and ϵ_t is an independent and identically distributed sequence with mean zero and variance one that is independent of η_t for all leads and lags. The TSV model incorporates the leverage effect by allowing the parameters of the log-volatility equation to change depending on whether past standardized returns are smaller or larger than the threshold δ.

Several restricted versions of the TSV model in Eqs. (1) and (2) have been considered previously in the literature. For example, Wirjanto et al. (2016) consider a TSV model with $\phi_0 = \sigma_{\eta 0} = 0$ and ϵ_t being Gaussian. Asai and McAleer (2004) further assume that $\delta = 0$. Several other authors also assume that $\delta = 0$ and ϵ_t are Gaussian; see for example Breidt (1996), as well as Lien (2005), who further