ARTICLE IN PRESS

Economic Modelling xx (xxxx) xxxx–xxxx

Contents lists available at ScienceDirect

Economic Modelling

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

Parameter instability, stochastic volatility and estimation based on simulated likelihood: Evidence from the crude oil market

Nima Nonejad

Department of Mathematical Sciences, Aalborg University and CREATES, Fredrik Bajers Vej 7G 9220, Aalborg Ø, Denmark

ARTICLE INFO

JEL:
C11
C22
C58
C63

Keywords:
Bayes
Crude oil
Metropolis-Hastings
Parameter instability
Particle filtering

ABSTRACT

Stochastic volatility models with fixed parameters can be too restrictive for time-series analysis due to instability in the parameters that govern conditional volatility dynamics. We incorporate time-variation in the model parameters for the plain stochastic volatility model as well its extensions with: Leverage, volatility feedback effects and heavy-tailed distributed innovations. With regards to estimation, we rely on one recently discovered result, namely, that when an unbiasedly simulated estimated likelihood (available for example through a particle filter) is used inside a Metropolis-Hastings routine then the estimation error makes no difference to the equilibrium distribution of the algorithm, the posterior distribution. This in turn provides an off-the-shelf technique to estimate complex models. We examine the performance of this technique on simulated and crude oil returns from 1987 to 2016. We find that (i): There is clear evidence of time-variation in the model parameters, (ii): Time-varying parameter volatility models with leverage/Student's t-distributed innovations perform best, (iii): The timing of parameter changes align very well with events such as market turmoils and financial crises.

1. Introduction

Ever since the contributions of Taylor (1986), Harvey et al. (1994) and Kim et al. (1998), stochastic volatility (SV) models have been increasingly used to model the volatility of financial and macroeconomic time-series. In this framework, conditional volatility is modeled as an unobserved process with an idiosyncratic error. Typically, we assume that conditional log-volatility follows an autoregression of order one, AR(1), which is a discrete time approach to the diffusion process used in option pricing, see Hull and White (1987). Generally, SV has proven to be more attractive than GARCH-type models introduced in Bollerslev (1986), where conditional volatility is a function of past squared innovations and conditional variances. Jacquiera et al. (1994) find that compared to GARCH, SV yields a more robust description of the autocorrelation pattern of the squared returns. Kim et al. (1998) show that the SV model provides a better in-sample fit than GARCH.

Typically, when we estimate SV models, we often find that the filtered conditional volatility process is very persistent as indicated by the estimated AR(1) coefficient close to one. We can argue that the nearly unit root behavior of the conditional volatility process is due to the failure of accounting for time-variation in the parameters that govern it. We can overcome this issue by incorporating a dynamic discrete latent state Markov process in the model such that volatility parameters can switch from one state to another in either a recurrent or non-recurrent fashion. Moreover, besides allowing for time-variation in the model parameters, we are also interested in enriching the plain SV equation to allow for: Leverage, volatility feedback effects and assuming a heavy-tailed distribution for the innovations, which are important features of financial time-series, see for example, Koopman and Uspensky (2002) and Chan and Grant (2016a). However, contrary to GARCH-type models, considering these features is more difficult in the SV framework due to several reasons. First, although conditional log-volatility follows an AR(1) process, it enters the model in a nonlinear fashion. Second, it cannot be observed resulting in an intractable likelihood function. Third, from a practical point of view, compared to the plain SV equation, incorporating leverage or volatility feedback effects complicates model estimation even further, see Chan and Grant (2016a). Finally, estimating time-varying parameter SV models is even harder because they involve two hidden processes, namely, the conditional volatility and the discrete Markov process.

In order to address these points, we provide a unified estimation framework by modifying the particle marginal Metropolis-Hastings (PMMH) sampler, see Andrieu et al. (2010) and Flury and Shephard (2011). In this context, we rely on using two important results from the aforementioned papers, namely, (i): For nonlinear state-space models that cannot be estimated using the Kalman filter, we can decompose the likelihood of the data, $y_{1:T}$, conditional on the model parameters, θ,

E-mail address: nimanonejad@gmail.com.

http://dx.doi.org/10.1016/j.econmod.2016.11.003
Received 13 June 2016; Received in revised form 31 October 2016; Accepted 1 November 2016
Available online xxxx
0264-9993/

$p(y_{1:T}|\theta)$, through the predictive composition and use simulation to unbiasedly estimate the likelihood contribution at time t, $p(y_t|\theta, y_{1:t-1})$, $t = 1,\ldots,T$, see Flury and Shephard (2011). Simulation from $p(y_t|\theta, y_{1:t-1})$ can be carried out using a particle filter (PF). (ii): When an unbiasedly estimated likelihood (for example from a PF) is used inside a Metropolis-Hastings (MH) algorithm then the estimation error makes no difference to the equilibrium distribution of the algorithm, i.e. the posterior distribution, $p(\theta|y_{1:T})$, allowing for exact inference even when $p(y_{1:T}|\theta)$ is simulated. We use (i) to generate the unobserved (discrete) states, $s_{1:T}$, from its conditional posterior. These states are modeled using either m-state Markov switching or change-point processes. (ii) is used to generate the model parameters within each regime, θ_k, $k = 1,\ldots,m$, through a MH step, replacing the exact likelihood with its unbiased simulation based estimator from a PF. At the same time, the estimation procedure automatically provides us with the observed-data likelihood, which we use to compute the deviance information criterion (DIC) of Spiegelhalter et al. (2002). DIC is then used to choose the number of breaks and perform model comparison within each specification.

We also demonstrate that our approach simplifies the estimation procedure greatly. First, we do not need to condition on the latent volatility process to generate the model parameters. Second, our approach requires limited design effort from the practitioner's part, especially if one desires to change some features in a particular model. For instance, allowing for heavy tails, leverage and volatility feedback effects requires minor modifications in the sampling algorithm. In each case, we slightly modify the particle filter and augment θ to include additional parameters without needing to make substantial changes in the codes. Moreover, we can directly compare model fit using the output from the estimation procedure through DIC. On the other hand, performing the same tasks following Vo (2009) and Chan and Grant (2016a) is more cumbersome. We must also mention that there are other perhaps more computationally sophisticated alternatives to our proposed framework. For instance, Kim (2016) proposes a particle Gibbs algorithm that generates $s_{1:T}$ and the latent volatilities in one-step through ancestor sampling, see Lindsten et al. (2014). Conditional on these processes, model parameters are generated using standard Gibbs sampling techniques. However, the approach of Kim (2016) requires more knowledge about particle filtering (ancestor sampling and partially deterministic sequential Monte Carlo) and computational effort. Our approach, while computationally valid and efficient is based on very simple ideas, namely, (i)–(ii). Furthermore, it requires only basic knowledge of particle filtering, Metropolis-Hasting and compared to papers such as Kim (2016) requires minor coding effort from the practitioner's part. Thus, we can efficiently estimate SV models without vast know how of simulation techniques.

One branch of applied econometrics where these techniques are very useful is with regards to analyzing time-variation in the parameters that govern the conditional volatility of crude oil prices. There are obviously several reason for this. For instance, for oil dependent nations, where crude oil volatility plays an important role in economic policy, unexpected volatility parameter shifts can imply huge losses (gains) and thus lower revenues (higher revenues) with drastic negative (positive) consequences on the economy. Being able to accurately model crude oil volatility is also crucial for decision making and risk management purposes[1]. Neglecting changes in the parameters that govern conditional volatility dynamics can result in biased estimates and thus poor forecasts, damaging investor's portfolio or exposing the investor to unnecessary higher risk. Indeed, a large amount of research devotes alot of attention to modeling the volatility of crude oil prices, see for instance, Bina and Vo (2007), Narayan and Narayan (2007), Fan et al. (2008), Hung et al. (2008), Agnolucci (2009), Kang et al. (2009), Oberndorfer (2009), and Wei et al. (2010). Most of these papers take a GARCH-based approach as compared to SV they are simpler and less computationally demanding to estimate[2]. These studies find that leverage, persistence and heavy-tailed distributed errors are important features of crude oil volatility. Research also concludes that allowing for time-variation in the parameters that govern conditional volatility dynamics also plays an important role in analyzing crude oil volatility, see Fong and See (2002), Fong and See (2003), Alizadeh et al. (2008), Aloui and Jammazi (2009), Vo (2009), Nomikos and Pouliasis (2011) and Arouri et al. (2012)[3].

We provide both a methodological and an economic contribution. From a methodological viewpoint, we contribute to the applied literature by providing a flexible framework that accounts for changing dynamics in the model parameters for a variety of SV-type processes. To our knowledge, no other attempts have been made to model time-varying SV models using our approach. From an economic view point, we investigate the relevance of time-variation in the model parameters in modeling and forecasting the conditional volatility of crude oil prices. This part is broadly related to Vo (2009). However, our approach is more general. First, we condition on both recurrent and non-recurrent regimes and determine, which technique provides the best fit. Second, we allow our SV models to accommodate more complex features. Our analyses are carried out in three phases. First, we present our algorithm. Second, we illustrate its properties and DIC computation on simulated time-series. The final phase provides posterior parameter estimates and forecast results.

The remaining of this paper is as follows: Stochastic volatility models are introduced in Section 2. Parameter instability is discussed in Section 3. Bayesian estimation is detailed in Section 4. Sections 5 and 6 present simulation and empirical results. The last section concludes. Appendices at the end of the paper provide details on a simple particle filter and a prior sensitivity analysis.

2. Stochastic volatility

We start with a brief discussion on the plain stochastic volatility model, which serves as the building block for other, more complex models within this family. We then move on to present several extensions. In Section 3, we incorporate time-variation in the model parameters that govern conditional volatility dynamics. The plain stochastic volatility (SV) model is given as

$$y_t = \mu + \sqrt{\gamma}\exp(h_t/2)\varepsilon_t, \quad \varepsilon_t \sim N(0, 1) \tag{2.1}$$

$$h_{t+1} = \phi h_t + \sigma\eta_t, \quad \eta_t \sim N(0, 1). \tag{2.2}$$

Here, y_t is the observed return at time t and $h_{1:T} = (h_1,\ldots,h_T)'$ are the unobserved log-volatilities, which follow (2.2). The parameter, $\gamma \geq 0$, is the average level of volatility. The remaining parameters in (2.1)–(2.2) are ϕ and σ, which denote the persistence and the conditional volatility of volatility, respectively. We follow Kim et al. (1998) and impose that $|\phi| < 1$, with the initial condition, $h_1 \sim N(0, \sigma^2/(1-\phi^2))$. We also

[1] It is well-established that oil price fluctuations influence aggregate economic activity and stock markets, see Hamilton (1983), Jones and Kaul (1996), Ciner (2001), Cologni and Manera (2008), Lardic and Mignon (2008), and Narayan et al. (2014). At the sector and individual firm levels, the reaction of returns to oil price changes is found to be heterogeneous, see Boyer and Filion (2007), Nandha and Faff (2008), Arouri and Nguyen (2010), Narayan and Sharma (2011) and Phan et al. (2015). We also refer the reader to Narayan et al. (2013) where it is shown how investors can use information on the futures market's ability to predict spot prices to devise trading strategies and obtain profits.

[2] Sadorsky (2005), Trolle and Schwartz (2009), Vo (2009), Larsson and Nossman (2011), Brooks and Prokopczuk (2013) and Chan and Grant (2016a) are the relatively few papers that study crude oil volatility using SV models.

[3] These papers combine GARCH and SV model with Markov-switching dynamics. On the other hand, Ewing and Malik (2010), Kang et al. (2011), Arouri et al. (2012) and Salisu and Fasanya (2013) use the same framework as Wilson et al. (1996), i.e. they first identify the break dates using techniques such as the ICSS algorithm of Inclán and Tiao (1994) or tests introduced in Narayan and Popp (2010) and Narayan and Liu (2015). They then incorporate them in the model by specifying dummy variables in the conditional mean/volatility equations.