Available online at www.sciencedirect.com
ScienceDirect
Computational Statistics & Data Analysis 52 (2008) 2863 – 2876
COMPUTATIONAL STATISTICS & DATA ANALYSIS
www.elsevier.com/locate/csda

Analysis of filtering and smoothing algorithms for Lévy-driven stochastic volatility models

Drew D. Creal*

Department of Econometrics, Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands

Available online 26 July 2007

Abstract

Filtering and smoothing algorithms that estimate the integrated variance in Lévy-driven stochastic volatility models are analyzed. Particle filters are algorithms designed for nonlinear, non-Gaussian models while the Kalman filter remains the best linear predictor if the model is linear but non-Gaussian. Monte Carlo experiments are performed to compare these algorithms across different specifications of the model including different marginal distributions and degrees of persistence for the instantaneous variance. The use of realized variance as an observed variable in the state space model is also evaluated. Finally, the particle filter's ability to identify the timing and size of jumps is assessed relative to popular nonparametric estimators.

Keywords: Lévy process; Stochastic volatility; Particle filter; Kalman filter

1. Introduction

This paper contributes to the literature on parametric modeling of volatility by comparing various filtering and smoothing algorithms for the non-Gaussian Ornstein–Uhlenbeck (OU) stochastic volatility models introduced by Barndorff-Nielsen and Shephard (2001a) and extended in Barndorff-Nielsen and Shephard (2003). The instantaneous variance in these models is driven by a non-Gaussian Lévy process, which requires approximate recursions for filtering and smoothing. I examine how well the particle filter compares to the Kalman filter at estimating the integrated variance under a number of different specifications of the model. I also consider estimation of jumps by the particle filter and smoother. The goal of the article is to build on some of the comments in the original paper. The filtering algorithms considered here are applicable regardless of the method used to estimate the model's parameters. Parameters for the BNS-SV model have been estimated using quasi-maximum likelihood in Barndorff-Nielsen and Shephard (2002) and Barndorff-Nielsen et al. (2004); the smooth particle filter of Pitt (2001); Markov chain Monte Carlo (MCMC) in Roberts et al. (2004), Gander and Stephens (2005), and Griffin and Steel (2006); minimum distance estimation in Todorov (2006); and sequential Monte Carlo samplers in Del Moral et al. (2006). MCMC can also be used to compute smoothed estimates, but that is beyond the scope of this article.

When there are no jumps, both filtering algorithms estimate increments of the integrated variance almost exactly, even at moderate intra-daily frequencies. When returns are sampled at low frequencies, the Kalman filter's performance is more robust than the particle filter. These differences can be large in practice depending on the marginal distribution of

* Tel.: +31 20 598 60 10; fax: +31 20 598 60 20.
E-mail address: dcreal@feweb.vu.nl.

0167-9473/$ - see front matter
doi:10.1016/j.csda.2007.07.009

the instantaneous variance. If returns are available at higher frequencies, the particle filter offers a moderate improvement over the Kalman filter. The size of the improvement also depends on the marginal distribution of the instantaneous variance. The particle filter struggles when the variance of the stationary distribution is larger. Small misspecifications of the marginal distribution can potentially lead to poorer estimates from the particle filter. These differences can be large enough to make the Kalman filter and smoother competitive at estimating the integrated variance at any frequency at which returns are sampled.

Finally, I also compare the particle filter and smoother's ability to estimate jumps versus the jump statistics developed in Barndorff-Nielsen and Shephard (2004). Results indicate that the particle filter identifies the jump times and sizes well.

2. Lévy-driven stochastic volatility models

2.1. Price processes with stochastic volatility and jumps

I assume that the log-price process $y^*(t)$ takes the form

$$\mathrm{d}y^*(t) = (\mu + \beta\sigma^2(t))\,\mathrm{d}t + \sigma(t)\,\mathrm{d}B(t) + \rho\,\mathrm{d}\bar{z}(\lambda t), \tag{1}$$

where μ is the drift and β is the risk premium. The instantaneous variance $\sigma^2(t)$ is assumed to be independent of the standard Brownian motion $B(t)$. Following Barndorff-Nielsen and Shephard (2001a), the instantaneous variance is specified as a non-Gaussian OU process

$$\mathrm{d}\sigma^2(t) = -\lambda\sigma^2(t)\,\mathrm{d}t + \mathrm{d}z(\lambda t), \tag{2}$$

where $z(t)$ is a Lévy process called the background driving Lévy process. The process $z(t)$ is also known as a subordinator. The timing of the subordinator ($\mathrm{d}z(\lambda t)$ versus $\mathrm{d}z(t)$) is intentionally designed to separate the marginal distribution of the instantaneous variance from its autocovariance function. Lévy processes are a general group of continuous time stochastic processes that have independent and stationary increments, see e.g. Cont and Tankov (2004). The compensated version of the subordinator, $\bar{z}(t) = z(t) - E\{z(t)\}$, in (1) accounts for the leverage effect with ρ determining the direction of the jumps in the price.

Regardless of the model for the instantaneous variance, the integrated variance is defined as

$$\sigma^{2*}(t) = \int_0^t \sigma^2(u)\,\mathrm{d}u.$$

Defining $\hbar$ to be the length of time between two periods of interest (e.g. 1 day), increments of the integrated variance

$$\sigma_n^2 = \sigma^{2*}(\hbar n) - \sigma^{2*}(\hbar(n-1))$$

are known as the actual variance and they are a critical ingredient in option pricing and risk management. For this model, the actual variance has a simple structure:

$$\sigma_n^2 = \lambda^{-1}\{z(\lambda\hbar n) - \sigma^2(\hbar n) - z(\lambda\hbar(n-1)) + \sigma^2(\hbar(n-1))\}. \tag{3}$$

It can be calculated by simulating the continuous time process

$$\begin{bmatrix}\sigma^2(\hbar n)\\ z(\lambda\hbar n)\end{bmatrix} = \begin{bmatrix}\exp(-\lambda\hbar)\sigma^2(\hbar(n-1))\\ z(\lambda\hbar(n-1))\end{bmatrix} + \eta_n, \tag{4}$$

where

$$\eta_n = \begin{bmatrix}\exp(-\lambda\hbar)\int_0^{\hbar}\exp(\lambda s)\,\mathrm{d}z(\lambda s)\\ \int_0^{\hbar}\mathrm{d}z(\lambda s)\end{bmatrix}. \tag{5}$$

There are numerous methods for simulating Lévy processes, see e.g. Rosínski (2001b) and Cont and Tankov (2004). In this paper, I use the series representation from Barndorff-Nielsen and Shephard (2001a) to simulate $\sigma^2(t)$ when it