



Moving average stochastic volatility models with application to inflation forecast



Joshua C.C. Chan*

Research School of Economics, Australian National University, Australia

ARTICLE INFO

Article history:

Received 27 March 2012

Received in revised form

15 October 2012

Accepted 23 May 2013

Available online 30 May 2013

JEL classification:

C11

C51

C53

Keywords:

State space

Unobserved components model

Precision

Sparse

Density forecast

ABSTRACT

We introduce a new class of models that has both stochastic volatility and moving average errors, where the conditional mean has a state space representation. Having a moving average component, however, means that the errors in the measurement equation are no longer serially independent, and estimation becomes more difficult. We develop a posterior simulator that builds upon recent advances in precision-based algorithms for estimating these new models. In an empirical application involving US inflation we find that these moving average stochastic volatility models provide better in-sample fitness and out-of-sample forecast performance than the standard variants with only stochastic volatility.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Since the pioneering works of Box and Jenkins, autoregressive moving average (ARMA) models have become standard tools for modeling and forecasting time series. The theoretical justification of these ARMA models, as is well-known, is the Wold decomposition theorem, which states that any zero mean covariance-stationary time series has an infinite moving average representation. One implication is that any such process can be approximated arbitrarily well by a sufficiently high order ARMA model. In practice, it is found that simple univariate ARMA models often outperform complex multivariate models in forecasting.

However, despite the theoretical justification and empirical success of this class of models, a voluminous literature has highlighted the importance of allowing for time-varying volatility in macroeconomic and financial data for estimation and forecasting. Standard ARMA models that assume constant variance are seemingly not flexible enough. One way to accommodate this time-variation in variance is via the GARCH model (Bollerslev,

1986). For example, Nakatsuma (2000) considers a linear regression model with ARMA-GARCH errors. Another popular way to allow for time-varying volatility is via the stochastic volatility (SV) model (e.g., Taylor, 1994; Kim et al., 1998). The popularity of this approach can be seen through the numerous extensions of the basic SV setup in recent years, such as the SV models with jump and Student's t error (Chib et al., 2002), SV with leverage (Jacquier et al., 2004; Omori et al., 2007), SV with asymmetric, heavy-tailed error (Nakajima and Omori, 2012), semiparametric SV models via the Dirichlet process mixture (Jensen and Maheu, 2010), etc., to name but a few examples.

Several recent studies have attempted to bridge these two literatures on ARMA and SV models, and have considered various flexible autoregressive models with stochastic volatility (e.g., Cogley and Sargent, 2005; Primiceri, 2005; Cogley et al., 2010). But there are few papers that investigate moving average models with SV. The purpose of this article is to fill this gap: we introduce a class of models that includes both the moving average and stochastic volatility components, where the conditional mean process has a flexible state space representation. As such, the setup includes a wide variety of popular specifications as special cases, including the unobserved components and time-varying parameter models. Of course, any invertible MA process can be approximated by a sufficiently high order AR model. In practice, however, forecasts based on these AR models—since they have many parameters to

* Correspondence to: Research School of Economics, ANU College of Business and Economics, LF Crisp Building 26, The Australian National University, Canberra ACT 0200, Australia. Tel.: +61 2 612 57358; fax: +61 2 612 50182.

E-mail addresses: joshuacc.chan@gmail.com, joshua.chan@anu.edu.au.

estimate—often compare poorly to parsimonious ARMA models (e.g., Stock and Watson, 2007; Athanasopoulos and Vahid, 2008). In our empirical work that involves quarterly inflation, we find that there is substantial support for the proposed models against their counterparts with only SV. The forecasting results suggest that addition of the MA component further improves the forecast performance of standard SV models, particularly at short forecast horizons.

A second contribution of this paper is to develop an efficient Markov chain Monte Carlo (MCMC) sampler for estimating this class of models. Since the conditional mean process has a state space form, estimation might appear to be straightforward. However, under our models the errors in the measurement equation are no longer serially independent due to the presence of the MA component. As such, application of Kalman filter-based methods would first require a suitable transformation of the data to make the errors serially independent. Instead of using the Kalman filter, we take a different approach: we extend the previous work on precision-based algorithms for state space models in Chan and Jeliazkov (2009) and McCausland et al. (2011), which are shown to be more efficient than Kalman filter-based methods. The idea of exploiting banded precision matrices can be traced back to Fahrmeir and Kaufmann (1991); see also Rue et al. (2009) and Ruiz-Cardenas et al. (2012). By exploiting the sparse structure of the covariance matrix of the observations, we develop an easy and fast method for estimating these new models.

A third contribution involves a substantive empirical application on modeling and forecasting US quarterly consumer price index (CPI) inflation. A vast literature on this topic has emerged over the last two decades; recent studies include Koop and Potter (2007), Stock and Watson (2007, 2010), Cogley and Sbordone (2008), Cogley et al. (2010), Clark and Doh (2011), Korobilis (2013), Koop and Korobilis (2012), among many others. One key finding in this literature is that both persistence and volatility in the inflation process have changed considerably over time. In particular, inflation volatility decreased gradually from the great inflation of the 1970s and throughout the great moderation, until it increased again and peaked at the aftermath of the global financial crisis. Empirically, it is often found that models with stochastic volatility provide substantially better point and density forecasts than those obtained from constant error variance models (e.g., Clark and Doh, 2011; Chan et al., 2012).

Another key finding in this literature is that for forecasting inflation, both at short and long horizons, it is often difficult to improve upon univariate models using only information in observed inflation (e.g., Stock and Watson, 2007, 2010; Chan et al., 2012). One reason for this lack of predictive power of a wide range of seemingly relevant variables—such as unemployment rate and GDP growth—might be because variables useful for forecasting change over time (e.g., oil price might be an important predictor for inflation in the 1970s but is less important in the 2000s) and/or over business cycle (e.g., some variables may predict well in expansions but not in recessions). In fact, Koop and Korobilis (2012) find evidence that the set of relevant predictors for inflation does change over time. Given these findings, we consider univariate time series models using only information in observed inflation. Additional explanatory variables, of course, can be incorporated if desired.

We focus on univariate MA-SV models in this paper, partly because for our empirical work these models are sufficient. We note that, the univariate framework developed here can be used to construct multivariate models in a straightforward manner. For example, in the multivariate SV models of Chib et al. (2006), SV is induced by a number of latent factors, each of which follows an independent univariate SV process. In this setup, we can, for example, replace the SV process with the univariate MA-SV process

introduced in this paper to construct a multivariate SV model with autocorrelated errors. We leave the multivariate case for future research.

The rest of this article is organized as follows. In Section 2 we introduce the general framework, and discuss how this state space form includes a variety of popular specifications as special cases. Section 3 develops an efficient posterior simulator to estimate this new class of models. Section 4 presents the empirical results for modeling and forecasting US CPI inflation. In the last section we conclude our findings and discuss the future research direction.

2. Moving average stochastic volatility models

The general framework we consider is the following q -th-order moving average model with stochastic volatility:

$$y_t = \mu_t + \varepsilon_t^y, \tag{1}$$

$$\varepsilon_t^y = u_t + \psi_1 u_{t-1} + \dots + \psi_q u_{t-q}, \quad u_t \sim \mathcal{N}(0, e^{h_t}), \tag{2}$$

$$h_t = \mu_h + \phi_h(h_{t-1} - \mu_h) + \varepsilon_t^h, \quad \varepsilon_t^h \sim \mathcal{N}(0, \sigma_h^2), \tag{3}$$

where we assume that $|\phi_h| < 1$. The errors u_t and ε_t^h are independent of each other for all leads and lags. We further assume that $u_0 = u_{-1} = \dots = u_{-q+1} = 0$. One can, of course, treat these initial error terms as parameters if desired, and the estimation procedures discussed in the next section can be easily extended to allow for this possibility. For typical situations where $T \gg q$, whether these errors are modeled explicitly or not makes little difference in practice.

Let $\boldsymbol{\mu} = (\mu_1, \dots, \mu_T)'$, $\mathbf{h} = (h_1, \dots, h_T)'$ and $\boldsymbol{\psi} = (\psi_1, \dots, \psi_q)'$. Then, it is easy to see that the conditional variance of y_t is given by

$$\text{Var}(y_t | \boldsymbol{\mu}, \boldsymbol{\psi}, \mathbf{h}) = e^{h_t} + \psi_1^2 e^{h_{t-1}} + \dots + \psi_q^2 e^{h_{t-q}}.$$

In other words, the conditional variance of y_t is time-varying through two channels: it is a moving average of the $q + 1$ most recent variances $e^{h_t}, \dots, e^{h_{t-q}}$, and the log-volatility h_t in turn evolves according to the stationary AR(1) process in (3). Unlike the standard SV models, y_t is serially correlated even after conditioning on the states. In fact, its conditional autocovariances are given by

$$\text{Cov}(y_t, y_{t-j} | \boldsymbol{\mu}, \boldsymbol{\psi}, \mathbf{h}) = \begin{cases} \sum_{i=0}^{q-j} \psi_{i+j} \psi_i e^{h_{t-i}}, & \text{for } j = 1, \dots, q, \\ 0, & \text{for } j > q, \end{cases}$$

where $\psi_0 = 1$. It is interesting to note that due to the presence of the log-volatility h_t , the autocovariances of y_t are also time-varying.¹

Now, by choosing a suitable conditional mean process μ_t , the model in (1)–(3) includes a variety of popular specifications, such as:

1. the autoregressive model:

$$\mu_t = \phi_0 + \phi_1 y_{t-1} + \dots + \phi_p y_{t-p};$$

2. the linear regression model:

$$\mu_t = \beta_0 + \beta_1 x_{1t} + \dots + \beta_k x_{kt},$$

where $\mathbf{x}_t = (x_{1t}, \dots, x_{kt})$ is a vector of explanatory variables;

3. the unobserved components model:

$$\begin{aligned} \mu_t &= \tau_t, \\ \tau_t &= \tau_{t-1} + \varepsilon_t^\tau, \quad \varepsilon_t^\tau \sim \mathcal{N}(0, \sigma_\tau^2); \end{aligned}$$

¹ On the other hand, the marginal variance and autocovariances of y_t unconditional on \mathbf{h} do not seem to have closed-form expressions.

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات