

CAMA

Centre for Applied Macroeconomic Analysis

Moving Average Stochastic Volatility Models with Application to Inflation Forecast

CAMA Working Paper 31/2013
May 2013

Joshua C.C. Chan,

Centre for Applied Macroeconomic Analysis (CAMA), Australian National University

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 U.S. 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.

Keywords

state space, unobserved components model, precision, sparse, density forecast

JEL Classification

C11, C51, C53

Address for correspondence:

(E) cama.admin@anu.edu.au

[The Centre for Applied Macroeconomic Analysis](#) in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

Moving Average Stochastic Volatility Models with Application to Inflation Forecast

Joshua C.C. Chan*
Centre for Applied Macroeconomic Analysis,
Australian National University

May 2013

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 U.S. 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.

Keywords: state space, unobserved components model, precision, sparse, density forecast

JEL classification codes: C11, C51, C53

*We thank seminar participants at Monash University, the 6-th Rimini Bayesian Econometrics Workshop, the 18-th Australasian Macroeconomics Workshop, and Workshop on Empirical Methods in Macroeconomic Policy Analysis for helpful comments and suggestions. All errors are, of course, our own. Postal address: Research School of Economics, ANU College of Business and Economics, LF Crisp Building 26, The Australian National University, Canberra ACT 0200, Australia. Email: joshua.chan@anu.edu.au. Phone: (+61) 2 612 57358. Fax: (+61) 2 612 50182.

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, Shephard, and Chib, 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, Nardari, and Shephard, 2002), SV with leverage (Jacquier, Polson, and Rossi, 2004; Omori, Chib, Shephard, and Nakajima, 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, Primiceri, and Sargent, 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 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 previous work on precision-based algorithms for state space models in Chan and Jeliazkov (2009) and McCausland, Miller, and Pelletier (2011), which are shown to be more efficient than Kalman filter-based methods. The idea of exploiting banded precision matrix can be traced back to Fahrmeir and Kaufmann (1991); see also Rue, Martino, and Chopin (2009) and Ruiz-Cardenas, Krainski, and Rue (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 U.S. 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 (2012), 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, Koop, Leon-Gonzalez, and Strachan, 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, Nardari, and Shephard (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 empirical results for modeling and forecasting U.S. CPI inflation. In the last section we conclude our findings and discuss 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:

$$y_t = \mu_t + \varepsilon_t^y, \quad (1)$$

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

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

where we assume $|\phi_h| < 1$. The errors u_t and ε_t^h are independent of each other for all leads and lags. We further assume that $u_0 = u_{-1} = \cdots = u_{-q+1} = 0$. 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

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

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

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

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:

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

2. the linear regression model:

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

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

3. the unobserved components model:

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

4. the time-varying parameter model:

$$\begin{aligned} \mu_t &= \beta_{0t} + \beta_{1t} x_{1t} + \cdots + \beta_{kt} x_{kt}, \\ \boldsymbol{\beta}_t &= \boldsymbol{\beta}_{t-1} + \boldsymbol{\varepsilon}_t^\beta, \quad \boldsymbol{\varepsilon}_t^\beta \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\beta), \end{aligned}$$

where $\boldsymbol{\beta}_t = (\beta_{0t}, \beta_{1t}, \dots, \beta_{kt})'$.

Some other flexible time-varying models recently introduced also fall within this general framework. Examples include an autoregressive unobserved components model discussed in Clark and Doh (2011), as well as various bounded trend inflation models proposed in Chan, Koop, and Potter (2013). The framework in (1)–(3) is a natural extension of the standard stochastic volatility setting. In particular, the latter is a special case of this general framework with $\psi_1 = \cdots = \psi_q = 0$. For identification purposes, we assume the usual invertibility conditions, i.e., the roots of the characteristic polynomial associated with the MA coefficients are all outside the unit circle.

It is well known that moving average models have a state space representation, and the likelihood function can be evaluated using the Kalman filter (Harvey, 1985). This approach can be slow, however, especially when we need to make tens of thousands functional evaluations of the likelihood in the MCMC algorithm. We therefore introduce in the next section a direct way to evaluate the likelihood function. By utilizing fast sparse matrix routines, the new approach is simple and easy to program. Another complication in our setting is that the MA component induces serial dependence in observations. Consequently, in order to apply conventional Kalman filter-based methods to simulate the states, one would first need to transform the data so that the errors in the new measurement equation are serially independent (e.g., one such transformation is suggested in Chib and Greenberg, 1994). Instead of using Kalman filter, we introduce a direct approach that builds upon previous work on precision-based algorithms in Chan and Jeliazkov (2009) for fitting this new class of models. McCausland et al. (2011) provide a careful comparison between Kalman-filter based and precision-based algorithms, and show that the latter algorithms are substantially more efficient.

3 Estimation

We introduce a direct approach for estimating the class of models in (1)–(3) that exploits the special structure of the problem, particularly that the covariance matrix of the joint distribution for $\mathbf{y} = (y_1, \dots, y_T)'$ is sparse, i.e., it contains only a few non-zero elements. We first introduce a fast and simple way to evaluate the likelihood function that is useful for both maximum likelihood and Bayesian estimation. It is followed by a detailed discussion on a new posterior simulator for estimating the log-volatilities and the parameters in the conditional mean process.

3.1 Likelihood Evaluation

To obtain the likelihood function, we first derive the joint distribution of the observations $\mathbf{y} = (y_1, \dots, y_T)'$. To this end, we rewrite (1)–(2) in matrix notations:

$$\mathbf{y} = \boldsymbol{\mu} + \mathbf{H}_\psi \mathbf{u}, \quad (4)$$

where $\boldsymbol{\mu} = (\mu_1, \dots, \mu_T)'$, $\mathbf{u} = (u_1, \dots, u_T)' \sim \mathcal{N}(\mathbf{0}, \mathbf{S}_y)$, $\mathbf{S}_y = \text{diag}(e^{h_1}, \dots, e^{h_T})$, and \mathbf{H}_ψ is a $T \times T$ lower triangular matrix with ones on the main diagonal, ψ_1 on first lower diagonal, ψ_2 on second lower diagonal, and so forth. For example, for $q = 2$, we have

$$\mathbf{H}_\psi = \begin{pmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ \psi_1 & 1 & 0 & 0 & \cdots & 0 \\ \psi_2 & \psi_1 & 1 & 0 & \cdots & 0 \\ 0 & \psi_2 & \psi_1 & 1 & \cdots & 0 \\ \vdots & & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & \psi_2 & \psi_1 & 1 \end{pmatrix}.$$

It is important to note that in general \mathbf{H}_ψ is a banded $T \times T$ matrix that contains only $(T - q/2)(q + 1) < T(q + 1)$ non-zero elements, which is substantially less than T^2 for typical applications where $T \gg q$. This special structure can be exploited to speed up computation. For instance, obtaining the Cholesky decomposition of a banded $T \times T$ matrix with fixed bandwidth involves only $\mathcal{O}(T)$ operations (e.g., Golub and van Loan, 1983, p.156) as opposed to $\mathcal{O}(T^3)$ for a full matrix of the same size. Similar computational savings can be generated in operations such as multiplication, forward and backward substitution by using block-banded or sparse matrix algorithms. These banded and sparse matrix algorithms are implemented in standard packages such as MATLAB, GAUSS and R.

Now, by a simple change of variable, it follows from (4) that

$$(\mathbf{y} \mid \boldsymbol{\psi}, \boldsymbol{\mu}, \mathbf{h}) \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Omega}_y),$$

where $\mathbf{h} = (h_1, \dots, h_T)'$, $\boldsymbol{\Omega}_y = \mathbf{H}_\psi \mathbf{S}_y \mathbf{H}_\psi'$. Since $\mathbf{S}_y = \text{diag}(e^{h_1}, \dots, e^{h_T})$ is a diagonal matrix, and \mathbf{H}_ψ is a lower triangular sparse matrix, the product $\boldsymbol{\Omega}_y$ is sparse. In fact,

it is a banded matrix with only a narrow band of non-zero elements around the main diagonal. Moreover, since $|\mathbf{H}_\psi| = 1$ for any $\boldsymbol{\psi} = (\psi_1, \dots, \psi_q)'$, we have $|\boldsymbol{\Omega}_y| = |\mathbf{S}_y| = \exp\left(\sum_{t=1}^T h_t\right)$. The log joint density of \mathbf{y} is therefore given by

$$\log p(\mathbf{y} | \boldsymbol{\psi}, \boldsymbol{\mu}, \mathbf{h}) = -\frac{T}{2} \log(2\pi) - \frac{1}{2} \sum_{t=1}^T h_t - \frac{1}{2} (\mathbf{y} - \boldsymbol{\mu})' \boldsymbol{\Omega}_y^{-1} (\mathbf{y} - \boldsymbol{\mu}). \quad (5)$$

It is important to realize that one need not obtain the $T \times T$ inverse matrix $\boldsymbol{\Omega}_y^{-1}$ in order to evaluate the log density in (5)—it would involve $\mathcal{O}(T^3)$ operations. Instead, it can be computed in three steps, each of which requires only $\mathcal{O}(T)$ operations. To this end, we introduce the following notations: given a lower (upper) triangular $T \times T$ non-singular matrix \mathbf{A} and a $T \times 1$ vector \mathbf{c} , let $\mathbf{A} \setminus \mathbf{c}$ denote the unique solution to the triangular system $\mathbf{A}\mathbf{x} = \mathbf{c}$ obtained by forward (backward) substitution, i.e., $\mathbf{A} \setminus \mathbf{c} = \mathbf{A}^{-1}\mathbf{c}$. Now, we first obtain the Cholesky decomposition \mathbf{C}_y of the banded matrix $\boldsymbol{\Omega}_y$ such that $\mathbf{C}_y \mathbf{C}_y' = \boldsymbol{\Omega}_y$, which involves only $\mathcal{O}(T)$ operations. Then compute

$$\mathbf{x}_1 = \mathbf{C}_y' \setminus (\mathbf{C}_y \setminus (\mathbf{y} - \boldsymbol{\mu}))$$

by forward followed by backward substitution, each of which requires $\mathcal{O}(T)$ operations since \mathbf{C}_y is also banded. By definition,

$$\mathbf{x}_1 = \mathbf{C}_y^{-1'} (\mathbf{C}_y^{-1} (\mathbf{y} - \boldsymbol{\mu})) = (\mathbf{C}_y \mathbf{C}_y')^{-1} (\mathbf{y} - \boldsymbol{\mu}) = \boldsymbol{\Omega}_y^{-1} (\mathbf{y} - \boldsymbol{\mu}).$$

Finally, compute

$$\mathbf{x}_2 = -\frac{1}{2} (\mathbf{y} - \boldsymbol{\mu})' \mathbf{x}_1 = -\frac{1}{2} (\mathbf{y} - \boldsymbol{\mu})' \boldsymbol{\Omega}_y^{-1} (\mathbf{y} - \boldsymbol{\mu}),$$

which gives the quadratic term in (5). Thus, given $\boldsymbol{\mu}$, $\boldsymbol{\psi}$ and \mathbf{h} , one can efficiently evaluate the likelihood function without the need of the Kalman filter.

3.2 Posterior Analysis

Now, we discuss an efficient posterior sampler for estimating the MA-SV model in (1)–(3). To keep the discussion concrete, we consider in particular the unobserved components specification; estimation for other conditional mean processes follows similarly. Specifically, the measurement equation is given by (4) with $\boldsymbol{\mu} = \boldsymbol{\tau} = (\tau_1, \dots, \tau_T)'$, whereas the transition equations are

$$\begin{aligned} \tau_t &= \tau_{t-1} + \varepsilon_t^\tau, & \varepsilon_t^\tau &\sim \mathcal{N}(0, \sigma_\tau^2), \\ h_t &= \mu_h + \phi_h (h_{t-1} - \mu_h) + \varepsilon_t^h, & \varepsilon_t^h &\sim \mathcal{N}(0, \sigma_h^2), \end{aligned}$$

with $|\phi_h| < 1$. The transition equation for τ_t is initialized with $\tau_1 \sim \mathcal{N}(\tau_0, \sigma_{0\tau}^2)$, where τ_0 and $\sigma_{0\tau}^2$ are some known constants. In particular, we set $\tau_0 = 0$. Moreover, the transition equation for h_t is initialized with $h_1 \sim \mathcal{N}(\mu_h, \sigma_h^2 / (1 - \phi_h^2))$.

We assume independent priors for $\boldsymbol{\psi}$, σ_τ^2 , μ_h , ϕ_h and σ_h^2 , i.e., $p(\boldsymbol{\psi}, \sigma_\tau^2, \mu_h, \phi_h, \sigma_h^2) = p(\boldsymbol{\psi})p(\sigma_\tau^2)p(\mu_h)p(\phi_h)p(\sigma_h^2)$. For $\boldsymbol{\psi}$, we consider a multivariate normal prior with support in the region where the invertibility conditions on $\boldsymbol{\psi}$ hold. As for other parameters, we assume the following independent priors:

$$\sigma_\tau^2 \sim \mathcal{IG}(\nu_\tau, S_\tau), \quad \mu_h \sim \mathcal{N}(\mu_{h0}, V_{\mu_h}), \quad \phi_h \sim \mathcal{N}(\phi_{h0}, V_{\phi_h})\mathbf{1}(|\phi_h| < 1), \quad \sigma_h^2 \sim \mathcal{IG}(\nu_h, S_h),$$

where \mathcal{IG} denotes the inverse-gamma distribution. Note that we impose the stationarity condition $|\phi_h| < 1$ through the prior on ϕ_h . Then posterior draws can be obtained by sequentially sampling from:²

1. $p(\boldsymbol{\tau} \mid \mathbf{y}, \mathbf{h}, \boldsymbol{\psi}, \sigma_\tau^2)$;
2. $p(\mathbf{h} \mid \mathbf{y}, \boldsymbol{\tau}, \boldsymbol{\psi}, \sigma_h^2, \mu_h, \phi_h)$;
3. $p(\boldsymbol{\psi}, \sigma_h^2, \sigma_\tau^2 \mid \mathbf{y}, \boldsymbol{\tau}, \mathbf{h}, \mu_h, \phi_h) = p(\boldsymbol{\psi} \mid \mathbf{y}, \boldsymbol{\tau}, \mathbf{h}) p(\sigma_h^2 \mid \mathbf{h}, \mu_h, \phi_h) p(\sigma_\tau^2 \mid \boldsymbol{\tau})$;
4. $p(\mu_h \mid \mathbf{h}, \sigma_h^2, \phi_h)$;
5. $p(\phi_h \mid \mathbf{h}, \mu_h, \sigma_h^2)$.

We first discuss an efficient way to sample from $p(\boldsymbol{\tau} \mid \mathbf{y}, \mathbf{h}, \boldsymbol{\psi}, \sigma_\tau^2)$. First note that by pre-multiplying (4) by \mathbf{H}_ψ^{-1} , we have

$$\tilde{\mathbf{y}} = \tilde{\boldsymbol{\tau}} + \mathbf{u},$$

where $\tilde{\mathbf{y}} = \mathbf{H}_\psi^{-1}\mathbf{y}$ and $\tilde{\boldsymbol{\tau}} = \mathbf{H}_\psi^{-1}\boldsymbol{\tau}$. In other words, the log density for $\tilde{\mathbf{y}}$ is

$$\log p(\tilde{\mathbf{y}} \mid \boldsymbol{\psi}, \tilde{\boldsymbol{\tau}}, \mathbf{h}) \propto -\frac{1}{2} \sum_{t=1}^T h_t - \frac{1}{2} (\tilde{\mathbf{y}} - \tilde{\boldsymbol{\tau}})' \mathbf{S}_y^{-1} (\tilde{\mathbf{y}} - \tilde{\boldsymbol{\tau}}), \quad (6)$$

where $\mathbf{S}_y = \text{diag}(e^{h_1}, \dots, e^{h_T})$. Hence it is more convenient to work with $\tilde{\boldsymbol{\tau}}$ instead of the original parameterization $\boldsymbol{\tau}$. Once we have a draw for $\tilde{\boldsymbol{\tau}}$, we simply pre-multiply it by \mathbf{H}_ψ to get a draw for $\boldsymbol{\tau}$.

Next, we derive the prior density for $\tilde{\boldsymbol{\tau}}$. To this end, we first obtain the prior density for $\boldsymbol{\tau}$. Rewrite the transition equation for τ_t in matrix notations:

$$\mathbf{H}\boldsymbol{\tau} = \boldsymbol{\varepsilon}^\tau,$$

where $\boldsymbol{\varepsilon}^\tau = (\varepsilon_1^\tau, \dots, \varepsilon_T^\tau)' \sim \mathcal{N}(\mathbf{0}, \mathbf{S}_\tau)$, $\mathbf{S}_\tau = \text{diag}(\sigma_{0\tau}^2, \sigma_\tau^2, \dots, \sigma_\tau^2)$, and \mathbf{H} is the first difference matrix

$$\mathbf{H} = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ -1 & 1 & 0 & \cdots & 0 \\ 0 & -1 & 1 & \cdots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & 0 & \cdots & -1 & 1 \end{pmatrix}.$$

²MATLAB code for estimating this unobserved components MA-SV model is available at <http://people.anu.edu.au/joshua.chan/>.

That is, $(\boldsymbol{\tau} | \sigma_\tau^2) \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega}_\tau)$, where $\boldsymbol{\Omega}_\tau^{-1} = \mathbf{H}'\mathbf{S}_\tau^{-1}\mathbf{H}$. Recall that $\sigma_{0\tau}^2$ is the variance for the initial underlying trend τ_1 , and is assumed to be a fixed hyper-parameter (although it is straightforward to treat it as a parameter).

It is important to realize that in this case the precision matrix $\boldsymbol{\Omega}_\tau^{-1}$ is sparse. Now, by a simple change of variable, we have $(\tilde{\boldsymbol{\tau}} | \sigma_\tau^2) \sim \mathcal{N}(\mathbf{0}, \mathbf{H}_\psi^{-1}\boldsymbol{\Omega}_\tau\mathbf{H}_\psi^{-1})$. Noting that $|\mathbf{H}| = |\mathbf{H}_\psi| = 1$ and $|\boldsymbol{\Omega}_\tau| = \sigma_{0\tau}^2(\sigma_\tau^2)^{T-1}$, the log prior density for $\tilde{\boldsymbol{\tau}}$ is therefore given by

$$\log p(\tilde{\boldsymbol{\tau}} | \sigma_\tau^2) \propto -\frac{T-1}{2} \log \sigma_\tau^2 - \frac{1}{2} \tilde{\boldsymbol{\tau}}' \mathbf{H}_\psi' \boldsymbol{\Omega}_\tau^{-1} \mathbf{H}_\psi \tilde{\boldsymbol{\tau}}, \quad (7)$$

Combining (6) and (7), and using standard results from linear regression (see, e.g., Koop, 2003), we obtain the log conditional density $\log p(\tilde{\boldsymbol{\tau}} | \tilde{\mathbf{y}}, \mathbf{h}, \boldsymbol{\psi}, \sigma_\tau^2)$ as follows:

$$\begin{aligned} \log p(\tilde{\boldsymbol{\tau}} | \tilde{\mathbf{y}}, \mathbf{h}, \boldsymbol{\psi}, \sigma_\tau^2) &\propto -\frac{1}{2} (\tilde{\mathbf{y}} - \tilde{\boldsymbol{\tau}})' \mathbf{S}_y^{-1} (\tilde{\mathbf{y}} - \tilde{\boldsymbol{\tau}}) - \frac{1}{2} \tilde{\boldsymbol{\tau}}' \mathbf{H}_\psi' \boldsymbol{\Omega}_\tau^{-1} \mathbf{H}_\psi \tilde{\boldsymbol{\tau}} \\ &\propto -\frac{1}{2} (\tilde{\boldsymbol{\tau}}' (\mathbf{S}_y^{-1} + \mathbf{H}_\psi' \boldsymbol{\Omega}_\tau^{-1} \mathbf{H}_\psi) \tilde{\boldsymbol{\tau}} - 2\tilde{\boldsymbol{\tau}}' \mathbf{S}_y^{-1} \tilde{\mathbf{y}}) \\ &\propto -\frac{1}{2} (\tilde{\boldsymbol{\tau}} - \hat{\boldsymbol{\tau}})' \mathbf{D}_{\tilde{\boldsymbol{\tau}}}^{-1} (\tilde{\boldsymbol{\tau}} - \hat{\boldsymbol{\tau}}), \end{aligned}$$

where $\mathbf{D}_{\tilde{\boldsymbol{\tau}}} = (\mathbf{S}_y^{-1} + \mathbf{H}_\psi' \boldsymbol{\Omega}_\tau^{-1} \mathbf{H}_\psi)^{-1}$, and $\hat{\boldsymbol{\tau}} = \mathbf{D}_{\tilde{\boldsymbol{\tau}}} \mathbf{S}_y^{-1} \tilde{\mathbf{y}}$. That is,

$$(\tilde{\boldsymbol{\tau}} | \tilde{\mathbf{y}}, \mathbf{h}, \boldsymbol{\psi}, \sigma_\tau^2) \sim \mathcal{N}(\hat{\boldsymbol{\tau}}, \mathbf{D}_{\tilde{\boldsymbol{\tau}}}).$$

Since $\mathcal{N}(\hat{\boldsymbol{\tau}}, \mathbf{D}_{\tilde{\boldsymbol{\tau}}})$ is typically high dimensional, sampling from it using a brute-force approach is time-consuming. Here we adopt the precision-based sampling method in Chan and Jeliazkov (2009) to obtain draws from $\mathcal{N}(\hat{\boldsymbol{\tau}}, \mathbf{D}_{\tilde{\boldsymbol{\tau}}})$ efficiently. To proceed, first note that the precision matrix $\mathbf{D}_{\tilde{\boldsymbol{\tau}}}^{-1} = (\mathbf{S}_y^{-1} + \mathbf{H}_\psi' \boldsymbol{\Omega}_\tau^{-1} \mathbf{H}_\psi)$ is sparse. Thus $\hat{\boldsymbol{\tau}}$ can be computed quickly using the same approach for evaluating the likelihood function discussed earlier: obtain the Cholesky decomposition $\mathbf{C}_{\tilde{\boldsymbol{\tau}}}$ of $\mathbf{D}_{\tilde{\boldsymbol{\tau}}}^{-1}$, and compute

$$\hat{\boldsymbol{\tau}} = \mathbf{C}_{\tilde{\boldsymbol{\tau}}} \setminus (\mathbf{C}_{\tilde{\boldsymbol{\tau}}} \setminus (\mathbf{S}_y^{-1} \tilde{\mathbf{y}}))$$

by forward and backward substitution. A draw from $\mathcal{N}(\hat{\boldsymbol{\tau}}, \mathbf{D}_{\tilde{\boldsymbol{\tau}}})$ can now be obtained as follows: sample T independent standard normal draws $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and return

$$\tilde{\boldsymbol{\tau}} = \hat{\boldsymbol{\tau}} + \mathbf{C}_{\tilde{\boldsymbol{\tau}}} \setminus \mathbf{z}.$$

Since $\tilde{\boldsymbol{\tau}}$ is an affine transformation of a normal random vector, it is also a normal random vector. It is easy to check that its expectation is $\hat{\boldsymbol{\tau}}$ and its covariance matrix is

$$\mathbf{C}_{\tilde{\boldsymbol{\tau}}}^{-1'} \mathbf{C}_{\tilde{\boldsymbol{\tau}}}^{-1} = (\mathbf{C}_{\tilde{\boldsymbol{\tau}}} \mathbf{C}_{\tilde{\boldsymbol{\tau}}}')$$

as desired. Finally, given the draw $\tilde{\boldsymbol{\tau}} \sim \mathcal{N}(\hat{\boldsymbol{\tau}}, \mathbf{D}_{\tilde{\boldsymbol{\tau}}})$, return $\boldsymbol{\tau} = \mathbf{H}_\psi \tilde{\boldsymbol{\tau}}$.

In Step 2 of the MCMC sampler, we sample from $p(\mathbf{h} | \mathbf{y}, \boldsymbol{\tau}, \boldsymbol{\psi}, \sigma_h^2, \mu_h, \phi_h)$. To proceed, first define $\mathbf{y}^* = \mathbf{H}_\psi^{-1}(\mathbf{y} - \boldsymbol{\tau})$. It follows from (4) that $\mathbf{y}^* = \mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{S}_y)$, where $\mathbf{S}_y = \text{diag}(e^{h_1}, \dots, e^{h_T})$. With this transformation, the auxiliary mixture sampling approach in

Kim, Shepherd, and Chib (1998) can be applied to draw \mathbf{h} efficiently; see also Koop and Korobilis (2010), p. 308–310, for a textbook treatment. Note that in Kim et al. (1998), a forward-backward smoothing algorithm is used; here it is replaced by the precision-based sampler.

For Step 3, note that $\boldsymbol{\psi}$, σ_h^2 , and σ_τ^2 are conditionally independent given \mathbf{y} , $\boldsymbol{\tau}$, and \mathbf{h} . Hence, we can sample each sequentially. Given the prior $p(\boldsymbol{\psi})$, it follows from (5) that

$$\begin{aligned} \log p(\boldsymbol{\psi} | \mathbf{y}, \boldsymbol{\tau}, \mathbf{h}) &\propto \log p(\boldsymbol{\psi}) + \log p(\mathbf{y} | \boldsymbol{\psi}, \boldsymbol{\tau}, \mathbf{h}) \\ &\propto \log p(\boldsymbol{\psi}) - \frac{1}{2} \mathbf{y}^* \mathbf{S}_y^{-1} \mathbf{y}^*, \end{aligned}$$

where we use the transformation $\mathbf{y}^* = \mathbf{H}_\psi^{-1}(\mathbf{y} - \boldsymbol{\tau})$. Hence, $\log p(\boldsymbol{\psi} | \mathbf{y}, \boldsymbol{\tau}, \mathbf{h})$ can be quickly evaluated for any $\boldsymbol{\psi}$ given \mathbf{y} , $\boldsymbol{\tau}$ and \mathbf{h} using the method discussed in the previous section. Since in typical applications $\boldsymbol{\psi}$ is low dimensional—only a few moving average terms are needed—one can maximize $\log p(\boldsymbol{\psi} | \mathbf{y}, \boldsymbol{\tau}, \mathbf{h})$ numerically and obtain the mode and the negative Hessian evaluated at the mode, denoted as $\hat{\boldsymbol{\psi}}$ and \mathbf{K} , respectively. Then, draws from $p(\boldsymbol{\psi} | \mathbf{y}, \boldsymbol{\tau}, \mathbf{h})$ are obtained using an independence-chain Metropolis-Hastings step with proposal density $\mathcal{N}(\hat{\boldsymbol{\psi}}, \mathbf{K}^{-1})$, for example. When $\boldsymbol{\psi}$ is high dimensional, one can avoid the high-dimensional numerical maximization by implementing, e.g., various adaptive MCMC samplers discussed in Andrieu and Thoms (2008).

Next, both $p(\sigma_h^2 | \mathbf{h}, \mu_h, \phi_h)$ and $p(\sigma_\tau^2 | \boldsymbol{\tau})$ are inverse-gamma densities, and can therefore be sampled using standard methods. In fact, we have

$$(\sigma_\tau^2 | \boldsymbol{\tau}) \sim \mathcal{IG} \left(\nu_\tau + (T - 1)/2, \tilde{S}_\tau \right), \quad (\sigma_h^2 | \mathbf{h}, \mu_h, \phi_h) \sim \mathcal{IG} \left(\nu_h + T/2, \tilde{S}_h \right),$$

where $\tilde{S}_\tau = S_\tau + \sum_{t=2}^T (\tau_t - \tau_{t-1})^2/2$ and $\tilde{S}_h = S_h + [(1 - \phi_h^2)(h_1 - \mu_h)^2 + \sum_{t=2}^T (h_t - \mu_h - \phi_h(h_{t-1} - \mu_h))^2]/2$. Lastly, Steps 4 and 5 are standard and can be performed, e.g., as described in Kim et al. (1998) by simply changing the priors.

We note that the estimation for other conditional mean processes can be implemented analogously; we provide the estimation details for the autoregressive MA-SV model in the appendix.

3.3 Computation Efficiency

In this section we briefly discuss some computation issues and the scalability of the algorithm introduced in the previous section. In a typical finance application, for example, one might have time series with thousands of observations. It is therefore important that the estimation method can handle data with large T . In what follows, we discuss the computational complexity of the proposed posterior sampler as T gets large. Again, for concreteness we consider in particular the unobserved components model.

First, the conditional distribution $p(\boldsymbol{\tau} | \mathbf{y}, \mathbf{h}, \boldsymbol{\psi}, \sigma_\tau^2)$ is Gaussian, and draws from this distribution can be obtained exactly using the precision-based method. In particular, obtaining

a draw requires the Cholesky decomposition of the banded precision matrix, as well as a few forward/backward substitutions. Each of these operations involves only $\mathcal{O}(T)$ operations. In Step 2 of the posterior sampler, we first transform the data $\mathbf{y}^* = \mathbf{H}_\psi^{-1}(\mathbf{y} - \boldsymbol{\tau})$, which involves $\mathcal{O}(T)$ operations as \mathbf{H}_ψ is a banded matrix. Then, we directly apply the auxiliary mixture sampling of Kim et al. (1998). Hence, this step is as efficient as the corresponding step in standard SV models.

Next, in Step 3 of the posterior sampler, the conditional distribution $p(\boldsymbol{\psi} | \mathbf{y}, \boldsymbol{\tau}, \mathbf{h})$ is non-standard, and we numerically maximize $\log p(\boldsymbol{\psi} | \mathbf{y}, \boldsymbol{\tau}, \mathbf{h})$ to obtain the mode $\hat{\boldsymbol{\psi}}$ and \mathbf{K} , the negative Hessian evaluated at the mode. We then implement a Metropolis-Hastings step with the proposal density $\mathcal{N}(\hat{\boldsymbol{\psi}}, \mathbf{K}^{-1})$. Evaluation of $p(\boldsymbol{\psi} | \mathbf{y}, \boldsymbol{\tau}, \mathbf{h})$ is done using the method described in Section 3.1, which again requires $\mathcal{O}(T)$ operations. This step is sufficiently efficient in typical applications where $\boldsymbol{\psi}$ is low dimensional. When $\boldsymbol{\psi}$ is high dimensional, one can avoid the high-dimensional numerical maximization by implementing, e.g., adaptive MCMC samplers such as those discussed in Andrieu and Thoms (2008). Lastly, the remaining steps—drawing from $p(\sigma_h^2 | \mathbf{h}, \mu_h, \phi_h)$, $p(\sigma_\tau^2 | \boldsymbol{\tau})$, $p(\mu_h | \mathbf{h}, \sigma_h^2, \phi_h)$ and $p(\phi_h | \mathbf{h}, \mu_h, \sigma_h^2)$ —require trivial computation efforts and can be done quickly even when T is large.




Figure 1: AUD/USD daily returns from January 2006 to December 2010.

To get a sense of how long it takes to estimate the unobserved components MA-SV model, we fit the model using daily returns on AUD/USD from January 2006 to December 2010 with a total of $T = 1,303$ observations (see Figure 1). The algorithm is implemented using MATLAB on a desktop with an Intel Core i7-870 @2.93 GHz processor. It takes about 132 seconds to obtain every 10,000 posterior draws.

4 Modeling and Forecasting U.S. Inflation Rate

Now we use the proposed MA-SV models to analyze the behavior of U.S. quarterly CPI inflation, and contrast the results with those obtained from the standard variants with only stochastic volatility. In addition, we also compare the forecast performance of the two classes of models at various forecast horizons. Since allowing for stochastic volatility is found to be empirically important, unless stated otherwise, models considered in this section all have stochastic volatility in the measurement equation.

4.1 Competing Models

We consider four popular specifications for modeling inflation, and for each we have two versions: with and without the MA component. The primary goal of this exercise is not to find the best model *per se*. Rather, our objective is to investigate if the addition of the MA component improves model-fit and forecast performance, and how it affects the estimates of the states and other parameters of interest. The first specification is the unobserved components model, which we reproduce here for convenience:

$$\begin{aligned} y_t &= \tau_t + u_t, & u_t &\sim \mathcal{N}(0, e^{h_t}), \\ \tau_t &= \tau_{t-1} + \varepsilon_t^\tau, & \varepsilon_t^\tau &\sim \mathcal{N}(0, \sigma_\tau^2), \\ h_t &= \mu_h + \phi_h(h_{t-1} - \mu_h) + \varepsilon_t^h, & \varepsilon_t^h &\sim \mathcal{N}(0, \sigma_h^2). \end{aligned}$$

We refer to this version of unobserved components model as the **UC** model. Stock and Watson (2007) extend this specification to include stochastic volatility in the transition equation for the underlying trend τ_t . Specifically, the variance for ε_t^τ —instead of being fixed to be a constant—is allowed to be time-varying: $\varepsilon_t^\tau \sim \mathcal{N}(0, e^{g_t})$. The log-volatility g_t , in turn, evolves as a stationary AR(1) process:³

$$g_t = \mu_g + \phi_g(g_{t-1} - \mu_g) + \varepsilon_t^g, \quad \varepsilon_t^g \sim \mathcal{N}(0, \sigma_g^2),$$

where $|\phi_g| < 1$. This version of unobserved components model, where both the measurement and transition equations have stochastic volatility, is referred to as the **UCSV** model. The final two specifications are autoregressive models:

$$y_t = \phi_0 + \phi_1 y_{t-1} + \cdots + \phi_p y_{t-p} + u_t,$$

where u_t has the same stochastic volatility specification as before. In addition, we impose the conditions that the roots of the characteristic polynomial associated with the AR coefficients all lie outside the unit circle, so that the AR process is stationary. We consider two lag lengths $p = 1$ and $p = 2$, and we refer them as **AR(1)** and **AR(2)** respectively.

For each of the four models—**UC**, **UCSV**, **AR(1)** and **AR(2)**—we include the MA-SV variants as specified in (1)–(2), and we refer them as **UC-MA**, **UCSV-MA**, **AR(1)-MA**

³In Stock and Watson (2007) the stochastic volatilities h_t and g_t evolve as random walks, which may be viewed as a limiting case considered here with $\phi_g = \phi_h = 1$.

and **AR(2)-MA** respectively. For these models with the moving average components, we set q , the number of moving average terms, to be one. Empirical evidence supporting this choice will be presented in Section 4.3. We summarize all eight specifications in Table 1.

Table 1: A list of competing models.

Model	Description
UC	unobserved components model
UC-MA	same as UC but with an MA(1) component
UCSV	same as UC but the state equation for τ_t has SV
UCSV-MA	same as UCSV but with an MA(1) component
AR(1)	autoregressive model with 1 lag
AR(1)-MA	same as AR(1) but with an MA(1) component
AR(2)	autoregressive model with 2 lags
AR(2)-MA	same as AR(2) but with an MA(1) component
UC-MA-NoSV	same as UC-MA but without SV
AR(1)-MA-NoSV	same as AR(1)-MA but without SV
AR(2)-MA-NoSV	same as AR(2)-MA but without SV

In the forecasting exercise, we also include three additional models without stochastic volatility for comparison. Specifically, **UC-MA-NoSV** is the same as **UC-MA**, but the error term in the measurement equation u_t now has constant variance: $u_t \sim \mathcal{N}(0, \sigma_y^2)$. The models **AR(1)-MA-NoSV** and **AR(2)-MA-NoSV** are defined similarly.

4.2 Data and Priors

The data consist of U.S. quarterly CPI inflation from 1947Q1 to 2011Q3. More specifically, given the quarterly CPI figures z_t , we use $y_t = 400 \log(z_t/z_{t-1})$ as the CPI inflation. A plot of the data is given in Figure 2. For easy comparison, we choose broadly similar priors across models. In particular, we use exactly the same priors for the common model parameters in each pair of models with and without the MA component. For the MA coefficient ψ_1 , we assume the truncated normal prior $\psi_1 \sim \mathcal{N}(\psi_0, V_\psi) \mathbf{1}(|\psi_1| < 1)$ so that the MA process is invertible. We set $\psi_0 = 0$ and $V_\psi = 1$. The prior distribution thus centers around 0 and has support within the interval $(-1, 1)$. Given the large prior variance, it is also relatively non-informative.

As discussed in Section 3.2, we assume independent inverse-gamma priors for σ_τ^2 and σ_h^2 : $\sigma_\tau^2 \sim \mathcal{IG}(\nu_\tau, S_\tau)$ and $\sigma_h^2 \sim \mathcal{IG}(\nu_h, S_h)$. We choose relatively small—hence relatively non-informative—values for the degrees of freedom parameters: $\nu_\tau = \nu_h = 10$. For the scale parameters, we set $S_\tau = 0.18$ and $S_h = 0.45$. These values imply $\mathbb{E} \sigma_\tau^2 = 0.141^2$ and $\mathbb{E} \sigma_h^2 = 0.224^2$. The chosen prior means reflect the desired smoothness of the corresponding state transition, and are comparable to those used in previous studies, such as Chan et al. (2013) and Stock and Watson (2007). As for μ_h and ϕ_h , their priors are respectively

normal and truncated normal: $\mu_h \sim \mathcal{N}(\mu_{h0}, V_{\mu_h})$ and $\phi_h \sim \mathcal{N}(\phi_{h0}, V_{\phi_h})\mathbf{1}(|\phi_h| < 1)$, with $\mu_{h0} = 0$, $V_{\mu_h} = 5$, $\phi_{h0} = 0.9$ and $V_{\phi_h} = 1$.




Figure 2: U.S. quarterly CPI inflation from 1947Q1 to 2011Q3.

For the models **UCSV** and **UCSV-MA** where both the measurement and state equations have stochastic volatility, we follow Stock and Watson (2007) and fix σ_h^2 and σ_g^2 . In particular, we set $\sigma_h^2 = \sigma_g^2 = 0.224^2$. For the autoregressive models, we assume a truncated normal prior for the AR coefficients restricted to the stationary region \mathbf{A}_ϕ : $\phi \sim \mathcal{N}(\phi_0, \mathbf{V}_\phi)\mathbf{1}(\phi \in \mathbf{A}_\phi)$, where $\phi = (\phi_0, \phi_1)'$ under the **AR(1)** and **AR(1)-MA** models, and $\phi = (\phi_0, \phi_1, \phi_2)'$ under the **AR(2)** and **AR(2)-MA** models. Further we set $\phi_0 = \mathbf{0}$ and $\mathbf{V}_\phi = 5 \times \mathbf{I}$. Finally, for the three models without stochastic volatility—**UCMA-NoSV**, **AR(1)-MA-NoSV** and **AR(2)-MA-NoSV**—the error variance of the measurement equation σ_y^2 is assumed to have an inverse-gamma prior $\sigma_y^2 \sim \mathcal{IG}(\nu_y, S_y)$ with $\nu_y = 10$ and $S_y = 9$. This implies $\mathbb{E}\sigma_y^2 = 1$, which is comparable to the stochastic volatility specifications where the prior mean for μ_h , the unconditional mean of the log-volatility h_t , is $\mu_{h0} = 0$ (and hence $\exp(\mu_{h0}) = 1$).

4.3 Full Sample Estimation Results

We present in this section the empirical results for the first eight models listed in Table 1, obtained using the full sample from 1947Q1 to 2011Q3. All the posterior moments and quantiles are based on 50,000 draws from the MCMC algorithm introduced in Section 3.2 after a burnin period of 5,000.

Figure 3 presents point estimates (posterior means) of the underlying inflation τ_t for four models, as well as credible intervals obtained under the MA variants. It is clear from the figure that there are large differences between estimates under the **UC** and **UC-MA**

models. In particular, by allowing for an extra channel for persistence through the moving average errors, the latter model produces much smoother estimates, which are more in line with the notion of a smooth, gradually changing underlying inflation. This finding is broadly consistent with those reported in earlier studies, such as Clark and Doh (2011) and Chan et al. (2013), who also find that by explicitly modeling short-run dynamics one often obtains smoother, more reasonable, underlying inflation estimates.




Figure 3: Posterior estimates and quantiles for τ_t under the **UC**, **UC-MA**, **UCSV** and **UCSV-MA** models. The posterior quantiles are obtained under the MA variants.

On the other hand, the underlying inflation estimates for the **UCSV** and **UCSV-MA** models are very similar, except for the early sample where those for the former appear to be more erratic. By allowing the volatility in the transition equation for τ_t to be time-varying, the model attributes much of the variation in observed inflation to variation in τ_t , with the consequence that it traces closely the actual inflation. As such, the moving average errors in the measurement equation play a lesser role in channeling the short-run dynamics of the series. The erratic underlying inflation estimates, which change rapidly in short periods, cast doubt on the overall appropriateness of the **UCSV** model. Regardless, for our purpose it suffices to note that the MA component does help to obtain smoother estimates for the underlying inflation, though to a lesser extent than in the **UC** model.

In Figure 4 we report the estimates for $\exp(h_t/2)$. Under all specifications we find substantial variation in the inflation volatility estimates, highlighting the importance of allowing for stochastic volatility in the measurement equation. In addition, these results also suggest that estimates for $\exp(h_t/2)$ are somewhat sensitive to modeling assumptions, especially the treatment of the error process in the **UCSV** model.




Figure 4: Posterior estimates for $\exp(h_t/2)$ under the eight models.

A key parameter of interest is the MA coefficient ψ_1 —if the posterior density for ψ_1 is concentrated around 0, it would indicate that the MA component might not be necessary. Table 2 reports the posterior means and standard deviations for ψ_1 , as well as the probabilities that ψ_1 is positive, under the four MA-SV models. It is clear from the table that the posterior means for ψ_1 are all quite large in absolute value. For instance, the MA coefficient is estimated to be about 0.463 in the **UC-MA** model, indicating that even after controlling for the underlying inflation and stochastic volatility, there is still considerable positive autocorrelation in the observed inflation. In fact, the posterior probability that ψ_1 is positive is estimated to be 1, i.e., out of the 50,000 posterior draws for ψ_1 obtained, none are less than zero.

Table 2: Posterior means, standard deviations, and the probabilities of being positive for ψ_1 under the four MA-SV models.

	UC-MA	UCSV-MA	AR(1)-MA	AR(2)-MA
$\mathbb{E}(\psi_1 \mathbf{y})$	0.463	0.307	-0.374	-0.378
$\sqrt{\text{Var}(\psi_1 \mathbf{y})}$	0.068	0.107	0.074	0.138
$\mathbb{P}(\psi_1 > 0 \mathbf{y})$	1.00	0.993	0.000	0.007

Another interesting point to note is that under both versions of unobserved components models, ψ_1 is estimated to be positive, whereas its estimates are negative under both the autoregressive models. This difference is perhaps not surprising: in the former two models the autocorrelation of the observed inflation is modeled only through the random walk underlying trend, which is seemingly not sufficient to capture all the observed positive

autocorrelation. In contrast, past observed inflation rates enter directly the conditional mean process in both the autoregressive models. The result appears to suggest that the AR components, based on actual inflation rates with one or two lags, “over-capture” the observed positive autocorrelation, which in turn induce a negative autocorrelation in the residual.

Next, we compute the marginal density $p(\psi_1 | \mathbf{y})$ for the four moving average models. Since the conditional density $p(\psi_1 | \mathbf{y}, \mathbf{h}, \boldsymbol{\mu})$ has support in $(-1, 1)$ and is known up to a normalizing constant, it can be evaluated on a grid (and normalized so that the area under the curve is one). Then, we estimate $p(\psi_1 | \mathbf{y})$ using the Monte Carlo average

$$\widehat{p(\psi_1 | \mathbf{y})} = \frac{1}{M} \sum_{i=1}^M p(\psi_1 | \mathbf{y}, \mathbf{h}^{(i)}, \boldsymbol{\mu}^{(i)})$$

by summing over the M posterior draws $\mathbf{h}^{(i)}$ and $\boldsymbol{\mu}^{(i)}$.




Figure 5: Density estimates for $p(\psi_1 | \mathbf{y})$ under the four MA-SV models.

We report in Figure 5 the estimates under the four models. For example, most of the mass for ψ_1 is concentrated between 0.2 and 0.7 in **UC-MA** and between -0.6 and -0.1 in **AR(1)-MA**. The posterior densities under the remaining two models tell a similar story: ψ_1 is unlikely to be a small value around zero. Taken together, these results suggest the relevance of extending the standard stochastic volatility models to include an MA component.

The results in Table 2 and Figure 5 may be viewed as suggestive evidence showing the relevance of the proposed MA component in the stochastic volatility models. In what follows, we perform a formal Bayesian model comparison exercise to compare each pair

of stochastic volatility models (i.e., with and without the MA component) using Bayes factors (see, e.g., Koop, 2003, p. 3–4). Since we are comparing nested models, the Bayes factor in favor of the model that has the MA component against the standard variant can be computed using the Savage-Dickey density ratio (Verdinelli and Wasserman, 1995):

$$\text{BF} = \frac{p(\psi_1 = 0)}{p(\psi_1 = 0 | \mathbf{y})}.$$

In other words, we simply need to evaluate the marginal prior and posterior densities for ψ_1 at 0. The ratio of the two values then gives the relevant Bayes factor. The numerator density is a univariate truncated normal, and can be easily evaluated. The denominator density is not of standard form, but we can estimate it using $p(\widehat{\psi_1 = 0} | \mathbf{y})$.

The results are reported in Table 3. For each of the four pairwise comparisons—**UC-MA** against **UC**, **UCSV-MA** against **UCSV**, **AR(1)-MA** against **AR(1)**, and **AR(2)-MA** against **AR(2)**—there is strong to overwhelming evidence that the data prefer the variant with the MA component. Remember that each pair of the models only differs in an extra parameter ψ_1 , and the stochastic volatility models are standard in the literature. Given the context, these full sample estimation results present strong evidence in favor of the proposed models against their counterparts without the MA component. Not only do the former models fit the data better, they also give more sensible underlying inflation estimates. In the next section, we present forecasting results that show the proposed models also provide more accurate point and density forecasts.

Table 3: Bayes factors in favor of the proposed models against the standard variants with only stochastic volatility.

UC-MA	UCSV-MA	AR(1)-MA	AR(2)-MA
1.78×10^7	5.41	1.75×10^3	6.47

So far we have fixed q , the number of MA terms, to be one. We now briefly discuss the choice of q in general. One natural way to proceed is to view the problem as a model comparison exercise. For example, to compare models with q and $q + 1$ MA terms, we simply need to compute the relevant Bayes factor. Specifically, since we have nested models, the Bayes factor in favor of the MA(q) model can be obtained using the Savage-Dickey density ratio $p(\psi_{q+1} = 0 | \mathbf{y})/p(\psi_{q+1} = 0)$, which can be estimated using the method described previously.

As an illustration, we investigate if there is empirical support for models with two MA terms. First, we plot the density estimates for $p(\psi_2 | \mathbf{y})$ under the four MA-SV models in Figure 6. It is clear that for each of the density, there is substantial mass around 0, indicating that ψ_2 is quite likely to be a small value. This can be viewed as evidence supporting models with only one MA term.




Figure 6: Density estimates for $p(\psi_2 | \mathbf{y})$ under the four MA-SV models.

Then, to formally compare the models with one and two MA terms, we calculate the Bayes factors in favor of the MA(1) specifications against the corresponding MA(2) models. The results are presented in Table 4. In each pairwise comparison, there is substantial to strong evidence that the data prefer the model with one MA term only.

Table 4: Bayes factors in favor of the MA(1) specifications against the corresponding MA(2) models.

UC-MA	UCSV-MA	AR(1)-MA	AR(2)-MA
12.78	3.04	8.62	12.52

4.4 Forecasting Results

We now perform a recursive out-of-sample forecasting exercise to evaluate the performance of the models listed in Table 1 for forecasting U.S. quarterly CPI inflation at various horizons. In addition to the four pairs of models—**UC**, **UCSV**, **AR(1)**, **AR(2)**, and their variants with an MA component—we also include three versions that have only constant variance for comparison. We use each of the eleven models to produce both point and density k -step-ahead iterated forecasts with $k = 1, 4, 8, 12$ and 16 . Specifically, given the data up to time t , denoted as $\mathbf{y}_{1:t}$, we implement the MCMC sampler in Section 3.2 to obtain posterior draws given $\mathbf{y}_{1:t}$. Then, we compute the predictive mean $\mathbb{E}(y_{t+k} | \mathbf{y}_{1:t})$ as the point forecast and the predictive density $p(y_{t+k} | \mathbf{y}_{1:t})$ as the density forecast. Next, we move one period ahead and repeat the whole exercise with data $\mathbf{y}_{1:t+1}$, and so forth. These forecasts are then evaluated for $t = t_0, \dots, T - k$, where t_0 is 1975Q1.

In practice, neither the predictive mean nor the predictive density of y_{t+k} can be computed

analytically. Instead, they are obtained using predictive simulation. More precisely, at every MCMC iteration, given the model parameters and states (up to time t), we simulate future states from time $t + 1$ to $t + k$ using the relevant transition equations. We also simulate future errors $u_s \sim \mathcal{N}(0, e^{h_s})$ or $u_s \sim \mathcal{N}(0, \sigma_y^2)$ for $s = t + 1, \dots, t + k - 1$. Given these draws, y_{t+k} is a normal random variable as specified in (1), and one can easily produce the point and density forecasts for y_{t+k} . Hence, we have a pair of forecasts (point and density) at every MCMC iteration. These forecasts are then averaged over all the posterior draws to produce estimates for $\mathbb{E}(y_{t+k} | \mathbf{y}_{1:t})$ and $p(y_{t+k} | \mathbf{y}_{1:t})$. The whole exercise is then repeated using data up to time $t + 1$ to produce $\mathbb{E}(y_{t+k+1} | \mathbf{y}_{1:t+1})$ and $p(y_{t+k+1} | \mathbf{y}_{1:t+1})$, and so forth.

Let y_{t+k}^o denote the observed value of y_{t+k} that is known at time $t + k$. The metric used to evaluate the point forecasts is the root mean squared forecast error (RMSFE) defined as

$$\text{RMSFE} = \sqrt{\frac{\sum_{t=t_0}^{T-k} (y_{t+k}^o - \mathbb{E}(y_{t+k} | \mathbf{y}_{1:t}))^2}{T - k - t_0 + 1}}.$$

To evaluate the density forecast $p(y_{t+k} | \mathbf{y}_{1:t})$, one natural measure is the predictive likelihood $p(y_{t+k} = y_{t+k}^o | \mathbf{y}_{1:t})$, i.e., the predictive density of y_{t+k} evaluated at the observed value y_{t+k}^o . Clearly, if the actual outcome y_{t+k}^o is unlikely under the density forecast, the value of the predictive likelihood will be small, and vice versa; see, e.g., Geweke and Amisano (2011) for a more detailed discussion of the predictive likelihood and its connection to the marginal likelihood. We evaluate the density forecasts using the sum of log predictive likelihoods:

$$\sum_{t=t_0}^{T-k} \log p(y_{t+k} = y_{t+k}^o | \mathbf{y}_{1:t}).$$

For this metric, a larger value indicates better forecast performance.

Table 5 presents the point forecast results for the eleven models. For easy comparison, we report the ratios of RMSFEs of a given model to those of **UC**. Hence, values smaller than unity indicate better forecast performance than **UC**. A few broad observations can be drawn from these forecasting results. First, except for **UCSV**, there is clear evidence that the proposed models perform better than the standard variants for short-horizon forecasts. For instance, in comparing one-quarter-ahead forecasts, the RMSFE for **UC-MA** is 92% of the value for **UC**. Remember that the latter is among the best forecasting models in the literature. Clearly, these results provide strong evidence that the addition of the MA component gives substantial benefits. Second, for medium-horizon forecasts (e.g., $k \geq 8$), each of the three pairs—**UCSV** and **UCSV-MA**, **AR(1)** and **AR(1)-MA**, **AR(2)** and **AR(2)-MA**—give almost identical RMSFEs. For **UC-MA**, however, it consistently gives better point forecasts compared to **UC**, even at four-year forecast horizon. This might seem surprising at first glance, as the MA component only models short-run dynamics. However, the results in Section 4.3, especially Figure 3, suggest one explanation: by including the MA component, one obtains smoother and more reasonable estimates for the underlying inflation, which in turn help produce better forecasts at longer horizons.

Table 5: Relative RMSFEs for forecasting quarterly CPI inflation relative to **UC**.

	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$
UC	1.00	1.00	1.00	1.00	1.00
UC-MA	0.92	0.98	0.94	0.92	0.93
UCSV	0.94	1.00	1.06	1.08	1.08
UCSV-MA	0.94	1.00	1.06	1.09	1.07
AR(1)	0.96	0.95	0.96	0.96	0.96
AR(1)-MA	0.93	0.95	0.97	0.97	0.96
AR(2)	0.96	0.96	0.97	0.96	0.95
AR(2)-MA	0.94	0.95	0.97	0.97	0.95
UC-MA-NoSV	0.96	1.03	0.97	0.95	0.96
AR(1)-MA-NoSV	0.94	0.99	1.00	0.99	0.99
AR(2)-MA-NoSV	0.94	1.00	1.01	0.99	1.00

The results also show the importance of allowing for stochastic volatility. For example, **UC-MA** dominates the version without stochastic volatility—**UC-MA-NoSV**—at all forecast horizons. Comparing **AR(1)-MA** and **AR(2)-MA** with the corresponding variants without stochastic volatility tells a similar story. Lastly, it is also of interest to note that **UCSV** performs better than **UC** at short-horizon forecasts but worse at longer horizons. This is consistent with the estimation results presented in Figure 3: by including stochastic volatility in the transition for τ_t , the estimates for underlying inflation under **UCSV** traces closely the actual inflation, which is good for short-horizon forecasts but not for longer horizons. However, by including an MA component, **UC-MA** outperforms **UCSV-MA** at all horizons.

Table 6: Sum of log predictive likelihoods for forecasting quarterly CPI inflation relative to **UC**.

	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
UC	0.0	0.0	0.0	0.0	0.0
UC-MA	6.5	6.3	15.0	20.5	20.9
UCSV	0.0	-4.4	-7.7	-12.1	-16.4
UCSV-MA	1.7	-2.1	-6.5	-11.2	-13.2
AR(1)	1.2	2.8	5.5	6.3	8.0
AR(1)-MA	6.9	8.1	10.4	11.4	12.1
AR(2)	3.2	5.6	8.5	8.9	10.2
AR(2)-MA	5.9	7.5	9.9	11.1	11.7
UC-MA-NoSV	-22.3	-16.5	-6.1	-4.5	-4.2
AR(1)-MA-NoSV	-18.8	-13.7	-11.2	-12.5	-9.8
AR(2)-MA-NoSV	-19.1	-18.9	-14.3	-14.2	-11.6

We present the results for the density forecasts in Table 6. For easy comparison, we first compute the sum of log predictive likelihoods of a given model, and from which we

subtract the corresponding value of **UC**. Hence, positive values indicate better forecast performance than **UC**. These results indicate that the addition of the MA component provides substantial benefits not only for short-horizon density forecasts, but also for longer horizons. Moreover, models which do not allow for time-varying variance perform quite badly compared to similar models with stochastic volatility, indicating the important role of stochastic volatility in producing good density forecasts.

5 Concluding Remarks and Future Research

With the aim of expanding the toolkit for analyzing time series data, we have introduced a new class of models that generalizes the standard stochastic volatility specification to include a moving average component. The addition of the MA component leads to models that are more robust to misspecification, which often translates into better forecast performance in practice. A drawback of the new models is that the estimation is more difficult, as the MA component induces serial dependence in the observations. In view of this, we introduce a novel algorithm that exploits the special structure of the models — that the covariance matrix for the observations is sparse, and therefore fast routines for sparse matrices can be used.

We illustrate the relevance of the new class of models with an empirical application involving U.S. CPI inflation. Our empirical results show that the data strongly favor the new models over the standard variants with only SV. In a recursive forecasting exercise, there is clear evidence that the new models deliver improvements in out-of-sample forecasts. A second finding is that by allowing for SV in the error variance and imposing stationary conditions on the conditional mean process, parsimonious ARMA models can be competitive against more complex models such as the unobserved components models. In this paper we have only considered univariate models. For future research it would be interesting to formulate multivariate versions of the proposed MA-SV models, using, e.g., factors with SV, and extend the new estimation methods to those settings.

Appendix

In this appendix we outline a posterior simulator for estimating the **AR-MA** model:

$$\begin{aligned} y_t &= \phi_0 + \phi_1 y_{t-1} + \cdots + \phi_p y_{t-p} + u_t + \psi_1 u_{t-1} + \cdots + \psi_q u_{t-q}, \\ h_t &= \mu_h + \phi_h (h_{t-1} - \mu_h) + \varepsilon_t^h, \end{aligned}$$

with $|\phi_h| < 1$, where $u_t \sim \mathcal{N}(0, e^{h_t})$ and $\varepsilon_t^h \sim \mathcal{N}(0, \sigma_h^2)$. The transition equation for h_t is initialized with $h_1 \sim \mathcal{N}(\mu_h, \sigma_h^2 / (1 - \phi_h^2))$.

Let $\phi = (\phi_0, \phi_1, \dots, \phi_p)'$, and we assume independent priors for ϕ , ψ , ϕ_h , μ_h and σ_h^2 . The priors for ψ , ϕ_h , μ_h and σ_h^2 are given in Section 3.2. As for ϕ , we take the following

truncated multivariate normal prior

$$\boldsymbol{\phi} \sim \mathcal{N}(\boldsymbol{\phi}_0, \mathbf{V}_\phi) \mathbf{1}(\boldsymbol{\phi} \in \mathbf{A}_\phi),$$

where $\boldsymbol{\phi}_0 = \mathbf{0}$, $\mathbf{V}_\phi = 5 \times \mathbf{I}$ and \mathbf{A}_ϕ is the set where the roots of the autoregressive polynomial associated with $\boldsymbol{\phi}$ are all outside the unit circle. The posterior simulator sequentially draws from: (1) $p(\boldsymbol{\phi} | \mathbf{y}, \mathbf{h}, \boldsymbol{\psi})$; (2) $p(\mathbf{h} | \mathbf{y}, \boldsymbol{\phi}, \boldsymbol{\psi}, \mu_h, \phi_h, \sigma_h^2)$; (3) $p(\boldsymbol{\psi}, \sigma_h^2 | \mathbf{y}, \boldsymbol{\phi}, \mathbf{h}, \mu_h, \phi_h) = p(\boldsymbol{\psi} | \mathbf{y}, \boldsymbol{\phi}, \mathbf{h}) p(\sigma_h^2 | \mathbf{h}, \mu_h, \phi_h)$; (4) $p(\mu_h | \mathbf{h}, \sigma_h^2, \phi_h)$; and (5) $p(\phi_h | \mathbf{h}, \mu_h, \sigma_h^2)$.

Steps 2-5 can be implemented as in the sampler given in Section 3.2, and here we focus on Step 1. To this end, we first write the observation equation in matrix form:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\phi} + \mathbf{H}_\psi \mathbf{u},$$

where \mathbf{X} is a $T \times (p + 1)$ matrix of intercepts and lagged observations appropriately defined. Now let $\tilde{\mathbf{y}} = \mathbf{H}_\psi^{-1} \mathbf{y}$ and $\tilde{\mathbf{X}} = \mathbf{H}_\psi^{-1} \mathbf{X}$, both of which can be computed quickly as \mathbf{H}_ψ is banded. Then,

$$(\tilde{\mathbf{y}} | \boldsymbol{\psi}, \boldsymbol{\phi}, \mathbf{h}) \sim \mathcal{N}(\tilde{\mathbf{X}}\boldsymbol{\phi}, \mathbf{S}_y),$$

where $\mathbf{S}_y = \text{diag}(e^{h_1}, \dots, e^{h_T})$. It follows that

$$(\boldsymbol{\phi} | \mathbf{y}, \mathbf{h}, \boldsymbol{\psi}) \sim \mathcal{N}(\hat{\boldsymbol{\phi}}, \mathbf{D}_\phi) \mathbf{1}(\boldsymbol{\phi} \in \mathbf{A}_\phi),$$

where $\mathbf{D}_\phi = (\tilde{\mathbf{X}}' \mathbf{S}_y^{-1} \tilde{\mathbf{X}} + \mathbf{V}_\phi^{-1})^{-1}$ and $\hat{\boldsymbol{\phi}} = \mathbf{D}_\phi \tilde{\mathbf{X}}' \mathbf{S}_y^{-1} \tilde{\mathbf{y}}$. A draw from the truncated density can then be obtained via the acceptance-rejection sampling (e.g., Kroese, Taimre, and Botev, 2011, chapter 3) with proposal density $\mathcal{N}(\hat{\boldsymbol{\phi}}, \mathbf{D}_\phi)$. This step is efficient when $\boldsymbol{\phi}$ is low-dimensional, as is often the case.

References

- C. Andrieu and J. Thoms. A tutorial on adaptive MCMC. *Statistics and Computing*, 18(4):343–373, 2008.
- G. Athanasopoulos and F. Vahid. VARMA versus VAR for macroeconomic forecasting. *Journal of Business and Economic Statistics*, 26(2):237–252, 2008.
- T. Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986.
- J. C. C. Chan and I. Jeliazkov. Efficient simulation and integrated likelihood estimation in state space models. *International Journal of Mathematical Modelling and Numerical Optimisation*, 1:101–120, 2009.
- J. C. C. Chan, G. Koop, R. Leon-Gonzalez, and R. Strachan. Time varying dimension models. *Journal of Business and Economic Statistics*, 30:358–367, 2012.
- J. C. C. Chan, G. Koop, and S. M. Potter. A new model of trend inflation. *Journal of Business and Economic Statistics*, 31:94–106, 2013.

- S. Chib and E. Greenberg. Bayes inference in regression models with ARMA(p, q) errors. *Journal of Econometrics*, 64(12):183 – 206, 1994.
- S. Chib, F. Nardari, and N. Shephard. Markov chain Monte Carlo methods for stochastic volatility models. *Journal of Econometrics*, 108(2):281–316, 2002.
- S. Chib, F. Nardari, and N. Shephard. Analysis of high dimensional multivariate stochastic volatility models. *Journal of Econometrics*, 134(2):341371, 2006.
- T. E. Clark and T. Doh. A Bayesian evaluation of alternative models of trend inflation. Working paper, Federal Reserve Bank of Cleveland, 2011.
- T. Cogley and T. J. Sargent. Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics*, 8(2):262 – 302, 2005.
- T. Cogley and A. Sbordone. Trend inflation, indexation, and inflation persistence in the new keynesian phillips curve. *American Economic Review*, 98:2101–2126, 2008.
- T. Cogley, G. Primiceri, and T. Sargent. Inflation-gap persistence in the U.S. *American Economic Journal: Macroeconomics*, 2:43–69, 2010.
- L. Fahrmeir and H. Kaufmann. On kalman filtering, posterior mode estimation and Fisher scoring in dynamic exponential family regression. *Metrika*, 38(1):37–60, 1991.
- J. Geweke and G. Amisano. Hierarchical Markov normal mixture models with applications to financial asset returns. *Journal of Applied Econometrics*, 26:1–29, 2011.
- G. H. Golub and C. F. van Loan. *Matrix computations*. Johns Hopkins University Press, Baltimore, 1983.
- A. C. Harvey. Trends and cycles in macroeconomic time series. *Journal of Business and Economic Statistics*, 3(3):216–227, 1985.
- E. Jacquier, N. G. Polson, and P. E. Rossi. Bayesian analysis of stochastic volatility models with fat-tails and correlated errors. *Journal of Econometrics*, 122(1):185212, 2004.
- M. J. Jensen and J. M. Maheu. Bayesian semiparametric stochastic volatility modeling. *Journal of Econometrics*, 157(2):306–316, 2010.
- S. Kim, N. Shepherd, and S. Chib. Stochastic volatility: Likelihood inference and comparison with ARCH models. *Review of Economic Studies*, 65(3):361–393, 1998.
- G. Koop. *Bayesian Econometrics*. Wiley & Sons, New York, 2003.
- G. Koop and D. Korobilis. Bayesian multivariate time series methods for empirical macroeconomics. *Foundations and Trends in Econometrics*, 3(4):267–358, 2010.
- G. Koop and D. Korobilis. Forecasting inflation using dynamic model averaging. *International Economic Review*, 2012. Forthcoming.

- G. Koop and S. M. Potter. Estimation and forecasting in models with multiple breaks. *Review of Economic Studies*, 74:763–789, 2007.
- D. Korobilis. VAR forecasting using Bayesian variable selection. *Journal of Applied Econometrics*, 2012. Forthcoming.
- D. P. Kroese, T. Taimre, and Z. I. Botev. *Handbook of Monte Carlo Methods*. John Wiley & Sons, New York, 2011.
- W. J. McCausland, S. Miller, and D. Pelletier. Simulation smoothing for state-space models: A computational efficiency analysis. *Computational Statistics and Data Analysis*, 55:199–212, 2011.
- J. Nakajima and Y. Omori. Stochastic volatility model with leverage and asymmetrically heavy-tailed error using GH skew Student’s t -distribution. *Computational Statistics & Data Analysis*, 2012. Forthcoming.
- T. Nakatsuma. Bayesian analysis of ARMA-GARCH models: A Markov chain sampling approach. *Journal of Econometrics*, 95(1):57–69, 2000.
- Y. Omori, S. Chib, N. Shephard, and J. Nakajima. Stochastic volatility with leverage: Fast and efficient likelihood inference. *Journal of Econometrics*, 140(2):425–449, 2007.
- G. E. Primiceri. Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, 72(3):821–852, 2005.
- H. Rue, S. Martino, and N. Chopin. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace. *Journal of the Royal Statistical Society Series B*, 71:319–392, 2009.
- R. Ruiz-Cardenas, E. Krainski, and H. Rue. Direct fitting of dynamic models using integrated nested laplace approximations—INLA. *Computational Statistics and Data Analysis*, 56(6):1808–1828, 2012.
- J. H. Stock and M. W. Watson. Why has U.S. inflation become harder to forecast? *Journal of Money Credit and Banking*, 39:3–33, 2007.
- J. H. Stock and M. W. Watson. Modeling inflation after the crisis. Working Paper 16488, National Bureau of Economic Research, October 2010.
- S. J. Taylor. Modelling stochastic volatility. *Mathematical Finance*, 4:183–204, 1994.
- I. Verdinelli and L. Wasserman. Computing Bayes factors using a generalization of the Savage-Dickey density ratio. *Journal of the American Statistical Association*, 90(430):614–618, 1995.