Has the Volatility of U.S. Inflation Changed and How?

Grassi, Stefano and Proietti, Tommaso

SEFEMEQ, University of Rome "Tor Vergata"

7 November 2008

Has the Volatility of U.S. Inflation Changed and How?

Stefano Grassi* **Tommaso Proietti**

S.E.F. e ME.Q., University of Rome "Tor Vergata"

Abstract

The local level model with stochastic volatility, recently proposed for U.S. by Stock and Watson (*Why Has U.S. Inflation Become Harder to Forecast?*, Journal of Money, Credit and Banking, Supplement to Vol. 39, No. 1, February 2007), provides a simple yet sufficiently rich framework for characterizing the evolution of the main stylized facts concerning the U.S. inflation. The model decomposes inflation into a core component, evolving as a random walk, and a transitory component. The volatility of the disturbances driving both components is allowed to vary over time. The paper provides a full Bayesian analysis of this model and readdresses some of the main issues that were raised by the literature concerning the evolution of persistence and predictability and the extent and timing of the great moderation. The assessment of various nested models of inflation volatility and systematic model selection provide strong evidence in favor of a model with heteroscedastic disturbances in the core component, whereas the transitory component has time invariant size. The main evidence is that the great moderation is over, and that volatility, persistence and predictability of inflation underwent a turning point in the late 1990s. During the last decade volatility and persistence have been increasing and predictability has been going down.

Keywords: Marginal Likelihood; Bayesian Model Comparison; Stochastic Volatility; Great Moderation; Inflation Persistence.

*Address for Correspondence: Via Columbia 2, 000133 Rome, Italy. Tel +393343534438. E-Mail: stefanograssi@hotmail.com

1 Introduction

Inflation's volatility has attracted a good deal of attention recently; the interest has been sparked by the debate on the Great Moderation that has been documented for real economic aggregates. Inflation stabilization is indeed a possible source of the reduction in the volatility of macroeconomic aggregates. The issue is also closely bound up with inflation persistence and predictability. In an influential paper Stock and Watson (2007), using a local level model with stochastic volatility, document that inflation is less volatile now than it was in the 1970s and early 1980s; moreover, persistence, which measure the long run effect of a shock, has declined, and predictability has increased.

There is still an ongoing debate about the statistical significance of inflation persistence and its stability over time, see Pivetta and Reis (2007), Cogley, Primicieri and Sargent (2008), Cecchetti et al. (2007), among others. Recently Bos, Koopman and Ooms (2008) analyzed a U.S. core inflation series (excluding food and energy) as a long memory process subject to heteroscedastic shocks and documented remarkable changes, taking place about at the time of the Great Moderation (1984), in the variance of the series and that of the volatility process, the fractional integration parameter (which is the measure of persistence adopted in that paper), in the short memory characteristics of the series.

In this paper we consider the simple unobserved components model of U.S. inflation considered in Stock and Watson (2007), referred to as the local level model with stochastic volatility (UC-SV). The model provides a simple but yet sufficiently rich framework for discussing the main stylized facts concerning inflation, such as the changes in persistence and predictability. The model postulates the decomposition of observed inflation into two components: the core component (or underlying inflation) which captures the trend in inflation, and the transitory component, which captures the deviations of inflation from its trend value. We will start from a specification such that both components are driven by disturbances whose variance evolves over time according to a stationary stochastic volatility process, and will attempt to assess the significance of the changing volatility in each of the components.

The contributions of this paper are the following: we provide a full Bayesian analysis, so that unlike the current literature, we do not assume that some of the parameters, namely the variances of the stochastic volatility components, are known. Secondly, we carry out systematic model selection by comparing the marginal likelihood implied by the different models of inflation volatility. The marginal likelihood is estimated according to the Chib and Jeliazkov (2001) algorithm.

The interesting final result is that we find strong support for the specification with stochastic volatility in the core component but not in both. We document that persistence is higher than in previous studies and is subject to a significant reduction only at the beginning of the 2000s, whereas predictability has increased somewhat at about the same time.

This paper is organized as follows. In Section 2 we present the local level model with stochastic volatility. Section 3 illustrates the Monte Carlo Markov Chain (MCMC) sampling scheme used to perform Bayesian inference for this model. In Section 4 we present and discuss the estimation results. In Section 5 we describe the Chib and Jeliazkov (2001) approach to the evaluation of the marginal likelihood. The results are used to select the final model among four competitors. In Section 6 we conclude the paper.

2 The UC-SV model

The paper focuses on the quarterly inflation rate constructed from the Consumer Price Index (All Urban Consumers, seasonally adjusted), made available by the U.S. Bureau of Labor Statistics. The quarterly index is obtained from the monthly index by computing the average of the three months that make up each quarter; if we denote the quarterly CPI by P_t, the annualized quarterly inflation rate is then computed as $400\Delta \ln P_t$ and is denoted y_t, $t = 1, ..., n$. The series is plotted in figure 1 and is available for the sample sample period 1960:1 –2008:3.

The most general specification of the UC-SV model with stochastic volatility represents inflation as the sum of an underlying level, denoted here by α_t, which evolves as a random walk, and a transitory component

$$\begin{aligned} y_t &= \alpha_t + \sigma_{\varepsilon t}\varepsilon_t, \qquad \varepsilon_t \sim N(0,1) \\ \alpha_t &= \alpha_{t-1} + \sigma_{\eta t}\eta_t, \qquad \eta_t \sim N(0,1) \end{aligned} \tag{1}$$

where ε_t and η_t are independent standard normal Gaussian disturbances and their size, $\sigma_{\eta t}$ and $\sigma_{\varepsilon t}$, respectively evolve over time according to a SV process. Denoting $h_{1t} = \ln \sigma_{\eta t}$ and $h_{2t} = \ln \sigma_{\varepsilon t}$

$$\begin{aligned} h_{1,t} &= \mu_1 + \phi_1 h_{1,t-1} + \kappa_{1,t}, \qquad h_{1,0} \sim N\left(0, \frac{\sigma^2_{\kappa_1}}{1-\phi_1^2}\right), \qquad \kappa_1 \sim N(0,1) \\ h_{2,t} &= \mu_2 + \phi_2 h_{2,t-1} + \kappa_{2,t}, \qquad h_{2,0} \sim N\left(0, \frac{\sigma^2_{\kappa_2}}{1-\phi_2^2}\right), \qquad \kappa_2 \sim N(0,1) \end{aligned} \tag{2}$$

The model encompasses the traditional stochastic volatility model that is widely used in finance (see for instance Shephard, 2006), which arises when the process α_t degenerates to a constant. The specification of the stochastic volatility processes differ from Stock and Watson (2007) and Cecchetti et al. (2008), who assume a random walk process for the log-variances $h_{it}, i = 1, 2$. When both standard deviations $\sigma_{\varepsilon t}$ and $\sigma_{\eta t}$ do not vary with time, the model reduces to the the traditional local level model. The latter has a IMA(1,1) reduced form $\Delta y_t = \xi_t + \vartheta \xi_{t-1}$ with parameter $\vartheta = \left[(q^2 + 4q)^{\frac{1}{2}} - 2 - q]/2\right]$, where $q = \sigma^2_\eta/\sigma^2_\epsilon$ denotes the signal to noise ratio.

The local level model has a long tradition and a well-established role in the analysis of economic time series, since it provides the model-based interpretation for the popular forecasting technique known as *exponential smoothing*, which is widely used in applied economic forecasting and fares remarkably well in forecast competitions; see Muth (1960) and the comprehensive reviews by Gardner (1985, 2006). In the sequel we shall also consider the cases when either $\sigma_{\varepsilon t}$ or $\sigma_{\eta t}$ is constant.

The UC-SV model can be considered as a IMA(1,1) with time-varying moving average parameter. Hence, the local measure of persistence that we consider is obtained as $(1 + \vartheta_t)$, where ϑ_t varies with time according to the values of the time-varying signal to noise ratio $q_t = \sigma^2_{\eta t}/\sigma^2_{\epsilon t}$ Cecchetti et al. (2007) use the implied time varying first order autocorrelation of Δy_t, as a measure of persistence.

Predictability can be defined in terms of the Granger and Newbold (1986, p. 310) forecastability index

$$\text{Pred}_t = 1 - \frac{\text{Var}(\xi_t|h_{it})}{\text{Var}(\Delta y_t|h_{it})} \tag{3}$$

In terms of the parameters of the UC-SV, the prediction error variance equals $\text{Var}(\xi_t|h_{it}) = \frac{\sigma^2_{\eta t}}{(1+\vartheta^2)}$, whereas the variance $\text{Var}(\Delta y_t|h_{it}) = \sigma^2_{\eta t} + 2\sigma^2_{\epsilon t}$.

3 Bayesian Estimation

This section provides an overview of the MCMC methodology adopted for the estimation of the UC-SV model. All inferences are based on a Gibbs sampling scheme, according to which samples are drawn componentwise from the full conditionals; for the components which cannot be sampled directly a Metropolis-Hasting sub-chain is used within the Gibbs sampling cycle. In particular, the posterior of the AR parameters ϕ_1, ϕ_2, is not available in closed form, see Bos and Shepard (2006) and Kim et al. (1998). More details on the specification of the prior distributions, the full conditionals and the Metropolis-within Gibbs steps are provided in Appendix A.

Let $\theta = (\mu_1, \mu_2, \phi_1, \phi_2, \sigma^2_{\kappa_1}, \sigma^2_{\kappa_2})$ denote the vector of hyperparameters, $h_i, i = 1, 2$, be the collection of the values of the latent stochastic volatility processes for $i = 1, 2$ and α and y denote the stack of core inflation and values of y_t.

The Gibbs sampling scheme can be sketched as follows:

1. Initialize h_i, θ

2. Draw a sample from $\theta, \alpha | y, h_i$

   a) Sample θ from $\theta | y, \alpha, h_i$ (see Appendix A).

   b) Sample α from $\alpha | y, \theta, h_i$ using the simulation smoother of Durbin and Koopman (2002).

3. Sample $h_i, i = 1, 2$, from $h_i | \alpha, y, \theta$ using an Independent Metropolis-Hastings algorithm;

4. Go to 2.

The most complex part of the algorithm deals with the stochastic volatility processes; we use a single move sampler based on the density:

$$h_{it} | h_{i,t+1}, h_{i,t-1}, y_t, \alpha_{t-1}, \alpha_t, \tag{4}$$

we implement a Independent Metropolis-Hastings algorithm, for a detailed description refer to Cappé et. al.(2007) and the Appendix. In order to sample from the full conditional we use the following results:

$$f(h_{i,t} | h_{i,t-1}, h_{i,t+1}, y_t, \alpha_t, \alpha_{t-1}) \propto f(h_{i,t} | h_{i,t-1}) f(y_t | \alpha_t, h_{1,t}) f(\alpha_t | \alpha_{t-1}, h_{2,t}) \tag{5}$$

In the Appendix A the necessary steps to implement the Independent Metropolis-Hastings are explained.

4 Estimation Results

We report the results of the Bayesian estimation for the model presented in section 2. We inizialized the MCMC sampler by setting all $h_{i,t} = 0$ and $\phi_i = 0.86$, $\sigma_i^2 = 0.07$ and $\mu = 0.6$. We iterated the algorithm on the log-volatilities for 1000 iterations and then the parameters and the log-volatilities for 15000 more times before recording the draws from a subsequent 25000 iterations. The programm is written in Ox v. 5.10 console (Doornik (2007)) using our source code. The time needed for all calculations (including the additional simulations required to evaluate the marginal likelihood with the Chib and Jeliazkov method) is about 35 minutes.

Figure 2 shows the inflation series with the posterior mean of the core component, and the two stochastic volatilities components for the irregular and the core. The second and third panels show that the volatility of the core component is increasing from 1960 to 1982, and it is slowly decreasing until 2000. After the year 2000 there an increase in the volatility of both components.

Figure 3 displays the evolution of the Signal to Noise Ratio, the Persistence parameter, defined as in section 2, the Prediction error variance and the Predictability measure. The graph reveals that the size of the random walk component increases during the 70s and it is lower in the 80s and is subject to a sharp fall around the year 2000. Persistence is roughly constant at values well below 1 and there is evidence for the presence of a break again around 2000. The robustness of these results will be discussed later. As far as predictability is concerned, the prediction error variance undergoes a decline after 1982 (this is consistent with the results of Bos, Koopman and Ooms, 2008), but has been increasing after the year 2000. In relative terms, the forecastability index shows only an increase in the recent years.
Table I reports some summary statistics concerning the posterior distribution of the parameters and some convergence diagnostics. We notice that the volatility of the core component is higher than the irregular one. The convergence properties of the chain are satisfactory, although the Geweke statistic for the parameter μ_1 is significant.

Table 1: Posterior, Median, Geweke statistic and Inefficiency factor for UC-SV model

Parameters	Mean	Median	Geweke statistic	Inefficinecy factor (taper 0.05)
μ_1	-0.0193	-0.0188	2.52	1.41
μ_2	-0.0209	-0.0204	0.98	3.32
ϕ_1	0.9883	0.9890	-0.09	5.92
ϕ_2	0.9856	0.9864	1.12	6.07
$\sigma^2_{\kappa_1}$	0.0442	0.0436	-0.07	55.26
$\sigma^2_{\kappa_2}$	0.0515	0.0507	1.35	11.66

5 Model Selection

Thus far the literature has focused on fitting the UC-SV model (sometimes with arbitrary restrictions on the parameters $\sigma^2_{\kappa i}$) and describing the statistical evidence. There is a potential danger that the UC-SV model could be overfitting the data, but little or no attention has been devoted to careful model selection.

In this paper we perform Bayesian model selection; the models under comparison are the following four variants of the local level model:

- M_1: the Local Level Model without SV disturbances (UC);
- M_2: the Local Level Model with a SV disturbance only on the transitory component (UC-SVt);
- M_3: the Local Level Model with a SV disturbance only on the core component (UC-SVc);
- M_4: the Local Level Model with two SV disturbances (UC-SV).

Bayesian model comparison entails the computation of posterior model probabilities, see Geweke (2005) for more details. If the models have the same prior probability, the ratio of the posterior mode probabilities is the Bayes factor, which is the ratio of the marginal likelihoods of two rival specifications. The main difficulty lies with the evaluation of the marginal likelihood. For this purpose we adopt the method proposed by Chib and Jeliazkov (2001), which is based on the MCMC output, and additional draws from given partial full conditionals.

Denoting by $p(y|\theta_k, M_k)$ the density function of the data under model M_k, with parameter vector θ_k, and by $p(\theta_k|M_k)$ the priors densities, the Chib and Jeliazkov(2001) approach is based on the following basic marginal likelihood identity:

$$m(y|M_k) = \frac{f(y|M_k, \theta_k)\pi(\theta_k|M_k)}{\pi(\theta_k|y, M_k)}, \qquad k = 1, 2, 3, 4. \tag{6}$$

The formal Bayesian approach for comparing model M_1, M_2, M_3 and M_4, is through the pairwise Bayes factor, defined as the ratio of marginal likelihoods:

$$B_{1,2} = \frac{m(y|M_1)}{m(y|M_2)} \qquad B_{2,3} = \frac{m(y|M_2)}{m(y|M_3)} \qquad B_{3,4} = \frac{m(y|M_3)}{m(y|M_4)}$$

which can also be interpreted as the posterior probability of model M_1,model M_2 , model M_3, and model M_4, when both models are, a priori, equally likely.

Taking the logarithms of (6) and evaluating this function at some hight density point θ_k^* we have:

$$\log m(y|M_k) = \log f(y|M_k, \theta_k^*) + \log \pi(\theta_k^*|M_k) - \log \pi(\theta_k^*|y, M_k) \tag{7}$$

The first terms of the RHS of equation (16) have a closed form expression, and can be evaluated, for the four models, by the Kalman filter; the second component is simply the product of the prior distribution for the parameters of each model. The last component, i.e. the normalized posterior density of the parameters, requires a specialized treatment. In Appendix B we provide the relevant details for its estimation, for the UC-SV more general specification.

The estimates of the marginal likelihood for our particular application are reported in the following table. The results clearly point out that the model that performs best is the local level

Table 2: Marginal likelihood for UC models of U.S. inflation

Models	$\log f(y\|M_k, \theta_k^*)$	$\log \pi(\theta_k^*\|M_k)$	$\pi(\theta_k^*\|y, M_k)$	Total
UC	-629.80	1.529	-18.298	-609.98
UC-SVt	-471.65	-10.396	9.612	-491.67
UC-SVc	-400.39	-14.556	11.777	-426.72
UC-SV	-372.21	-44.394	33.328	-449.94

model with stochastic volatility in the core component. The variation in the transitory one is by and large insignificant. The UC-SV has the highest conditional likelihood, but receives a high 'penalty" from the term $\log \pi(\theta_k|M_k)$. As a result the posterior odds of model UC-SV against UC-SVc are close to zero. Hence, we conclude that the model with two stochastic volatility components is likely to overfit the data.

Thus, our preferred model is the UC-SVc specification; table 3 and figures 4-6 report the main estimation results for this model. In particular, figure 4 displays the posterior mean of the core component and the mean, median, and 2.5% and 97.5% percentiles of the posterior distribution of the the volatility of the core component.

Table 3: Posterior, Median, Geweke statistic and Inefficinecy factor for LLM with SV in the core component

Parameters	Mean	Median	Geweke statistic	Inefficinecy factor (taper 0.05)
μ_2	-0.0199	-0.0196	1.1185	1.67
ϕ_2	0.9806	0.9817	-1.28412	16.77
$\sigma^2_{\kappa_2}$	0.0410	0.0397	0.8093	14.73
σ^2_{ϵ}	0.5904	0.5911	0.8762	43.54

6 Conclusion

Using a model that provides a simple yet effective decomposition of U.S. inflation into a core component and a transitory one, with stochastic volatility in the disturbances driving the two components, Bayesian model selection enabled us to conclude that inflation's volatility is subject to significant changes over time, but the volatility affects only the core disturbances, not the transitory component.

The volatility of the core has been decreasing substantially after 1982, reaching a very low level during the 1990s, but has been subject to an increase since the end of the 1990s. The estimated volatility pattern support the view that a turning point took place and the great moderation is over. The persistence implied by the model has been decreasing during the years of the great moderation and it stayed at historical lows in the mid nineties. Recently, persistence has been increasing again. Correspondingly, the predictability of inflation increased during the great moderation up to maximum that took place in the mid 1990s and has been going down ever since.

7 APPENDIX A The Metropolis- within- Gibbs sampling scheme

This Appendix illustrates the prior and posterior distributions used in our analysis. For the prior distribution we assume an independent structure between each block of variables and and within each block so that $\pi(\theta, \alpha, h_1, h_2) = \pi(\theta)\pi(\alpha)\pi(h_1)\pi(h_2)$, and, for instance,

$$\pi(\theta) = \pi(\mu_1|c_1, d_1)\pi(\mu_2|c_2, d_2)\pi(\phi_1|a_1, b_1)\pi(\phi_2|a_2, b_2)\pi(\sigma^2_{\kappa_1}|\gamma_1, \beta_1)\pi(\sigma^2_{\kappa_2}|\gamma_2, \beta_2).$$

The prior distributions and their hyperparameters are reported in table 4.

The posterior densities are available in closed form for the core level of inflation (for which

Table 4: Specification of the prior distributions

θ	Prior	Hyperparameters	
μ_i	$N(c_i, d_i^2)$	$c_i = 0.00$	$d_i = 10.00$
ϕ_i	$Beta(a_i, b_i)$	$a_i = 20.50$	$b_i = 1.50$
$\sigma^2_{\kappa_1}$	$IG(\gamma_1, \beta_1)$	$\gamma_1 = 20$	$\beta_1 = 0.20$
$\sigma^2_{\kappa_2}$	$IG(\gamma_2, \beta_2)$	$\gamma_2 = 20$	$\beta_2 = 0.20$

samples are drawn by a multimove sampler known as the simulation smoother, here implemented according to the algorithms presented in Durbin and Koopman (2002)), and for some elements of the vector θ for which we can exploit conditional conjugacy.

1. Given the choice of the prior distribution, the full conditional density of the parameter ϕ_1 (and similarly ϕ_2) is not available in closed form; therefore, to sample from the full conditional we employ a random walk Metropolis-Hastings sampling algorithm, which has the merit of enforcing the stationarity of the stochastic volatility process. If $\phi_i^{(j-1)}$ denotes the current value of the chain at the j-th iteration, we sample a new proposal $\phi_i^{(j)} = \phi_i^{(j-1)} + w_j$, where w_j is drawn a normal distribution with mean 0 and variance 1. If the proposal is within the stationary region then it is accepted with probability $\min\{1, g(\phi_i^{(j)})/g(\phi_i^{(j-1)})$ where

$$g(\phi_i) = \pi(\phi_i) f(h_i|\mu_i, \phi_i, \sigma^2_{\kappa_i})$$

   and, apart from a constant term,

$$\log f(h_i|\mu_i, \phi_i, \sigma^2_{\kappa_i}) = -\frac{h^2_{i,0}}{2\sigma^2_{\kappa_i}} + \frac{1}{2}\log(1-\phi_i^2) - \frac{\sum_{t=1}^{n-1}(h_{i,t+1} - \phi_i h_{i,t} - \mu_i)^2}{2\sigma^2_{\kappa_i}}. \tag{8}$$

   Different sampling schemes, illustrated in Kim, Shephard and Chib (1998), were also adopted for comparison, but the results were unaffected.

2. Using a Normal prior, the full conditional distribution of the parameters μ_i is $N(\hat{C}_i, \hat{D}_i)$ where:

$$\hat{C}_i = \hat{D}_i \left(\frac{C_i}{D_i^2} + \frac{1}{\sigma^2_{\kappa_i}} \sum_{t=1}^{T} (h_{i,t} - \phi_i h_{i,t-1}) \right) \qquad \hat{D}_i = \left(\frac{1}{d_i^2} + \frac{T}{\sigma^2_{\kappa_i}} \right)^{-1} \tag{9}$$

3. Using a conjugate Inverse Gamma prior, the full conditional of the variances of volatility processes are:

$$\sigma^2_{\kappa_i}|y,\alpha,h_i,\phi_i,\mu_i \sim IG\left\{\frac{n}{2}+\alpha_i,\beta_i+\frac{h^2_{i,0}+\sum_{t=1}^{n-1}(h_{i,t+1}-\mu_i-\phi_i h_{i,t})^2}{2}\right\}$$

4. To sample from $h_{1t}|h_{1,t-1},h_{1,t+1},y_t,\alpha_t,\theta$, we adopt the single move Metropolis-Hastings simulation step, based on the factorization:

$$f(h_{1t}|h_{1,t-1},h_{1,t+1},y_t,\alpha_t,\theta) \propto f(h_{1t}|h_{1,t-1},h_{1,t+1},\theta)f(y_t|\alpha_t,h_{1t}). \tag{10}$$

It can be shown that

$$f(h_{1t}|h_{1,t-1},h_{1,t+1},\theta) = f(h_{1t}|h_{1,t-1},\theta)f(h_{1,t+1}|h_{1t},\theta) \tag{11}$$

is a Gassian density with mean

$$h^*_{1t} = \frac{\mu(1-\phi)+\phi(h_{1,t-1}+h_{1,t+1})}{(1+\phi_i^2)}$$

and variance

$$v_i^2 = \frac{\sigma^2_{\kappa_i}}{1+\phi_i^2}$$

(see Jacquier, Polson and Rossi, 1994). Independent proposals $h_{1t}^{(j)}$ can be made from this Gaussian density; their acceptance probability is $\min\{1, g(h_{1t}^{(j)})/g(h_{1t}^{(j-1)})\}$, where

$$g(h_{1t}) = \exp\left[-\left\{\frac{(h_{1,t+1}-\mu_1-\phi_1 h_{1,t})^2}{2\sigma^2_{\kappa_1}}+\frac{(h_{1,t}-\mu_1-\phi_1 h_{1,t-1})^2}{2\sigma^2_{\kappa_1}}\right\}\right]\frac{1}{\exp(h_1/2)}\exp\left[-\frac{(y_t-\alpha_t)^2}{2\exp(h_1)}\right] \tag{12}$$

for $t=1,\ldots,n$, whereas

$$g(h_{1,0}) = \exp\left\{-\frac{(h_{1,1}-\mu_1-\phi_1 h_{1,0})^2}{2\sigma^2_{\kappa_1}}-\frac{(1-\phi_1^2)h^2_{1,0}}{2\sigma^2_{\kappa_1}}\right\}$$

and for $t=n$

$$g(h_{1,n}) = \exp\left\{-\frac{(h_{1,n}-\mu_1-\phi_1 h_{1,n-1})^2}{2\sigma^2_{\kappa_1}}\right\}$$

A similar sampling scheme is adopted for h_2.

8 APPENDIX B The Chib and Jeliazkov algorithm

This Appendix illustrates the steps of the Chib and Jeliazkov (2001) algorithm that are necessary to estimate the posterior density $\pi(\theta|y)$ for the UC-SV model at a high density point θ^*. The latter is the component of the basic marginal likelihood identity that is not automatically available from the MCMC output.

The estimate

$$\hat{\pi}(\theta^*|y) = \prod_{k=1}^{K} \hat{\pi}(\theta_k^*|y, \theta_1^*, \dots, \theta_{k-1}^*)$$

where the elements of the vector θ are $\{\mu_1, \phi_1, \sigma^2_{\kappa_1}, \mu_2, \phi_2, \sigma^2_{\kappa_1}\}$.

Let $z = (h_1, h_2, \alpha)$. The algorithm goes as follows:

- From the MCMC sample evaluate the posterior mean of μ_1 and set μ_1^* equal to this value. A Monte Carlo estimate of the first multiplicative factor, $\pi(\theta_1^*|y) = \pi(\mu_1^*|y)$, is obtained from the output of the MCMC sampling scheme by the technique known as Rao-Blackwellization.

- For estimating $\pi(\theta_2^*|y, \theta_1^*) = \pi(\phi_1^*|y, \mu_1^*)$ run a reduced Metropolis-Hastings within Gibbs chain for the following subset of parameters $\{\phi_1, \sigma^2_{\kappa_1}, \mu_2, \phi_2, \sigma^2_{\kappa_2}, z\}$, where the value of μ_1 is kept fixed at μ_1^*.

- Estimate the value of the density $\pi(\theta_2^*|y, \theta_1^*) = \pi(\phi_1^*|y, \mu_1^*)$, using the following steps:

  1. Simulate G draws from the posterior of $\{\phi_1^{(g)}, \sigma^{2,(g)}_{\kappa_1}, \mu_2^{(g)}, \phi_2^{(g)}, \sigma^{2,(g)}_{\kappa_2}, z^{(g)}\}, g = 1, \dots, G$, by the same MCMC methods presented in appendix A, conditional on μ_1^*.
  2. Compute the posterior mean of ϕ_1 by averaging across the draws $\phi_1^{(g)}$ and denote it ϕ_1^*.
  3. Include ϕ_1^* in the conditioning set and sample J draws from the conditional distibutions:

$$\begin{aligned}
&\pi(\sigma^2_{\kappa_1}|y, z, \phi_1^*, \mu_1^*, \mu_2, \sigma^2_{\kappa_2}, \phi_2), \qquad \pi(z|y, \sigma^2_{\kappa_1}, \mu_1^*, \phi_1^*, \mu_2, \phi_2, \sigma^2_{\kappa_2}),\\
&\pi(\mu_2|y, z, \mu_1^*, \phi_1^*, \sigma^2_{\kappa_1}, \phi_2, \sigma^2_{\kappa_2}), \qquad \pi(\phi_2|y, z, \mu_1^*, \phi_1^*, \sigma^2_{\kappa_1}, \mu_2, \sigma^2_{\kappa_2}),\\
&\pi(\sigma^2_{\kappa_2}|y, z, \mu_1^*, \phi_1^*, \sigma^2_{\kappa_1}, \mu_2, \phi_2).
\end{aligned}$$

  These iterations provide the sample $\{\sigma^{2(j)}_{\kappa_1}, \mu_2^{(j)}, \phi_2^{(j)}, \sigma^{2(j)}_{\kappa_2}, z^{(j)}\}_{j=1}^{J}$. Furthermore, at each iteration we generate

$$\phi_1^{(j)} \sim q(\phi_1^*, \phi_1|y, z^{(j)}, \mu_1^*, \sigma^{2,(j)}_{\kappa_1}, \mu_2^{(j)}, \phi_2^{(j)}, \sigma^{2,(j)}_{\kappa_2})$$

where $q(\theta_j, \theta'_j|u)$ is the proposal density for the transition from θ_j to θ'_j conditional on u. As a result, the collection $\{\phi_1^{(j)}, \sigma_{\kappa_1}^{2(j)}, \mu_2^{(j)}, \phi_2^{(j)}, \sigma_{\kappa_2}^{2(j)}, z^{(j)}\}_{j=1}^J$ is are multiple (correlated) draws from the distribution:

$$\pi(\sigma_{\kappa_1}^2, \mu_2, \phi_2, \sigma_{\kappa_2}^2, z|y, \mu_1^*, \phi_1^*) \times q(\phi_1^*, \phi_1|y, z, \mu_1, \sigma_{\kappa_1}^2, \mu_2, \phi_2, \sigma_{\kappa_2}^2).$$

4. Denoting the probability of a move by

$$\alpha(\phi_1, \phi'_1|u) = \min\left\{1, \frac{f(y|\phi_1^*, \varsigma, z)\pi(\phi_1^*, \varsigma)}{f(y|\phi_1^{(g)}, \varsigma, z)\pi(\phi_1^{(g)}, \varsigma)} \frac{q(\phi_1^*, \phi_1^{(g)}|y, \varsigma, z)}{q(\phi_1^{(g)}, \phi_1^*|y, \varsigma, z)}\right\},$$

where ς is the collection of parameters $(\mu_1^*, \sigma_{\kappa_1}^2, \mu_2, \phi_2, \sigma_{\kappa_2}^2)$. The required marginal density at ϕ_1^* can now be estimated as

$$\hat{\pi}(\phi_1^*|y) = \frac{G^{-1}\sum_g \alpha(\phi_1^{(g)}, \phi_1^*|y, z^{(g)}, \mu_1^*, \sigma_{\kappa_1}^{2(g)}, \mu_2^{(g)}, \phi_2^{(g)}, \sigma_{\kappa_2}^{2(g)}) \times q(\phi_1^{(g)}, \phi_1^*|y, z^{(g)}, \mu_1^*, \sigma_{\kappa_1}^{2(g)}, \mu_2^{(g)}, \phi_2^{(g)}, \sigma_{\kappa_2}^{2(g)})}{J^{-1}\sum_j \alpha(\phi_1^*, \phi_1^{(j)}|y, z^{(j)}, \mu_1^*, \sigma_{\kappa_1}^{2(j)}, \mu_2^{(j)}, \phi_2^{(j)}, \sigma_{\kappa_2}^{2(j)})}$$

- Run a reduced Gibbs on the following parameters $\{\sigma_{\kappa_1}^2, \mu_2, \phi_2, \sigma_{\kappa_2}^2, z\}$ and calculate $\sigma_{\kappa_1}^{2,(*)}$
- Run a reduced Gibbs and calculate the ϕ_2^* with the same procedure describe before noticing that the $\phi_1^*, \mu_1^*, \sigma_{\kappa_1}^{2,(*)}$ are fixed.
- Run a reduced Gibbs on the following parameters $\{\mu_2, \sigma_{\kappa_2}^2, z\}$ and calculate μ_2^*;
- Run a reduced Gibbs on the following parameters $\{\sigma_{\kappa_2}^2, z\}$ and calculate $\sigma_{\kappa_2}^{2,(*)}$

References

[1] Bos, C. S., and Shephard, N. (2006). *Inference for Adaptive Time Series Models: Stochastic Volatility and Conditionally Guassian State Space Form*, Econometric Reviews, 25, 219–244.

[2] Bos, C. S., Koopman, S.J. and Ooms, M. (2008). *Long Memory Modelling of Inflation with Stochastic Variance and Structural Breaks*, Discussion paper TI 07–099/4, Tinbergen Institute.

[3] Brot, C. and Ruiz, E., (2008) *Testing for conditional heteroskedastic in the components of inflation*, Banco de España, WP n. 0812.

[4] Doornik, J.A (2007). *Ox: An Object-Oriented Matrix Programming Language*, Timberlake Consultants Press, London.

[5] Cappé, O., Moulines, E. and Ryden, T. (2007). *Inference in Hidden Markow Models*, Springer

[6] Cecchetti, G., Hooper, P., Kasman, B.C., Shoenholtz, K. and Watson, M. (2007). *Understanding the evolving Inflation Process*, U.S. Monetary Policy Forum.

[7] Chib, S. (1995). *Marginal Likelihood from the Gibbs Output*, Journal of the American Statistical Association Vol. 90, 1313–1321.

[8] Chib, S. and Jeliazkov, I. (2001). *Marginal Likelihood from the Metropolis-Hastings Output*, Journal of the American Statistical Association, 96, 270–281.

[9] Cogley, T., Primicieri, G. and Sargent, T.J. (2008). *Inflation-Gap persistence in the U.S.*, NBER Working Paper No. 13749.

[10] Durbin, J. and Koopman, S.J. (2002). *A simple and efficient simulation smoother for state space time series analysis*, Biometrika, 89, 603-616

[11] Geweke, J.F. (2005). *Contemporary Bayesian econometrics and statistics*, Wiley, New York.

[12] Kim, S., Shepard, N. and Chib, S. (1998). *Stochastic Volatility: Likelihood Inference and Comparison with ARCH model*. Review of Economic Studies, 65, 361-393.

[13] Liu, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. Springer.

[14] Pivetta, F. and R. Reis (2007). *The persistence of inflation in the United States*, Journal of Economic Dynamics and Control, 31, 1326–1358

[15] Stock, J.H., and Watson, M. (2007). *Why Has U.S. Inflation Become Harder to Forecast?*, Journal of Money, Credit and Banking, 39, 3–34.

[16] West, M. and Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models (2 ed.)*. New York: Springer-Verlag.

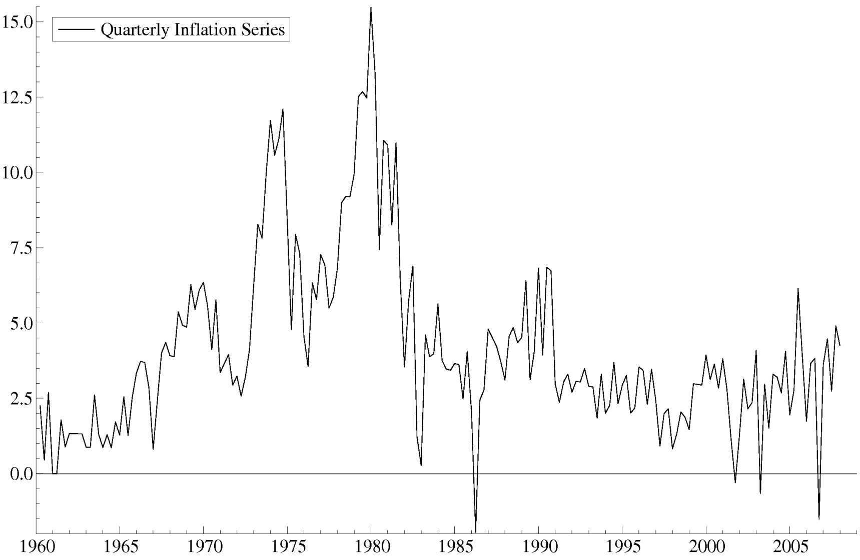

Figure 1: Quarterly Inflation series

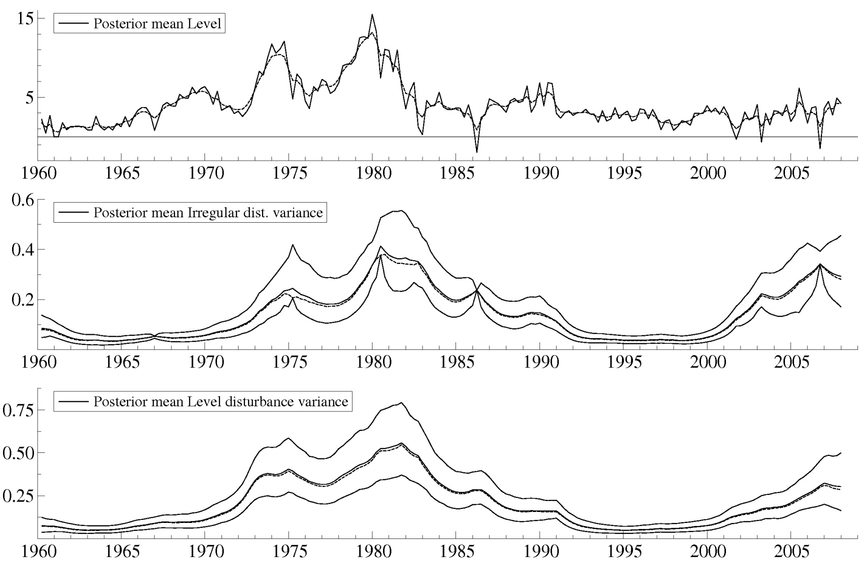

Figure 2: Upper: Inflation and the posterior mean component; Middle: Irregular Volatility component with confidence interval; Bottom: Core Volatility component with confidence interval

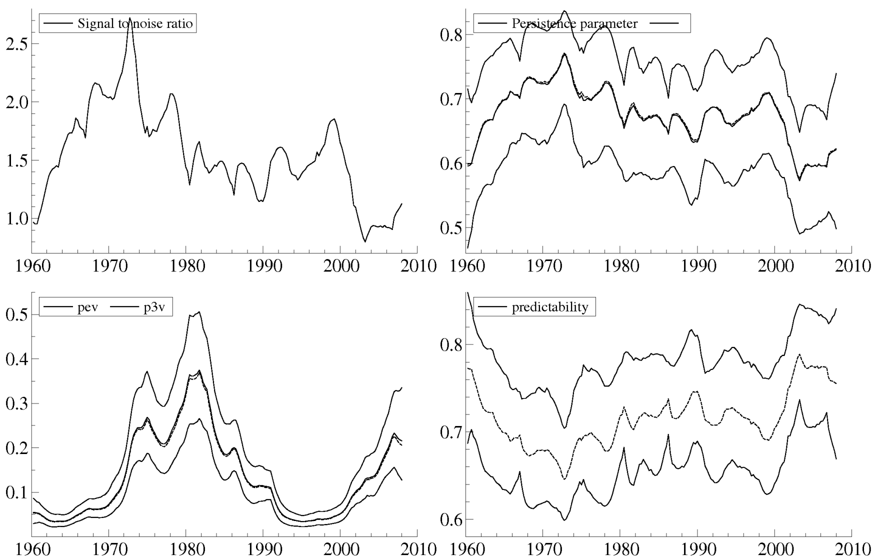

Figure 3: Upper left: Signal to noise ratio; Upper Right: Persistence Parameter; Bottom left: Prediction error variance; Bottom right: Predictability

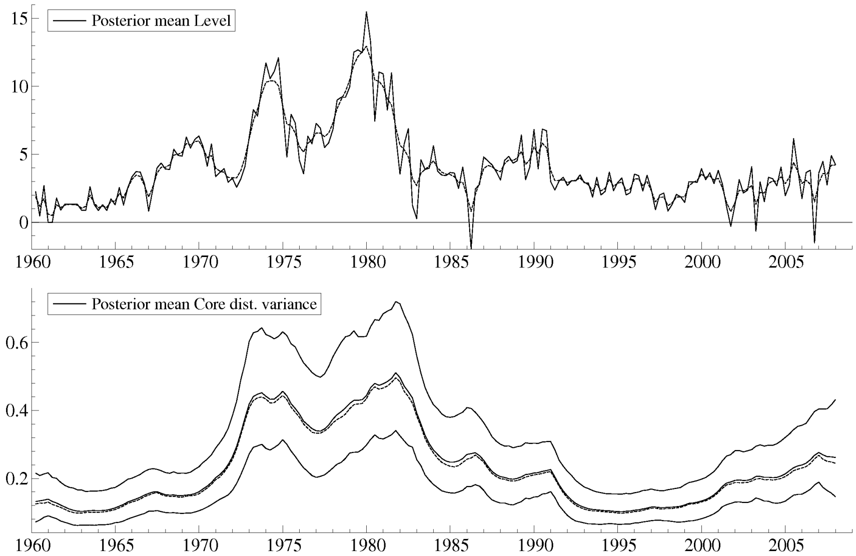

Figure 4: Upper: Quarterly inflation and its posterior mean level; Buttom: Volatility of the core component

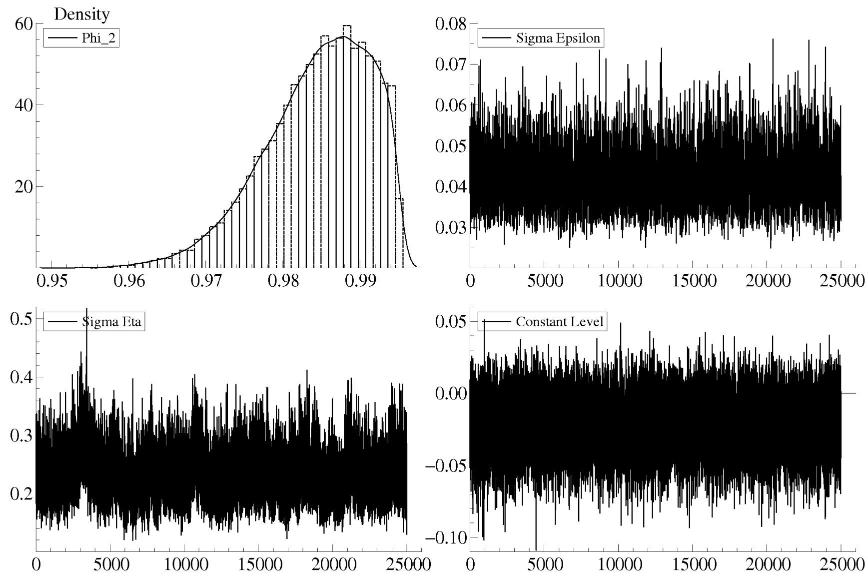

Figure 5: Posterior of the Volatility of volatility for the core component; Simulation against iterations for a LLM with SV in the core component.

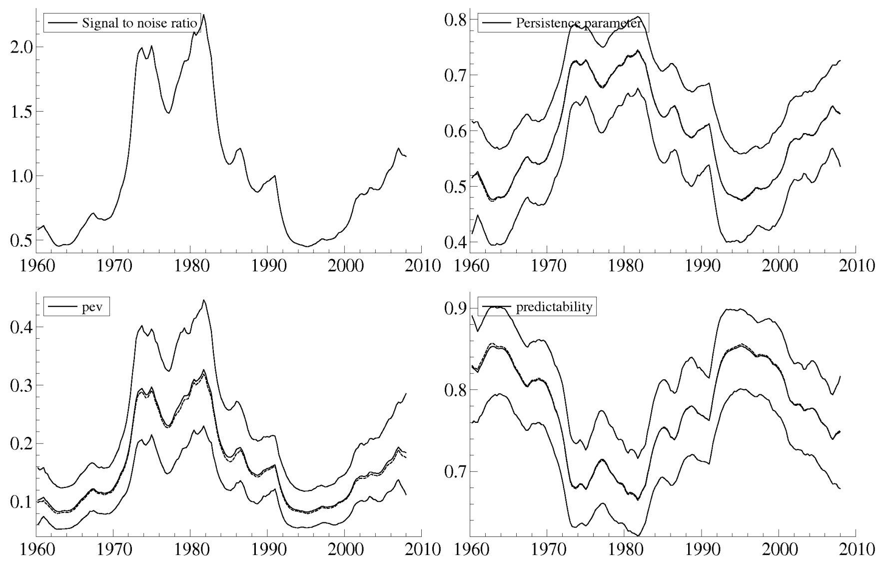

Figure 6: Upper left: Signal to noise ratio ; Upper Right: Persistence Parameter; Bottom left: Prediction error variance; Bottom right: Predictability

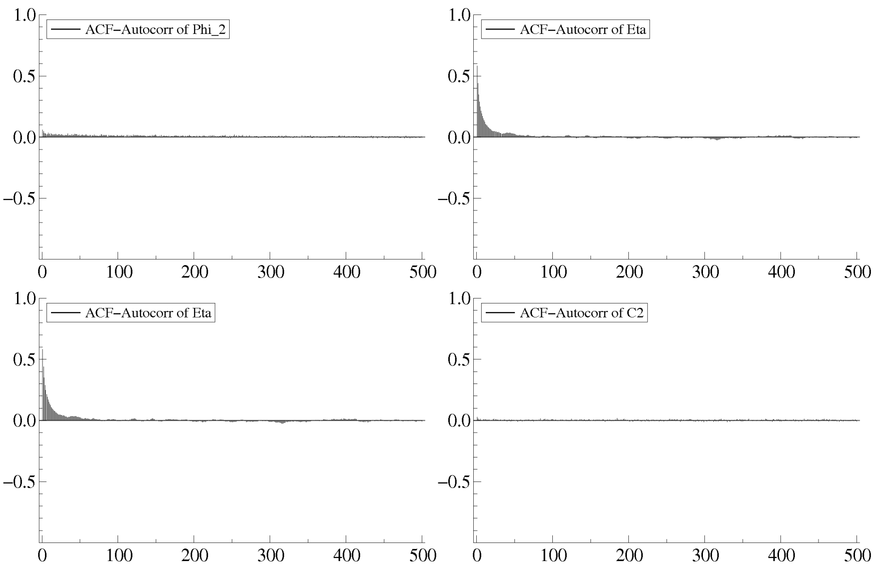

Figure 7: Upper left: Autocorrelation of ϕ_2; Upper right: Autocorrelation of σ_η^2; Buttom left: Autocorrelation of σ_ϵ^2; Buttom right: Autocorrelation of μ_2.