Journal of Time Series Econometrics

Volume 2, *Issue* 1 | 2010 | *Article* 6

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

Stefano Grassi, *Università degli Studi di Roma Tor Vergata*
Tommaso Proietti, *Università degli Studi di Roma Tor Vergata*

Recommended Citation:
Grassi, Stefano and Proietti, Tommaso (2010) "Has the Volatility of U.S. Inflation Changed and How?," *Journal of Time Series Econometrics*: Vol. 2: Iss. 1, Article 6.
DOI: 10.2202/1941-1928.1050

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

Stefano Grassi and Tommaso Proietti

Abstract

The local level model with stochastic volatility, recently proposed for U.S. Inflation 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 permanent 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 permanent 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 around 1995. During the last decade, volatility and persistence have been increasing and predictability has been going down.

KEYWORDS: marginal likelihood, Bayesian model comparison, auxiliary particle filter, stochastic volatility, great moderation, inflation persistence

Author Notes: This paper was presented at the 3rd Italian Congress of Econometrics and Empirical Economics, University of Ancona, January 30-31, 2009, and at the EEA-ESEM meeting in Barcelona, August 23-27, 2009. The authors wish to thank Charles Bos and an anonymous referee for their suggestions that helped to improve the content of the paper.

1 Introduction

Inflation's volatility has attracted a great 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 (2007), Cecchetti, Hooper, Kasman, Shoenholtz, and Watson (2007), among others. Recently Bos, Koopman, and Ooms (2007) 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 volatility of the series and the fractional integration parameter (which is the measure of persistence adopted in that paper).

In this paper we reconsider the unobserved components model of U.S. inflation estimated in Stock and Watson (2007), referred to as the local level model with stochastic volatility (UC-SV). The model provides a simple 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 permanent 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 permanent component, but not in both. We

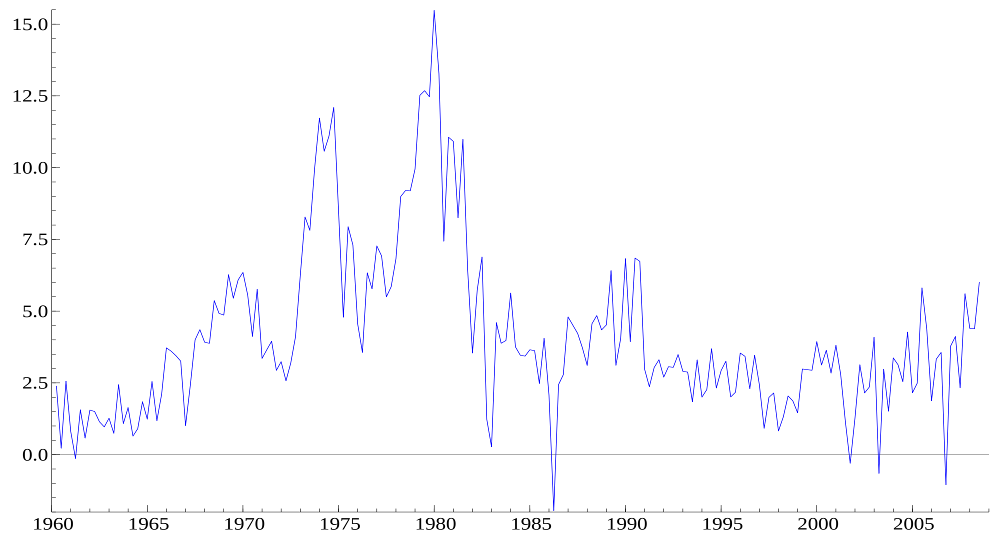

Figure 1: Quarterly U.S. Inflation, $y_t = 400\Delta \ln \text{CPI}_t$

document that persistence is higher than in previous studies and is subject to a significant increase starting from the second half of the 90's, whereas predictability has decreased 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. Section 6 concludes 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 series by CPI_t, the annualized quarterly inflation rate, denoted y_t, $t = 1, \ldots, T$, is computed as $y_t = 400\Delta \ln \text{CPI}_t$. The series is plotted in figure 1 and is available for the sample period 1960:q1 –2008:q3.

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:

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

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 $h_{1,t} = \ln \sigma^2_{\varepsilon t}$ and $h_{2,t} = \ln \sigma^2_{\eta t}$,

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

The model encompasses the traditional stochastic volatility model that is widely used in finance (see for instance Shepard, 2006), which arises when the process α_t degenerates to a constant.

The specification of the stochastic volatility processes differ only slightly from Stock and Watson (2007) and Cecchetti, Hooper, Kasman, Shoenholtz, and Watson (2007), who assume a random walk process for the log-variances $h_{i,t}, i = 1,2$. In fact, their specification is encompassed by (2), which is a more canonical specification of a volatility model (see for instance Jacquier, Polson, and Rossi, 1994, and Kim, Shepard, and Chib, 1998), since it arises as the discrete-time approximation to the Ornstein-Uhlenbeck continuous time process used in finance, and ensures the stationarity of η_t and ε_t, provided that $|\phi_i| < 1, i = 1,2$. As a matter of fact, when the autoregressive coefficients ϕ_i are close to unity and the constants $\mu_i, i = 1,2$, are close to zero, specification (2) is virtually indistinguishable from a random walk.

When both variances $\sigma^2_{\varepsilon t}$ and $\sigma^2_{\eta t}$ 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 $\xi_t \sim \text{NID}(0, \sigma^2)$. The structural parameters are related to the reduced form parameters by the two equations $\sigma^2(1+\vartheta^2) = \sigma^2_\eta + 2\sigma^2_\varepsilon$, $\sigma^2\vartheta = -\sigma^2_\varepsilon$, which are obtained by equating the autocovariances at lags 0 and 1, respectively; from these we obtain the moving average parameter $\vartheta = \left[(q^2+4q)^{\frac{1}{2}} - 2 - q\right]/2$, where $q = \sigma^2_\eta/\sigma^2_\varepsilon$ is the signal to noise ratio, and the prediction error variance (p.e.v.), $\sigma^2 = -\sigma^2_\varepsilon/\vartheta$. Notice that ϑ is restricted within the range [-1,0]. 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 $\sigma^2_{\varepsilon t}$ or $\sigma^2_{\eta t}$, or both, are constant.

The UC-SV model can be considered as an IMA(1,1) model with time-varying p.e.v. and moving average parameter. This suggests taking, as a local measure of persistence, $P_t = 1 + \vartheta_t$, where ϑ_t varies with time according to the values of the time-varying signal to noise ratio $q_t = \sigma^2_{\eta t}/\sigma^2_{\varepsilon t}$. The quantity P_t decreases linearly from 0 to 1 as ϑ increases from -1 to 0. Cecchetti, Hooper, Kasman, Shoenholtz, and Watson (2007) use the implied time varying first order autocorrelation of Δy_t, as a measure of persistence; the local autocorrelation (i.e. conditional on $\sigma^2_{\eta t}$ and $\sigma^2_{\varepsilon t}$) is $\rho_t(1) = -1/(q_t + 2) = \vartheta_t/(1+\vartheta_t^2)$. Alternatively, we can use the (conditional) normalized spectral generating function at the zero frequency, which is

$$P_t^* = \frac{\sigma^2_{\eta t}}{\sigma^2_{\eta t} + 2\sigma^2_{\varepsilon t}} = 1 + \frac{2\vartheta_t}{1+\vartheta_t^2} = 1 + 2\rho_t(1).$$

This measure decreases monotonically from 0 to 1 as ϑ increases from -1 to 0.

As a measure of local predictability we can take the prediction error variance, conditional on $\{h_{i,t}, i = 1,2, t = 1,\dots,T\}$, which is defined as

$$\sigma_t^2 = -\frac{\sigma^2_{\varepsilon t}}{\vartheta_t}.$$

A relative measure of predictability can be defined in terms of the (Granger and Newbold, 1986, p. 310) forecastability index:

$$\text{Pred}_t = 1 - \frac{\text{Var}(\xi_t|h_{i,t})}{\text{Var}(\Delta y_t|h_{i,t})} = \frac{\vartheta_t^2}{1+\vartheta_t^2}. \tag{3}$$

In terms of the parameters of the UC-SV, the prediction error variance equals $\text{Var}(\xi_t|h_{i,t}) = \frac{\sigma^2_{\eta t}}{(1+\vartheta)^2}$, whereas the variance $\text{Var}(\Delta y_t|h_{i,t}) = \sigma^2_{\eta t} + 2\sigma^2_{\varepsilon t}$. The above measure ranges between 0 ($\vartheta_t = 0$) and 0.5 ($\vartheta_t = -1$), and it is negatively related to the persistence of the process. As a matter of fact, as ϑ_t ranges from -1 to 0, predictability decreases from its maximum, 0.5, to zero.

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,

DOI: 10.2202/1941-1928.1050

according to which samples are drawn componentwise from the full conditionals; for the components which cannot be sampled directly a Metropolis-Hasting subchain is used within the Gibbs sampling cycle. In particular, the posterior of the AR parameters, ϕ_1 and ϕ_2, is not available in closed form; see Kim, Shepard, and Chib (1998) and Bos and Shephard (2006). 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, and α and y denote the stack of the values of permanent inflation and the series, respectively. 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 proposed by Durbin and Koopman (2002).
3. Sample $h_i, i = 1, 2$, from $h_i | \alpha, y, \theta$, using a Random Walk Metropolis-Hastings algorithm.
4. Go to 2.

The most complex part of the algorithm deals with the simulation of the stochastic volatility processes. We adopt a single move sampler based on the density:

$$h_{i,t} | h_{i,t+1}, h_{i,t-1}, y_t, \alpha_{t-1}, \alpha_t. \quad (4)$$

For this purpose, we implement a Random Walk Metropolis-Hastings algorithm, described in detail in Appendix A; see also Cappé, Moulines, and Rydén (2007). In order to sample from the full conditional we use the following results:

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

4 Estimation Results

This section reports on the main estimation results for the model presented in section 2. The MCMC sampler was initialized by setting all $h_{i,t} = 0$ and $\phi_i = 0.86$, $\sigma^2_i = 0.07$ and $\mu_i = 0.6$. We iterated the sampler for a burn-in period consisting of 12,500 iterations, before recording the draws from a subsequent 25,000 iterations. The programm is written in Ox v. 5.10 console (Doornik, 2007) using our own 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 displays the inflation series with the posterior mean of the permanent component, and the interval estimates of two stochastic volatility components for the irregular and the permanent disturbances, $\sigma_{\varepsilon t}$ and $\sigma_{\eta t}$. The third panel shows that the volatility of the permanent component has been increasing from the 60ies up to 1982, and then is slowly decreasing. The volatility of the transitory component (central panel) is much more stable, instead. In the sequel of the paper we will address the question as to whether it can be considered as time invariant.

The estimates of the latent volatility processes are comparable to the corrected estimates obtained by Stock and Watson (2007) and displayed in their Figure 2, referring to CPI (all items), on page 8, panels (a) and (b), of the document available at *http://www.princeton.edu/∼mwatson*. In particular, the estimated standard deviation of the permanent component shows two distinctive peaks in 1975 and 1981, and changes substantially over time; on the contrary, the volatility of the irregular component is much less evolutive. The difference that arise are due to the different sample considered and to the fact that Stock and Watson estimate a restricted version of the model (in particular, $\phi_i = 1, \mu_i = 0, \sigma^2_{\kappa i} = 0.2, i = 1,2$; notice that the variance of the volatility shocks is not estimated).

Figure 3 displays the evolution of the Monte Carlo estimates of the posterior mean of the signal to noise ratio, q_t, of the persistence parameter, P_t, the prediction error variance and the predictability measure, Pred_t. The graph reveals that the size of the random walk component increases during the 70s, when the trend dynamics become more sustained, and it is lower in the 80s. Persistence is time varying at values well below 1 and there is evidence for a strong decreasing tendency during the 80s. 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, 2007). In relative terms, the forecastability index shows an increase in the 80s.

Table 1 reports some summary statistics concerning the posterior distribution of the parameters and some convergence diagnostics. As for the assessment of convergence, we report the Geweke statistics: let $\theta^{(j)}$ denote the j-th sample of the sampling scheme for the generic parameter θ after the initial burn-in period. Let also $\bar{\theta}_a$ denote the average of the first n_a draws, $\bar{\theta}_b$ the average of the last n_b draws at the end of the convergence period, which are taken sufficiently remote to prevent any overlap, the Geweke's convergence statistic (Geweke, 1992, 2005) is defined as

$$C_G = \frac{\bar{\theta}_a - \bar{\theta}_b}{\sqrt{V_{L,a}/n_a + V_{L,b}/n_b}},$$

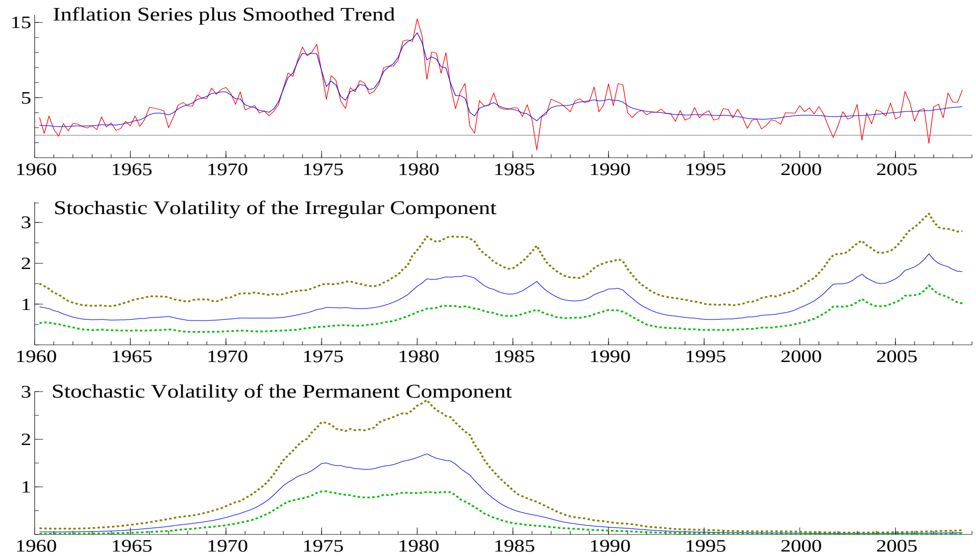

Figure 2: Upper: Inflation and posterior mean of permanent component; Middle: Irregular Volatility component with 95 percent credible interval. Bottom: Permanent Volatility component with 95 percent credible interval.

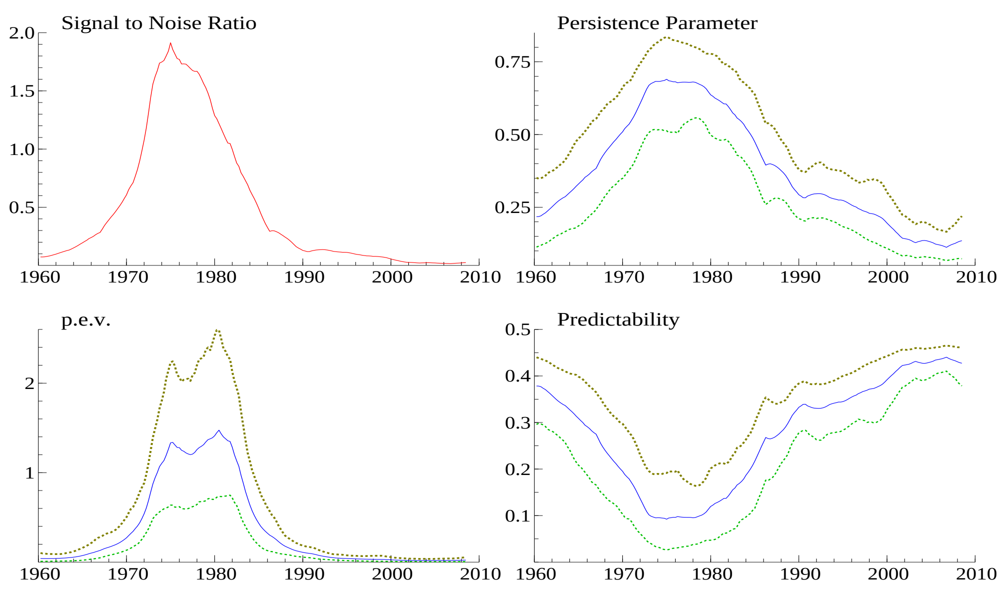

Figure 3: Upper left: Signal to noise ratio. Upper Right: Persistence, with 95 percent credible interval. Bottom left: Prediction error variance, with 95 percent credible interval. Bottom right: Predictability, with 95 percent credible interval.

where

$$V_{L,k} = c_{0,k} + 2\sum_{j=1}^{n_k-1} w_j c_{j,k}, \quad k = a, b,$$

is the long run variance of the parameter sample path for the n_k draws, based on a weighted combination of the autocovariances of the draws at lag j, $c_{j,k}$, with weights w_j that are decreasing in j and ensure that $V_{L,k} \geq 0$. A customary choice is the set of linearly declining weights $w_j = \frac{l-j}{l+1}$, where l is the truncation parameter.

The inefficiency factor is $\text{INEF} = 1 + 2\sum_{j=1}^{n-1} w_j \hat{\rho}_j$, where $\hat{\rho}_j$ is the sample autocorrelation of the draws at lag j. This can be interpreted as a normalized measure of persistence of the draws. Large values imply that the draws are strongly and positively autocorrelated (the spectral power is concentrated at the origin), so that the chain explores the parameter space very slowly and the additional information content of a draw is small.

The values reported in table 1 highlight that the convergence assessment of the chain are not fully satisfactory, since the Geweke statistic for some parameters, like μ_1 and ϕ_2, are strongly significant.

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

Parameters	Mean	Median	Geweke' G_C	INEF
μ_1	-0.0017	-0.0015	-11.20	137.1
μ_2	-0.0253	-0.0252	-2.30	30.92
ϕ_1	0.9356	0.9372	1.18	13.00
ϕ_2	0.9885	0.9905	15.06	307.7
$\sigma^2_{\kappa_1}$	0.0491	0.0482	-2.38	24.43
$\sigma^2_{\kappa_2}$	0.0487	0.0479	-0.18	63.41

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 estimation result. There is a potential danger that the UC-SV model could be overfitting the data, but little or no attention has been devoted to this issue.

We thus turn our attention to Bayesian model selection. The models under comparison are the following four variants of the local level model:

- M_1: the Local Level Model with homoscedastic 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 permanent 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 $f(y|\theta_k, M_k)$ the conditional density of the data, given M_k and the parameter vector θ_k, and by $\pi(\theta_k|M_k)$, $\pi(\theta_k|y, M_k)$, the prior and posterior densities, respectively, of θ_k, the Chib and Jeliazkov (2001) approach is based on the following basic marginal likelihood identity:

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

where $m(y|M_k)$ is the marginal likelihood of model M_k.

The formal Bayesian approach for comparing any two rival specifications, M_k and M_r, is through the pairwise Bayes factor, defined as the ratio of marginal likelihoods:

$$B_{k,r} = \frac{m(y|M_k)}{m(y|M_r)},$$

which can also be interpreted as the posterior odds ratio the two models, when they are assumed to be equally likely a priori.

Taking logarithms of (6) and evaluating this function at some hight density point θ_k^*, such as the mean of the posterior density $\pi(\theta_k|y, M_k)$, 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 conditional likelihood appearing as the first term on the right hand side is evaluated with the support of the Kalman filter for the linear Gaussian homoscedastic local level model (M_1); for the other specifications, featuring stochastic volatility in at least one of the components, it is evaluated by sequential Monte Carlo methods (particle filters). Full details are provided in Appendix B. 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 C we provide the relevant details for its estimation, with particular reference to UC-SV specification.

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

Models	$\log f(y\|M_k, \theta_k^*)$	$\log \pi(\theta_k^*\|M_k)$	$\log \pi(\theta_k^*\|y, M_k)$	Total
UC	-369.56	-11.48	-0.12	-380.93
UC-SVt	-367.80	-8.83	7.24	-383.87
UC-SVc	-366.71	-2.51	-13.5	-355.72
UC-SV	-356.10	-3.06	20.81	-379.96

The results, reproduced in table 2, clearly point out that the model that performs best is the local level model with stochastic volatility in the permanent 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 $\log \pi(\theta_k^*|y, M_k)$. 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 over-fit the data.

Hence, our preferred model is the UC-SVc specification; table 3 and figures 4-5 report the main estimation results for this model. In particular, figure 4 displays the posterior mean of the permanent component, along with the 95% credible intervals. The bottom panel displays the posterior mean and the interval estimates of the process $\sigma_{\eta t}$. The plot illustrates that the volatility of the permanent component is subject to a steep decline in the years 1982-1995, whereas the trend is reversed after 1995. The first panel of figure 5 displays the evolution of the posterior mean of the signal to noise ratio, $\sigma_{\eta t}^2/\sigma_\varepsilon^2$. The persistence parameter, plotted in the top right panel of figure 5, declined during the great moderation, but has been increasing since 1995. The trend in predictability (see the bottom panels of figure 5) is the mirror image of persistence: predictability increases during the great moderation, but declines at the end of the sample.

Finally, figure 6 displays the nonparametric estimates of the posterior density of the parameters of the permanent volatility process and the irregular variance, and table 3 presents summary statistics concerning the distribution of the parameters and the convergence of the MCMC sampling scheme. We notice in particular that the Geweke's convergence diagnostics are fully satisfactory.

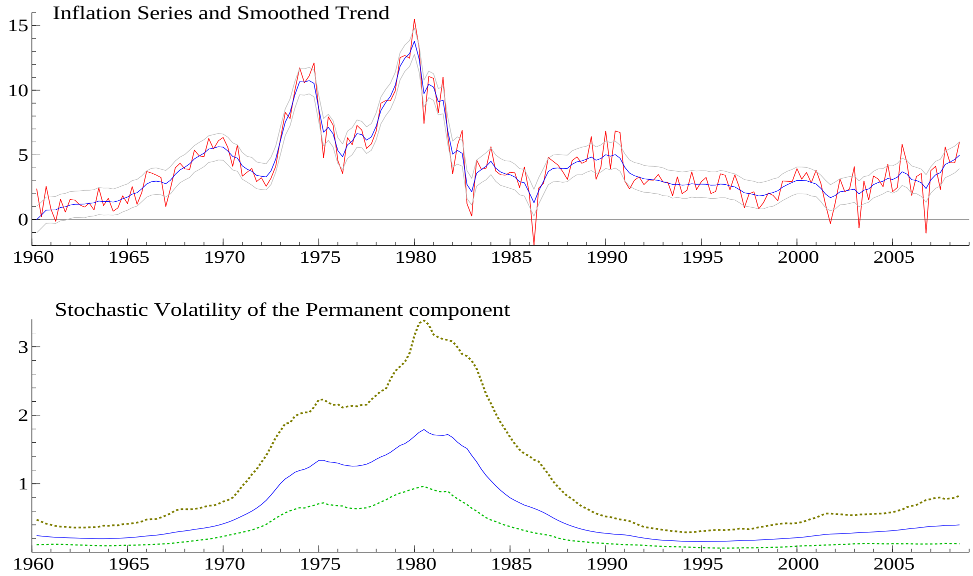

Figure 4: Upper: Quarterly inflation and its posterior mean level; Bottom: Volatility of the permanent component with 95% percent credible interval.

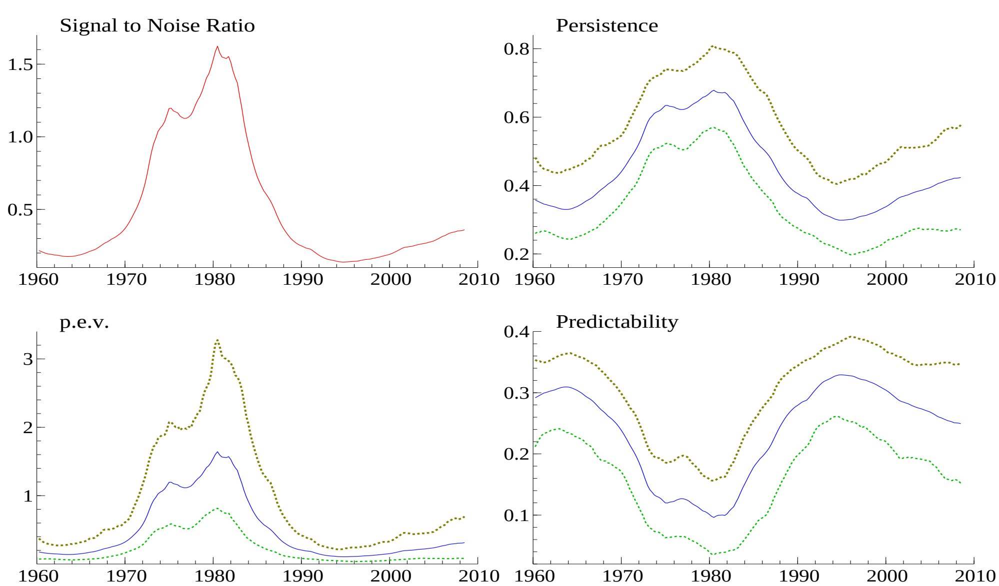

Figure 5: Upper left: Signal to noise ratio; Upper Right: Persistence Parameter with 95 percent credible interval. Bottom left: Prediction error variance with 95 percent credible interval; Bottom right: Predictability with 95 percent credible interval.

Table 3: Posterior mean, median, Geweke's statistic and Inefficiency factor for UC-SVc model.

Parameters	Mean	Median	Geweke's G_C	INEF
μ_2	-0.0233	-0.0229	-0.93	26.61
ϕ_2	0.9791	0.9801	-0.13	69.74
$\sigma^2_{\kappa_2}$	0.0472	0.0463	-0.82	32.16
σ^2_{ε}	1.2509	1.2403	1.80	135.00

6 Conclusions

The paper has provided a full Bayesian analysis of the local level model with stochastic volatility proposed by Stock and Watson (2007) for the U.S. quarterly CPI inflation rate. The model provides a simple yet effective decomposition of U.S. inflation into a permanent 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 permanent disturbances, not the transitory component.

The volatility of the permanent has been decreasing substantially after 1982, reaching a minimum around 1995, but has been increasing ever since, albeit at a small rate. The estimated volatility pattern support the view that a turning point took place around the mid-90ies and the great moderation is likely to be over. As correctly pointed out by a Referee, this result deserves further investigation as for its economic interpretation and implications. There are two possible explanations as to why it went undetected in previous analyzes: first and foremost, previous studies were conducted on a much shorter sample; for instance, the sample period consider by Stock and Watson (2007) ended in the fourth quarter of 2004, whereas our series ends in the 3rd quarter of 2008. The series, displayed in figure 1, does indeed display higher volatility at the end of the sample. Secondly, there are two substantial differences in model specification and estimation, that may play a role: on the one hand, our final specification, suggested by Bayesian model selection, is such that the volatility of the transitory component is constant. Also, the parameters of the permanent disturbance volatility process are estimated, rather than fixed.

The persistence implied by the model has been decreasing during the years of the great moderation and it stayed at historical lows up to the mid-90ies. Recently, persistence has been rising again. Correspondingly, the predictability of inflation increased during the great moderation up to maximum occurring around 1995 and it has been going down ever since.

DOI: 10.2202/1941-1928.1050

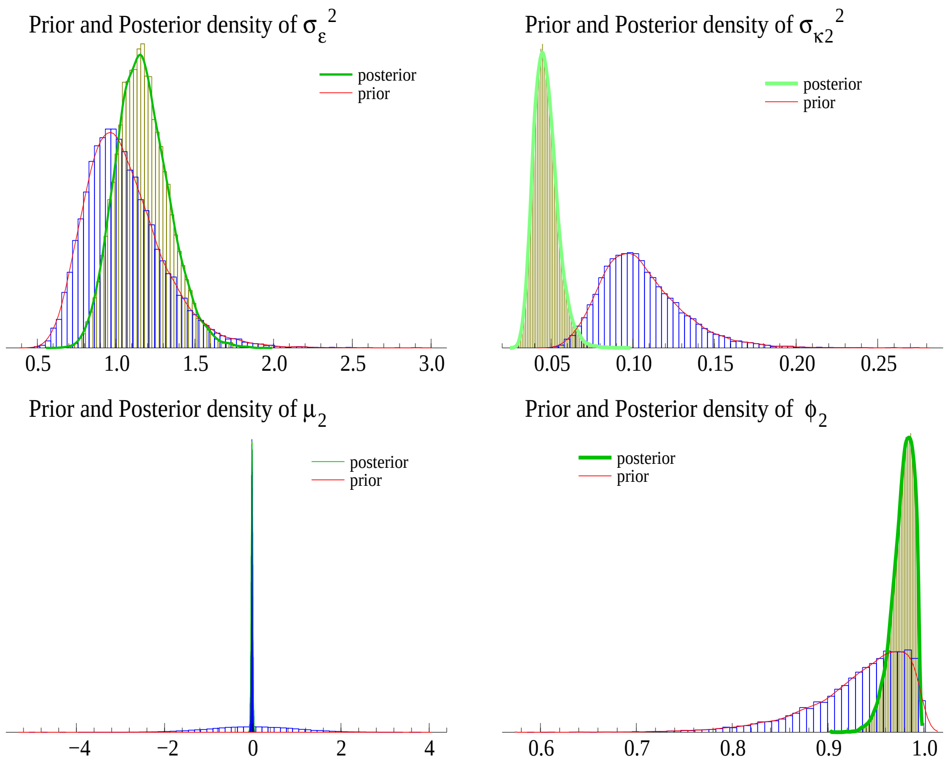

Figure 6: Prior and posterior densities of the parameters of the permanent volatility process and the irregular variance.

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

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	$Beta(a_i, b_i)$	$a_i = 20.50$	$b_i = 1.50$
$\sigma^2_{\kappa_1}$	$IG(\gamma_1, \beta_1)$	$\gamma_1 = 20.00$	$\beta_1 = 0.20$
$\sigma^2_{\kappa_2}$	$IG(\gamma_2, \beta_2)$	$\gamma_2 = 20.00$	$\beta_2 = 0.20$

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 Metropolis-Hastings sampling algorithm, similar to the one described in Kim, Shepard, and Chib (1998), which enforces the stationarity of the stochastic volatility process. Another possibility is to use a random walk Metropolis-Hasting that can be sketched as follows: 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 0.1. 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,
$$\log f(h_i|\mu_i, \phi_i, \sigma^2_{\kappa_i}) = -\frac{h^2_{i,0}}{2\sigma^2_{\kappa_i}} + \frac{1}{2}\log(1-\phi_i^2) - \frac{\sum_{t=1}^{T-1}(h_{i,t+1} - \phi_i h_{i,t} - \mu_i)^2}{2\sigma^2_{\kappa_i}}. \tag{8}$$
2. Using a Normal prior, the full conditional distribution of the parameters μ_i is $N(\hat{C}_i, \hat{D}_i)$ where:
$$\hat{C}_i = \hat{D}_i \quad \frac{C_i}{D_i^2} + \frac{1}{\sigma^2_{\kappa_i}} \sum_{t=1}^{T} (h_{i,t} - \phi_i h_{i,t-1}) \Bigg), \qquad \hat{D}_i = \left(\frac{1}{d_i^2} + \frac{T}{\sigma^2_{\kappa_i}}\right)^{-1}. \tag{9}$$
3. Using a conjugate Inverse Gamma prior, the full conditional of the variances of volatility processes are:
$$\sigma^2_{\kappa_i}|y, \alpha, h_i, \phi_i, \mu_i \sim IG\left\{\frac{T}{2} + \alpha_i, \beta_i + \frac{h^2_{i,0} + \sum_{t=1}^{T-1}(h_{i,t+1} - \mu_i - \phi_i h_{i,t})^2}{2}\right\}. \tag{10}$$

DOI: 10.2202/1941-1928.1050

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

$$f(h_{1,t}|h_{1,t-1}, h_{1,t+1}, y_t, \alpha_t, \theta) \propto f(h_{1,t}|h_{1,t-1}, h_{1,t+1}, \theta) f(y_t|\alpha_t, h_{1,t}). \quad (11)$$

It can be shown that

$$f(h_{1,t}|h_{1,t-1}, h_{1,t+1}, \theta) = f(h_{1,t}|h_{1,t-1}, \theta) f(h_{1,t+1}|h_{1,t}, \theta) \quad (12)$$

is a Gaussian density with mean

$$h^*_{1,t} = \frac{\mu_1(1-\phi_1) + \phi_1(h_{1,t-1} + h_{1,t+1})}{(1+\phi_1^2)}$$

and variance

$$v_1^2 = \frac{\sigma^2_{\kappa_1}}{1+\phi_1^2}$$

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

$$g(h_{1,t}) = \exp\left[-\left\{\frac{(h_{1,t+1} - \mu_1 - \phi_1 h_{1,t})^2}{2\sigma^2_{\kappa_1}} + \frac{(h_{1,t} - \mu_1 - \phi_1 h_{1,t-1})^2}{2\sigma^2_{\kappa_1}}\right\}\right] \times \times \frac{1}{\exp(h_1/2)} \exp\left[-\frac{(y_t - \alpha_t)^2}{2\exp(h_1)}\right] \quad (13)$$

for $t = 1, \ldots, T-1$, 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 $t = T$,

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

A similar sampling scheme is adopted for $h_{2,t}$.

APPENDIX B: Auxiliary Particle Filter

For evaluating the conditional likelihood, $f(y|\theta_k, M_k)$, for the SV specifications, we implemented an auxiliary particle filter (see Pitt and Shephard, 1999). The latter estimates the one-step-ahead predictive densities which enter the factorization: $f(y|\theta_k, M_k) = \prod_t f(y_{t+1}|Y_t, \theta_k, M_k)$, where $Y_t = \{y_1, \ldots, y_t\}$, and the predictive density is evaluated by sequential Monte Carlo methods as follows:

$$f(y_{t+1}|Y_t, \theta_k, M_k) = \frac{1}{M}\sum_{i=1}^{M} w_{1,t}^{(i)} \times \frac{1}{R}\sum_{j=1}^{R} w_{2,t}^{(j)}. \tag{14}$$

Here M denotes the number of particles, $w_{1,t}$ are the so-called first stage weights, R is the number of daughter particles (see below) and $w_{2,t}$ are the so-called second stage weights.

All the inferences will be conditional on (θ_k, M_k); henceforth, for notational simplicity we will omit these conditioning elements. After initializing the weights $w_{1,0} = \frac{1}{M}$ and drawing samples $z_0^{(i)}, i = 1, \ldots, M$, from the initial distribution of the random vector $z_t = (\alpha_t, h_{1,t}, h_{2,t})$, at time $t = 0$, with

$$\alpha_0 \sim N(0, 1000) \qquad h_{1,0} \sim N\left(0, \frac{\sigma_{\kappa_1}^2}{1-\phi_1^2}\right) \qquad h_{2,0} \sim N\left(0, \frac{\sigma_{\kappa_2}^2}{1-\phi_2^2}\right), \tag{15}$$

we iterate, for $t = 1, \ldots, T$, the following steps:

1. Set the first stage weights, $w_{1,t} \equiv \frac{1}{M}$.
2. Predict the unobserved states one-step-ahead, and update the weights, by

$$\begin{aligned} \bar{z}_{t+1}^{(i)} &= E(z_{t+1}|z_t^{(i)}), \\ w_{1,t}^{(i)} &= w_{1,t}^{(i)} \times f(y_{t+1}|\bar{z}_{t+1}^{(i)}), \end{aligned} \tag{16}$$

   where $f(y_{t+1}|\bar{z}_{t+1}^{(i)})$ is a Gaussian density with mean $\bar{\alpha}_{t+1}^{(i)}$ and variance $\exp(h_{1,t+1})$. The $w_{1,t}^{(i)}$ are the first stage weights described in Pitt and Shephard (1999).
3. Resample the particles $z_t^{(i)}$ with replacement R times (by multinomial resampling). Let $\tilde{z}_t^{(i)}$ denote the resampled particles.

4. Sample $z_{t+1}^{(i)}, i = 1, \ldots, R$, from $z_{t+1}|\tilde{z}_t^{(i)}, y_{t+1}$, using the approach by Godsill and Clapp (2001), which is based on the factorization:

$$f(z_{t+1}|z_t, y_{t+1}) = f(\alpha_{t+1}|h_{2,t+1}, \alpha_t, y_{t+1}) f(h_{1,t+1}|h_{1,t}) f(h_{2,t+1}|h_{2,t}) \quad (17)$$

where $f(h_{j,t+1}|h_{j,t}), j = 1,2$, are Gaussian densities with mean $\mu_j + \phi_j h_{j,t}$ and variance $\sigma^2_{\kappa_j}$, and

$$\alpha_{t+1}|h_{2,t+1}, \alpha_t, y_{t+1} \sim \mathrm{N}(m, S^2)$$

with

$$S^2 = \left(\frac{1}{\exp(h_{2,t+1})} + \frac{1}{\exp(h_{1,t+1})}\right)^{-1} \quad m = S^2\left(\frac{y_{t+1}}{\exp(h_{1,t+1})} + \frac{\alpha_t}{\exp(h_{2,t+1})}\right). \quad (18)$$

5. Compute the second stage weights:

$$w_{2,t}^{(i)} = \frac{f(y_{t+1}|z_{t+1}^{(i)}) f(z_{t+1}^{(i)}|\tilde{z}_t^{(i)})}{f(y_{t+1}|\bar{z}_{t+1}^{(i)}) f(\bar{z}_{t+1}^{(i)}|\tilde{z}_t^{(i)})}. \quad (19)$$

6. Resample M particles by multinomial resampling, with probabilities proportional to $w_{2,t}^{(i)}$.
7. Go to step 1.

APPENDIX C: 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 is constructed as follows: denoting $\theta = \{\theta_j, j = 1, \ldots, J\}$ the vector containing the hyperparameters, where the elements of the vector θ are $\{\mu_1, \phi_1, \sigma^2_{\kappa_1}, \mu_2, \phi_2, \sigma^2_{\kappa_1}\}$, consider the factorization of the joint conditional density:

$$\hat{\pi}(\theta^*|y) = \prod_{j=1}^{J} \hat{\pi}(\theta_j^*|y, \theta_1^*, \ldots, \theta_{j-1}^*).$$

Further, let $z = (h_1, h_2, \alpha)$.

The Chib and Jeliazkov (2001) algorithm takes the following steps:

- 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 $\{\phi_1^{(g)}, \sigma^{2,(g)}_{\kappa_1}, \mu_2^{(g)}, \phi_2^{(g)}, \sigma^{2,(g)}_{\kappa_2}, z^{(g)}\}$, $g = 1, \ldots, G$, 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:

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

     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 $\{\phi_1^{(j)}, \sigma^{2(j)}_{\kappa_1}, \mu_2^{(j)}, \phi_2^{(j)}, \sigma^{2(j)}_{\kappa_2}, z^{(j)}\}_{j=1}^J$ is are multiple (correlated) draws from the distribution:

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

  4. Denoting the probability of a move by

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

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

DOI: 10.2202/1941-1928.1050

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

- Run a reduced Gibbs sampling scheme 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 sampling scheme 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 sampling scheme on the following parameters $\{\mu_2, \sigma_{\kappa_2}^2, z\}$ and calculate μ_2^*.
- Run a reduced sampling scheme Gibbs on the following parameters $\{\sigma_{\kappa_2}^2, z\}$ and calculate $\sigma_{\kappa_2}^{2,(*)}$.

References

Bos, C. S., S. J. Koopman, and M. Ooms (2007): "Long Memory Modelling of Inflation with Stochastic Variance and Structural Breaks," *Tinbergen Institute Discussion paper*, N. 099/4.

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

Cappé, O., E. Moulines, and T. Rydén (2007): *Inference in Hidden Markow Models*, New York, USA: Springer-Verlag.

Cecchetti, G., P. Hooper, B. C. Kasman, K. Shoenholtz, and M. Watson (2007): "Understanding the Evolving Inflation Process," *U.S. Monetary Policy Forum*, February.

Chib, S. and I. Jeliazkov (2001): "Marginal Likelihood from the Metropolis-Hastings Output," *Journal of the American Statistical Association*, 90, 270–281.

Cogley, T., G. Primicieri, and T. J. Sargent (2007): "Inflation-Gap Persistence in the U.S." *NBER Working Paper*, N. 13749.

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

Durbin, J. and S. J. Koopman (2002): "A Simple and Efficient Simulation Smoother for State Space Time Series Analysis," *Biometrika*, 89, 603 – 616.

Gardner, E. (1985): "Exponential Smoothing: the State of the Art," *Journal of Forecasting*, 4, 1–28.

Gardner, E. (2006): "Exponential Smoothing: the State of the Art. Part II." *International Journal of Forecasting*, 22, 637 – 666.

Geweke, J. F. (1992): *Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments*, Bayesian Statistics 4, Oxford, UK: Oxford University Press.

Geweke, J. F. (2005): *Contemporary Bayesian econometrics and statistics*, New York, USA: Wiley.

Godsill, S. J. and T. Clapp (2001): "Improvement Strategies for Monte Carlo Particle Filters," in A. Doucet, N. De Freitas, and N. Gordon, eds., *Sequential Monte Carlo Methods in Practice*, New York, USA: Springer-Verlag.

Granger, C. W. J. and P. Newbold (1986): *Forecasting in Business and Economics*, Orlando, USA: Academic Press, Second edition.

Jacquier, E., N. G. Polson, and P. E. Rossi (1994): "Bayesian Analysis of Stochastic Volatility Models," *Journal of Business & Economic Statistics*, 12, 371 – 389.

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

Muth, J. F. (1960): "Optimal Properties of Exponentially Weighted Forecasts," *Journal of the American Statistical Association*, 55, 299 – 306.

Pitt, M. K. and N. Shephard (1999): "Filtering via Simulation: Auxiliary Particle Filters," *Journal of the American Statistical Association*, 94, 590 – 599.

Pivetta, F. and R. Reis (2007): "The Persistence of Inflation in the United States," *Journal of Economic Dynamics and Control*, 31, 1326 – 1358.

Shepard, N. (2006): *Stochastic Volatility. Selected readings*, Advanced Text in Econometrics, Oxford, UK: Oxford University Press.

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

DOI: 10.2202/1941-1928.1050