THE RIMINI CENTRE FOR ECONOMIC ANALYSIS

WP 43_14

Eric Eisenstat
University of Bucharest, Romania
RIMIR

Rodney Strachan
University of Queensland, Australia
The Rimini Centre for Economic Analysis, Italy

MODELLING INFLATION VOLATILITY

Copyright belongs to the author. Small sections of the text, not exceeding three paragraphs, can be used provided proper acknowledgement is given.

The *Rimini Centre for Economic Analysis* (RCEA) was established in March 2007. RCEA is a private, nonprofit organization dedicated to independent research in Applied and Theoretical Economics and related fields. RCEA organizes seminars and workshops, sponsors a general interest journal *The Review of Economic Analysis*, and organizes a biennial conference: *The Rimini Conference in Economics and Finance* (RCEF) . The RCEA has a Canadian branch: *The Rimini Centre for Economic Analysis in Canada* (RCEA-Canada). Scientific work contributed by the RCEA Scholars is published in the RCEA Working Papers and Professional Report series.

The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Rimini Centre for Economic Analysis.

The Rimini Centre for Economic Analysis
Legal address: Via Angherà, 22 – Head office: Via Patara, 3 - 47921 Rimini (RN) – Italy
www.rcfea.org - secretary@rcfea.org

Modelling Inflation Volatility

Eric Eisenstat*
University of Bucharest and RIMIR

Rodney W. Strachan
University of Queensland

October 22, 2014

*Corresponding author: Faculty of Business Administration, University of Bucharest, Blvd. Regina Elisabeta nr. 4-12, Bucharest (Romania), eric.eisenstat@faa.ro. Rodney Strachan is a Fellow of the Rimini Centre for Economic Analysis and the Centre for Applied Macroeconomic Analysis. We thank the co-editor and two referees for their valuable suggestions that helped greatly improve the paper. All errors are our own. The online appendix to this paper is available on the corresponding author's website: `http://www.rimir.ro/eric/research.html`.

Abstract

This paper discusses estimation of US inflation volatility using time varying parameter models, in particular whether it should be modelled as a stationary or random walk stochastic process. Specifying inflation volatility as an unbounded process, as implied by the random walk, conflicts with priors beliefs, yet a stationary process cannot capture the low frequency behaviour commonly observed in estimates of volatility. We therefore propose an alternative model with a change-point process in the volatility that allows for switches between stationary models to capture changes in the level and dynamics over the past forty years. To accommodate the stationarity restriction, we develop a new representation that is equivalent to our model but is computationally more efficient. All models produce effectively identical estimates of volatility, but the change-point model provides more information on the level and persistence of volatility and the probabilities of changes. For example, we find a few well defined switches in the volatility process and, interestingly, these switches line up well with economic slowdowns or changes of the Federal Reserve Chair. Moreover, a decomposition of inflation shocks into permanent and transitory components shows that a spike in volatility in the late 2000s was entirely on the transitory side and a characterized by a rise above its long run mean level during a period of higher persistence.

Summary: This paper considers the relative virtues of modelling inflation volatility as a stationary or non-stationary (random walk) process, both of which are conceptually problematic: random walks have increasing probability bounds while stationary processes cannot capture the large, low frequency movements we observe. We propose a change-point process that captures low-frequency movements while being always bounded in probability. To implement it, we develop a new methodology for sampling bounded state vectors. Our results suggest that for empirical purposes, either stationary or non-stationary processes are adequate, but our new process provides interesting inference on the changes in level and persistence of volatility.

Keywords: Inflation volatility, monetary policy, time varying parameter model, Bayesian estimation, Change-point model.

JEL Classification: C11, C22, E31

1 Introduction

The literature on modelling inflation is voluminous as inflation has an important place in many macroeconomic issues. For example, it is central to studies of the transmission of monetary policy shocks (Cogley and Sargent (2001 and 2005), Primiceri (2005), Sargent, Williams and Zha (2006), and Koop, Leon-Gonzalez and Strachan (2009)), there has been a resurgence in interest in the Phillips curve (King and Watson (1994), Staiger, Stock and Watson (1997), Koop, Leon-Gonzalez and Strachan (2010)), and there is a large literature devoted to forecasting inflation (e.g., Ang, Bekaert, and Wei (2007), Stock and Watson (2007 and 2009), D'Agostino, Gambetti, and Giannone (2009), Croushore (2010), Clark and Doh (2011), Chan (2013) and Wright (2013)).

Time varying parameter models of macroeconomic variables such as inflation have proven useful on a range of questions of interest to policymakers and the state space representation for these model has been a popular choice of specification. While there has been much attention to modelling the conditional mean of inflation, recently there has also been increasing interest in the variance with some evidence that the variance changes more (Primiceri (2005)) and more often (Koop, *et al.* (2009)) than do the mean coefficients. Therefore a feature that has proven important in such models is to allow for heteroscedasticity and a common specification in macroeconomics of this is stochastic volatility using a random walk for the state equation for log volatility (see for example, Cogley and Sargent (2005), Primiceri (2005) and Koop, *et al.* (2009)). This specification is attractive because of its parsimony, ease of computation and the smoothness it induces in the estimated volatility over time.

While the random walk specification is useful for practical reasons, it can be criticised as inappropriate since it implies that the range of likely values for volatility increases over time and is in the limit unbounded[1], which is clearly not what we observe. An alternative specification for stochastic volatility, which is commonly used in finance, is a stationary autoregressive model for the log volatility. Such a model implies inflation is bounded in probability at all horizons and has an easily derived stationary distribution. This property is appropriate for many financial processes where the variance shows only brief deviations far from its mean and then rapid mean reversion.

The behaviour of US inflation volatility, however, is not well described by a stationary, quickly mean reverting process. Although it is an unobserved latent process, common patterns have emerged in estimates presented in the literature on the behaviour of this process over time. Representative estimates of the volatility of inflation are presented in Figure 1. The pattern is an increase in the level of volatility that persisted during the 1970s and early 1980s, followed by a decline towards a lower lever over the late 1980s and early 1990s, and finally another increase in the 2000s. Other estimates in the literature differ in the detail, but what is generally evident in estimates of the volatility of inflation are large, low frequency movements. While the random walk model of volatility can be criticized for being incoherent, this model could be viewed as an approximation to the true process. In contrast, when we consider this behaviour, a time invariant stationary model of volatility does not appear appropriate for modelling inflation volatility.

[FIGURE 1 HERE]

This paper makes several contributions. i) We discuss the relative advantages and disadvantages of the random walk and stationary specifications of inflation volatility.

[1] By this we mean that the process is bounded in probability for any finite horizon, but these probability bounds widen over time such that in the limit, the process is not bounded in probability.

ii) This discussion leads us to present a change-point model of log inflation volatility that switches between stationary models with different levels and dynamics for inflation volatility. This model meets the theoretical concern that volatility should be bounded, but also permits the model to capture the occasional, large movements in the volatility level that have been observed over the past forty years. iii) While the specification and sampler are based upon the model of Koop and Potter (2007) for changes in the measurement equation, we develop a new specification that speeds computation when stationarity constraints are imposed and we expect that this algorithm will prove useful in a wider range of settings. iv) We compare outputs from our model with those from the random walk and stationary specifications and find that estimates of volatility differ little among the specifications and the estimated parameter values from the stationary model are close to the nonstationary region. The results suggest either a random walk or stationary model is a practically sensible specification to use for estimating inflation volatility.

Another contribution is v) a characterisation of regimes of inflation volatility since 1960. An advantage of our model over the random walk and stationary specifications is that it provides much more information on the level and persistence of inflation volatility. Our model also informs us on points at which inflation regimes change. The dates at which the model switches strongly suggest that the changes occur soon after economic slowdowns or a new Federal Chair appointment.

Finally, vi) we consider the volatility of the permanent component of inflation using core inflation and a permanent-transitory decomposition. Both approaches suggest that the recent rise we see in inflation volatility is in the transitory component of volatility.

In the next section, Section 2, we discuss and compare the attractive and less attractive properties of the random walk and stationary specifications for inflation volatility. In Section 3 we introduce our change-point model for the latent volatility process. This model is an adaptation of the model proposed by Koop and Potter (2007) for the measurement equation. However, we apply it to a latent process, log volatility, with the addition that we impose stationarity restrictions. The stationarity restriction complicates the estimation and so we present an new specification that simplifies and speeds estimation. In Section 4 we present and discuss the empirical results from the three models. Section 5 concludes the paper.

2 The Random Walk and Stationary Models

For both discussion and estimation, we will use a common measurement equation specification for inflation, y_t. In general, we write this as

$$y_t = \mu_t + \exp\left(h_t/2\right)\epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0,1).$$

We specify μ_t later and focus here on the specification of h_t. For a general specification of h_t we will use the following process:

$$h_t = \eta + \rho\left(h_{t-1} - \eta\right) + \nu_t, \quad \nu_t \sim \mathcal{N}\left(0, \sigma_h^2\right), \quad h_0 \sim \mathcal{N}\left(\underline{h}, V_{h_0}\right).$$

For the stationary specification (SP), common in finance, we assume $|\rho| < 1$ and $\underline{h} = \eta$. The random walk specification (RW), common in macroeconometrics, is obtained if we set $\rho = 1$ such that η drops out of the model.

Taking seriously the idea that the state equation is the prior, an unbiasedness argument has been made for the random walk as "the coefficients today have a distribution that is centered over last period's coefficients" (Koop, Léon-González and Strachan (2011)). On might equally, however, conclude RW is not a coherent prior as the variance of the log volatility grows linearly in time as $V(h_t|h_0) = t\sigma_h^2$. By contrast, the variance for SP converges to the constant $\frac{\sigma_h^2}{1-\rho^2}$ in the limit.

Correlation among the states, which is induced by the state equation, permits transmission of information in y_s about h_s to be transmitted to h_t and the strength of this transmission will depend upon the strength of the correlation. The process SP implies a correlation between h_t and h_{t-q} of $\rho^q\sqrt{\frac{1-\rho^{2(t-q)}}{1-\rho^{2t}}}$ and for RW this is $\sqrt{\frac{t-q}{t}}$. As t increases with fixed q, the correlation in RW converges to one suggesting that the time varying parameter model with a random walk state equation converges to a time invariant model. While this is prima facie incongruent, it has the advantage of increasing the transmission of information from earlier states to later ones. This property compensates for the increasing variance of, or uncertainty about, the later states.

For the purposes of estimation, the RW specification can be thought of as a parsimonious approximation to a SP specification with a high persistence. Further, the random walk implies greater smoothness and the stronger correlation structure induced by the RW specification aides in estimation of latent states.

3 A New Model of Inflation Volatility

A goal of this paper is to investigate the support for an alternative model that captures the feature of persistent shifts in the level of inflation volatility, but also implies that volatility has a stationary distribution at any point in time. To specify a model that is stationary at all times but permits changes in the level of volatility, we employ a change-point model based upon that developed in Koop and Potter (2007). In this section, we present our change-point model of inflation volatility. Existing techniques could be used to estimate this model, but they turn out to be computationally slow.

We therefore present an alternative representation that results in a more efficient sampler.

A significant difference between the model we present and the models used in Koop and Potter (2007) is that we apply the change-point process to the parameters governing the evolution of the volatility, which is a latent process. That is, the state equation for h_t is given as

$$h_t = \eta_{s_t} + \rho_{s_t}\left(h_{t-1} - \eta_{s_{t-1}}\right) + \nu_t, \quad \nu_t \sim \mathcal{N}\left(0, \sigma^2_{h,s_t}\right), \tag{1}$$

where $s_t \in \{1, \cdots, M\}$ indicates the regime at period t and M is the maximum number of regimes. Stationarity in each regime is imposed by assuming $|\rho_{s_t}| < 1$ for all s_t. Following Koop and Potter (2007), let

$$\begin{aligned}
\eta_m &= \eta_{m-1} + \xi_{\eta_m}, & \xi_{\eta_m} &\sim \mathcal{N}\left(0, \sigma^2_\eta\right), \\
\rho_m &= \rho_{m-1} + \xi_{\rho_m}, & \xi_{\rho_m} &\sim \mathcal{N}\left(0, \sigma^2_\rho\right), \\
\ln\sigma^2_{h,m} &= \ln\sigma^2_{h,m-1} + \xi_{\sigma_m}, & \xi_{\sigma_m} &\sim \mathcal{N}\left(0, \sigma^2_\sigma\right),
\end{aligned}$$

for $m = 1, \ldots, M$ and define the vectors $\eta = (\eta_1, \ldots, \eta_M)'$, $\rho = (\rho_1, \ldots, \rho_M)'$, $\sigma^2_h = \left(\sigma^2_{h,1}, \ldots, \sigma^2_{h,M}\right)'$, $h = (h_1, \ldots, h_T)'$, and $s = (s_1, \ldots, s_T)'$. These vectors are important as we do not use a Kalman filter based algorithm. Instead we use the more efficient precision based samplers (see Rue (2001), Chan and Jeliazkov (2009) and McCausland, Millera, and Pelletier (2011)).

An important feature of the Koop and Potter (2007) approach is the explicit specification of a prior on the duration of each regime. Define the time of the change-point from one regime to the next as $\tau_m = \{t : s_{t+1} = m+1, s_t = m\}$ such that the duration is defined as $d_m = \tau_m - \tau_{m-1}$. A hierarchical prior is specified for the duration. At the first level, the duration is *a priori* a Poisson process with mean λ_m. The parameter λ_m has a Gamma distribution $\mathcal{G}\left(\underline{\alpha}_\lambda, \beta_\lambda\right)$ in which $\underline{\alpha}_\lambda$ is fixed and the *rate* parameter β_λ is given a Gamma distribution $\mathcal{G}(\underline{\xi}_1, \underline{\xi}_2)$. This setup has a number of advantages and addresses several issues in modelling change-point processes as discussed in detail in Koop and Potter (2007). The notation we use for these parameters is identical to that of Koop and Potter (2007) and we refer the reader to that paper for further details. For our purposes we note that this structure implies the prior mean duration is

$$\underline{d}_m = E\left(d_m\right) = 1 + \underline{\alpha}_\lambda \left(\frac{\underline{\xi}_2}{\underline{\xi}_1 - 1}\right).$$

For η and σ_h^2, the above model implies no particular complication and we can complete the specification of the priors on these two parameter vectors with

$$\begin{aligned} \eta_0 &\sim \mathcal{N}(\kappa_{\eta_0}, V_{\eta_0}), & \ln \sigma_{h,0}^2 &\sim \mathcal{N}(\kappa_{\sigma_0}, V_{\sigma_0}), \\ \sigma_\eta^2 &\sim \mathcal{IG}(\gamma_\eta, \delta_\eta), & \sigma_\sigma^2 &\sim \mathcal{IG}(\gamma_\sigma, \delta_\sigma), \end{aligned}$$

where the κ's, V's, γ's and δ's are given constants. For a given volatility regime m to be stationary, however, we need to impose $|\rho_m| < 1$. In our approach, we specify a joint prior on $\left(\sigma_\rho^2, \rho_0, \rho\right)$, namely

$$\begin{aligned} p\left(\sigma_\rho^2, \rho_0, \rho\right) &\propto p\left(\sigma_\rho^2\right) p\left(\rho_0\right) p\left(\rho|\rho_0, \sigma_\rho^2\right) 1(|\rho| < \iota_M), \qquad (2) \\ p\left(\sigma_\rho^2\right) &= \mathcal{IG}_{\sigma_\rho^2}\left(\gamma_\rho, \delta_\rho\right), \\ p\left(\rho_0\right) &= \mathcal{N}_{\rho_0}\left(\kappa_{\rho_0}, V_{\rho_0}\right), \\ p\left(\rho|\rho_0, \sigma_\rho^2\right) &= \mathcal{N}_\rho\left(\rho_0 \iota_M, \sigma_\rho^2 \left(H'H\right)^{-1}\right), \end{aligned}$$

where ι_M denotes a $M \times 1$ vector of ones and $|\rho| < \iota_M$ is intended to mean that each element of ρ is less than one in absolute value. The the $M \times M$ matrix H is

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

The above specification yields simple computation to the extent that conditional on ρ, the hyper-parameters $\left(\rho_0, \sigma_\rho^2\right)$ can be sampled in a standard way. However, the conditional distribution

$$\rho|\sigma_\rho^2, \rho_0, \eta, h, \sigma_h^2, s \sim \mathcal{N}_{|\rho|<\iota_M}\left(\overline{\rho}, \overline{V}_\rho\right), \tag{3}$$
$$\overline{\rho} = \overline{V}_\rho\left(\frac{\rho_0}{\sigma_\rho^2}\widetilde{\iota}_M + X'\widetilde{\Sigma}_h^{-1}(h - \widetilde{\eta})\right),$$
$$\overline{V}_\rho = \left(\frac{1}{\sigma_\rho^2}H'H + X'\widetilde{\Sigma}_h^{-1}X\right)^{-1},$$

where $\widetilde{\eta} = (\eta_{s_1}, \ldots, \eta_{s_T})'$ and $\widetilde{\Sigma}_h = \text{diag}(\sigma^2_{h,s_1}, \ldots, \sigma^2_{h,s_T})$, is *multivariate truncated normal*.[2] Sampling from this distribution directly is difficult, and any type of accept-reject approach - i.e., draw ρ from an unrestricted multivariate normal distribution until we get a ρ that satisfies $|\rho| < \iota_M$ - would not work well because the constraints would be binding with a high probability for parameters in out of sample regimes.[3] To address this difficulty, we propose an approach that exploits the particular covariance

[2] $\widetilde{\iota}_M$ is the $M \times 1$ vector $\widetilde{\iota}_M = (1, 0, \ldots, 0)'$, and X is a $T \times M$ matrix with the element at row t, column m given by

$$x_{t,m} = \begin{cases} h_{t-1} - \eta_{s_{t-1}} & \text{if} \quad s_t = m \\ 0 & \text{otherwise} \end{cases}.$$

Clearly, as $X'X$ is diagonal, the precision for ρ is sparse and banded. Note that in the above definition, we assume $s_0 = 1$ for completeness.

[3] This includes recent algorithms such as the one used in Chan, *et al.* (2013).

structure in ρ. Specifically, we augment the parameter space with a $M \times 1$ latent vector of static factors $f = (f_1, \ldots, f_M)$, such that

$$
\begin{aligned}
p(\sigma_\rho^2, \rho_0, f, \rho) &\propto p\left(\sigma_\rho^2\right) p\left(\rho_0\right) p\left(f|\sigma_\rho^2\right) p\left(\rho|\rho_0, f, \sigma_\rho^2\right) \mathbb{1}(|\rho| < \iota_M), \quad (4)\\
p\left(\sigma_\rho^2\right) &= \mathcal{IG}_{\sigma_\rho^2}\left(\gamma_\rho, \delta_\rho\right),\\
p\left(\rho_0\right) &= \mathcal{N}_{\rho_0}\left(\kappa_{\rho_0}, V_{\rho_0}\right),\\
p\left(f|\sigma_\rho^2\right) &= \mathcal{N}_f\left(0, \sigma_\rho^2 I_M\right),\\
p\left(\rho|\rho_0, f, \sigma_\rho^2\right) &= \mathcal{N}_\rho\left(\rho_0 \iota_M + H^{-1} A f, 0.25\sigma_\rho^2 I_M\right),
\end{aligned}
$$

where A is a lower-triangular matrix such that $AA' = I_M - 0.25HH'$. Note that $I_M - 0.25HH'$ is guaranteed to be *positive-definite* and therefore A can be easily computed by the Cholesky decomposition. This is, in fact, closely related to the Stern (1992) decomposition and it is straightforward to verify that integrating (4) over f yields the original prior in (2). The latter implies two things: (i) the priors are equivalent, and (ii) all parameters besides ρ - including the hyper-parameters $\left(\rho_0, \sigma_\rho^2\right)$ - can be sampled marginally of f exactly as before.

The only role for the draws of f is to permit efficient sampling of ρ. Given a draw of f, the conditional (on f) distribution for ρ is

$$
\begin{aligned}
p\left(\rho|f, \sigma_\rho^2, \rho_0, \eta, \sigma_{h,m}^2, h, s\right) &= \prod_{m=1}^{M} p\left(\rho_m|f, \sigma_\rho^2, \rho_0, \eta, \sigma_{h,m}^2, h, s\right), \quad (5)\\
\rho_m|f, \sigma_\rho^2, \rho_0, \eta, \sigma_{h,m}^2, h, s &\sim \mathcal{N}_{|\rho_m|<1}\left(\overline{\rho}_m, \overline{V}_{\rho_m}\right),\\
\overline{\rho}_m &= \overline{V}_{\rho_m}\left(\frac{4}{\sigma_\rho^2} k_m + \frac{1}{\sigma_{h,m}^2} x_m'(h - \widetilde{\eta})\right), \quad (6)\\
\overline{V}_{\rho_m} &= \frac{\sigma_\rho^2 \sigma_{h,m}^2}{4\sigma_{h,m}^2 + \sigma_\rho^2 x_m' x_m}. \quad (7)
\end{aligned}
$$

In (6)-(7), k_m refers to the m-th element of the vector $k = \rho_0 \widetilde{\iota}_M + H'Af$, while x_m is the m-th column of X. These quantities are straightforward to compute and, hence, sampling ρ from independent univariate truncated normal distributions is straightforward (e.g. Robert (1995)).

Likewise, there is no difficulty in simulating f conditional on ρ. Noting that

$$\left(I_M + 4A' \left(HH'\right)^{-1} A\right)^{-1} = I_M - A'A,$$

the appropriate conditional distribution may be written as

$$p\left(f|\rho, \sigma_\rho^2, \rho_0, s\right) \sim \mathcal{N}\left(4D(\rho - \rho_0 \iota_M), \sigma_\rho^2 \left(I_M - A'A\right)\right), \qquad (8)$$
$$D = (I_M - A'A)(H^{-1}A)'.$$

Because $A'A$ is sparse and banded, simulation from (8) is fast even for large M. Moreover, the quantities D and $I_M - A'A$ involve only known constants and therefore need to be computed only once before commencing the MCMC. Details regarding the full Gibbs sampler are provided in the online appendix.

4 Empirical Results

In this section we report the estimates of inflation volatility from the random walk, stationary and change-point models. We first focus on the exact model in (1) and consider two alternative specifications for μ_t:

1. a constant-coefficient AR(4) equation restricted to stationarity;

2. bounded trend inflation (Chan, *et al.* (2013)) with the time-varying trend constrained to $[0, 6]$ and the time-varying AR coefficient to $[0, 1)$.

Specifically, the bounded trend inflation model we estimate is

$$y_t = \psi_t + \phi_t(y_{t-1} - \psi_{t-1}) + \exp(h_t/2)\epsilon_t, \qquad \epsilon_t \sim N(0,1), \tag{9}$$

with

$$\begin{aligned} p\left(\sigma_\psi^2, \psi_0, \psi\right) &\propto p\left(\sigma_\psi^2\right) p\left(\psi_0\right) p\left(\psi | \psi_0, \sigma_\psi^2\right) 1(|\psi| < 6\iota_T), \qquad (10) \\ p\left(\sigma_\psi^2\right) &= \mathcal{IG}_{\sigma_\psi^2}\left(\gamma_\psi, \delta_\psi\right), \\ p\left(\psi_0\right) &= \mathcal{N}_{\psi_0}\left(\kappa_{\psi 0}, V_{\psi 0}\right), \\ p\left(\psi | \psi_0, \sigma_\psi^2\right) &= \mathcal{N}_\psi\left(\psi_0 \iota_M, \sigma_\psi^2 \left(H'H\right)^{-1}\right), \end{aligned}$$

and

$$\begin{aligned} p\left(\sigma_\phi^2, \phi_0, \phi\right) &\propto p\left(\sigma_\phi^2\right) p\left(\phi_0\right) p\left(\phi | \phi_0, \sigma_\phi^2\right) 1(|\phi| < \iota_T), \qquad (11) \\ p\left(\sigma_\phi^2\right) &= \mathcal{IG}_{\sigma_\phi^2}\left(\gamma_\phi, \delta_\phi\right), \\ p\left(\phi_0\right) &= \mathcal{N}_{\phi_0}\left(\kappa_{\phi 0}, V_{\phi 0}\right), \\ p\left(\phi | \phi_0, \sigma_\phi^2\right) &= \mathcal{N}_\phi\left(\phi_0 \iota_M, \sigma_\phi^2 \left(H'H\right)^{-1}\right), \end{aligned}$$

For sampling purposes, we use the approach discussed in Section 3 to draw the constrained $\phi_1, \ldots, \phi_T$ and *accept-reject* to sample $\psi_1, \ldots, \psi_T$. Despite the slightly different prior construction, our estimates of these parameters are nearly identical to those reported in Chan, *et al.* (2013).

In the online appendix, we also estimate two additional specifications of μ_t, i.e.

1. time-varying intercept with constant AR(4) coefficients restricted to stationarity and the state equation for the intercept specified as a stationary AR(1) process;

2. a fully time-varying, unrestricted AR(2).

Among these alternatives, we find that they all yield similar results in terms of inference on inflation volatility. Of course, a variety of other specifications for the conditional mean are also possible—for example, one could include explanatory variables (i.e., along the lines of those recommended by Groen, *et al.*, 2013). Adding explanatory variables to the conditional mean is easily handled in our framework and requires only trivial extensions of the sampling algorithm.

The data are the quarterly inflation rate, computed as $400\ln\left(CPI_t/CPI_{t-1}\right)$ where CPI_t is U.S. total CPI data. The period covers 1947Q2-2013Q2, which after losing four lags for the mean equation (i.e. the maximum among the specifications considered), gives $T = 261$.[4] Using the change-point model, we are able to show estimates of the evolving persistence and level of volatility, as well as the probabilities of switching regimes at each point in time.

In the results reported below, we always set $M = 30$ and $\underline{\alpha}_\lambda = \underline{\xi}_1 = \underline{\xi}_2 = 30$, which implies a prior mean duration of $\underline{d}_m \approx 32$. Corresponding to this, we set $\delta_\eta = \delta_\rho = 0.25, \delta_\sigma = 0.01$ and $\gamma_\eta = \gamma_\rho = \gamma_\sigma = 5$.[5] The estimates of the parameters governing the behaviour of the latent inflation volatility process, such as ρ_m and η_m, are sensitive to the prior expected duration. In the online appendix, therefore, we further consider two additional combinations of hyper-parameters controlling the regime-search algorithm implying two different expected regime durations, i.e. $\underline{\alpha}_\lambda = \underline{\xi}_1 = \underline{\xi}_2 \in \{60, 120\}$ and $\underline{d}_m \in \{62, 122\}$.

Our general finding is that as the prior expected duration increases, the model will switch regimes less often and increasingly approximate the time invariant stationary

[4]The data were downloaded from the Federal Reserve Economic Data, St. Louis Fed.

[5]We also set throughout $\kappa_{\eta_0} = \kappa_{\rho_0} = 0$, $V_{\eta_0} = V_{\rho_0} = 10$ and $\kappa_{\sigma_0} = -2.45, V_{\sigma_0} = 0.29$. The latter implies a log-normal prior on $\sigma^2_{h,0}$ that is approximately equivalent to $\mathcal{IG}\left(5, 0.4\right)$—i.e., with $E\left(\sigma^2_{h,0}\right) = 0.1$ and $V\left(\sigma^2_{h,0}\right) = 3.3 \times 10^{-3}$.

model which has $\underline{d}_m = d_m \geq T$ and $M = 1$. At the other extreme, the random walk model has $\underline{d}_m = d_m = 1$ and $M = T$. Although the random walk does not technically nest within either the single or multiple stationary regime models since the support of ρ_m excludes the point $\rho_m = 1$, if the data prefers the random walk model this parameter will approach 1 and provide a good approximation to the random walk.

We obtain the following results using the MCMC algorithm detailed in the online appendix by taking 55,000 draws from the posterior distribution and discarding the first 5,000 as burn-in. We then *thin* the sample by recording every 5th draw to obtain a total of 10,000 upon which we base all inference. The inefficiency factors reported in Table 1 indicate that the chains are sufficiently well mixing.

[INSERT TABLE 1 HERE]

In Figure 2 we report the posterior estimates, $E(h_t|y)$, of the log volatility from the multiple stationary regimes (MSR), the single stationary regime (SSR) and the random walk (RW) models, for each of the conditional mean specifications outlined above. It is immediately apparent that the estimates differ very little across both the conditional mean and volatility specifications. Comparing the extremes of the single stationary regime model and the random walk model, we see that the estimates are close and there is considerable overlap of the error bands. The estimated parameters of the stationary model are close to the random walk model with the estimated posterior mean of ρ at $E(\rho|y) = 0.92$. The estimates from the random walk model are slightly smoother than from the stationary model. This is to be expected but it is not a dominant or distinguishing feature of the estimates.

[FIGURE 2 HERE]

An interesting observation arising from Figure 2 is that all specifications estimate a spike in volatility towards the end of the sample, after fairly constant, low volatility levels in the 1990s. Notably, inflation volatility in the past five years appears to exceed even the levels of the 1970s. It is therefore of interest to further examine the underlying process driving these estimated dynamics, and to this end we next consider the estimates of the level, η_m, and persistence, ρ_m, of inflation from the change-point model. Figure 3 shows the posterior estimates of these along with estimated probabilities of regime change points. We present estimates for both the AR(4) (left column) and bounded trend (right column) conditional mean specifications. However, it is clear that both specifications yield very similar estimates of η_m, ρ_m and the regime change point probabilities.

Looking first at the mean level of inflation, η_m, there is clear evidence of movements in the mean of inflation volatility. In particular, there is a fall in volatility in the 1960s, a rise from 1970 to the early 1980s, another decline from this point until the mid to late 1990s followed by another increase. The estimated mean level of volatility around 2008 is similar to that of the early 1980s, which raises the question of whether the great moderation had passed by this time, at least in this variable.

[FIGURE 3 HERE]

An editor conjectured that this increase in unconditional mean levels (as well as the actual volatility spike) is entirely driven by commodity prices (recall that we are using *total* CPI to measure volatility). In light of this, we also estimated our inflation volatility model using *core* CPI (e.g. excluding food and energy) and found substantial support for this conjecture—estimated mean levels remain low after the 1980s throughout the remainder of the sample and, although there is slight increase in inflation volatility in the late 2000s, it is still far lower than in the levels estimated for the 1970s. Moreover, core inflation volatility appears to be on the decline in the last 10 quarters of the sample. Details regarding these findings are reported in the online appendix and we focus on total inflation for the remaining of the paper. Later in this section, however, we consider a permanent/transitory decomposition of inflation. Interestingly, we find similarities in the behaviour of the volatility of the permanent component and that of core CPI.

Comparing the plots in Figure 2 with those in Figure 3 (top row), we see that there are extended periods when inflation volatility is above or below its mean. In particular, in periods of low volatility (e.g., during the 1990s) the volatility is lower than its mean and in periods of high volatility (e.g., late 1970s and early 1980s) the volatility is higher than its mean. Even assuming a change-point model of stationary inflation volatility, we see long periods of persistent deviations from the mean.

Figure 3 also reports the estimates for the volatility persistence parameter, ρ_m. The general pattern in the mean $E(\rho_m|y)$ is one of a rise in the level of persistence since the 1960s with a slight drop after 2008-2009. The association of persistence with level is not strong, but it does appear that the persistence falls when the mean level of volatility increases. However, the error bands for ρ_m are reasonably wide.

The bottom row depicts estimated probabilities of a change in regime at each point in time. In these plots, we also make evident periods in which the Federal Reserve Bank Chair was changed (dotted vertical lines) as well as NBER recorded recessions (shaded grey bands). A general observation is that the probability of a regime change at any point in time is always quite low, usually below 10%. Specific regime changes are therefore difficult to identify, which is likely due to the fact that the regime-switching structure is very high up in the prior hierarchy in our specification.

Nevertheless, there is still considerable movement in these plots and several spikes in probabilities can be observed, particularly at 1951Q3, 1973Q1, 1983Q3, 1991Q3 and 2008Q4. We note that these dates loosely align with significant economic events, usually following economic recoveries. The spikes in 1961, 1983, 1991 and 2008 appear at the end of recessions while the spike in 1951 does not appear close to a peak or trough and the spike in 1973 precedes the next downturn by 9 months. It is difficult to conclude much of a systematic relationship between recessions and changes in inflation volatility regimes from these results, except that the changes in inflation volatility appear associated with peaks or troughs in growth. These changes, however, appear to affect more the level of inflation volatility than its persistence.

Another interesting perspective on the regime-switches appears when we compare them with the terms of each of the Federal Chairperson. The change in the volatility regime appears to occur a few years after a new Chairperson takes office and the relationship is more consistent than that with the recessions. We can only make coincidental observations from our model but we feel that the cause of these break points does deserve further investigation.

To assess the performance of our approach in a more complex setting, and in light of some of the findings above, we also estimate the *unobserved components stochastic volatility* (UC-SV) model of Stock and Watson (2007). Specifically, the UC-SV model decomposes inflation shocks into *transitory shocks* and *permanent shocks*. In terms of our notation, the model can be written

$$\begin{aligned} y_t &= \mu_t + \exp(h_{y,t}/2)\epsilon_{y,t}, & \epsilon_{y,t} &\sim \mathcal{N}(0,1), \\ \mu_t &= \mu_{t-1} + \exp(h_{\mu,t}/2)\epsilon_{\mu,t}, & \epsilon_{\mu,t} &\sim \mathcal{N}(0,1), \end{aligned} \tag{12}$$

where Stock and Watson (2007) let the *transitory shock volatility*, $h_{y,t}$, and the *permanent shock volatility*, $h_{\mu,t}$, evolve as (independent) random walks.[6] Applying our

[6]Interestingly, Grassi and Proietti (2010) estimated the UC-SV model with stationary AR(1) processes specified for the log-volatilities, but found persistence parameters to be near 1.

change-point process, however, we specify

$$h_{y,t} = \eta_{y,s_t} + \rho_{y,s_t} \left(h_{y,s_t} - \eta_{y,s_t} \right) + \nu_{y,t}, \qquad \nu_{y,t} \sim \mathcal{N}\left(0, \sigma^2_{h_y,s_t}\right), \tag{13}$$

$$h_{\mu,t} = \eta_{\mu,s_t} + \rho_{\mu,s_t} \left(h_{\mu,s_t} - \eta_{\mu,s_t} \right) + \nu_{\mu,t}, \qquad \nu_{\mu,t} \sim \mathcal{N}\left(0, \sigma^2_{h_\mu,s_t}\right). \tag{14}$$

Note that in terms of this formulation, both the $h_{y,t}$ and $h_{\mu,t}$ regimes are determined by a single change-point process. The results reported in Figure 4 are based on this specification.[7]

To sample the model in (12)-(14), we again use our main MCMC algorithm, modified slightly to account for the two types of volatilities. The UC-SV model is more complex and so requires longer runs to achieve the same level of simulation accuracy. Consequently, we base the estimates below on a run of 220,000 iterations, with the first 20,000 draws discarded as burn-in; these are again thinned to a sample of 10,000 by recording every 20th draw. Inefficiency factors are reported in Table 2.

[INSERT TABLE 2 HERE]

Examining the unconditional means, we estimate that the permanent shock volatility level, η_{μ,s_t}, has dropped over the course of the sample period and remains low even towards the end of the sample. Likewise, the estimated permanent volatility itself, $h_{\mu,t}$, is constantly low through the end of the sample, remaining at the levels reached in the early 1990s. The spike in volatility around 2008 that we find with previous estimates is apparently entirely on the transitory side, as it is $h_{y,t}$ that is estimated to jump towards the end of the sample. Incidentally, the unconditional mean of the transitory shock log-volatility, η_{y,s_t}, is consistently rising, although this increase is

[7] We have also estimated a variant of this model letting $h_{y,t}$ and $h_{\mu,t}$ follow separate change-point processes. The results we obtained, but did not report, are nearly identical to those generated by a common change-point process. However, computation is noticeably slower with this generalization.

very gradual and remains far below the actual transitory shock volatility level in 2008.

The estimated permanent shock persistence parameter, ρ_{μ,s_t}, appears to decline over the sample period, reaching around 0.5 towards 2013. Conversely, the transitory shock persistence parameter, ρ_{y,s_t}, increases over time, with the highest level (around 0.8) occurring in 2008, followed by a decline. Therefore, these results suggest that the spike in inflation volatility around 2008 is characterized by persistent growth of the transitory shock volatility above its long run mean level. Interestingly, estimates with core inflation data discussed above (i.e. of a single volatility shock) appear to resemble the results we obtain for permanent shock volatility.

[INSERT FIGURE 4 HERE]

Finally, we provide some summary measures of model preference. In Table 3 we report the deviance information criteria (DIC) (Spiegelhalter, *et al.* (2002)), which has previously been used in a related set up to compare stochastic volatility models (e.g. Berg, *et al.* (2004)) and is generally used to assess the ability of the model to predict future data that would be generated by the same mechanism as the existing data. It is particularly useful with large complex models with correlated latent parameters, such as state space models, where simply counting the number of parameters as a measure of model complexity is not appropriate due to the prior correlation structure.

[TABLE 3 HERE]

For the multiple stationary regime model, we consider three different prior specifications of varying prior mean duration and compare these to the single stationary regime and random walk models. For each of these five volatility specifications, we also compare two specifications of the unconditional mean (AR(4) and bounded trend) as well as the UC-SV model. A smaller value of the DIC for one model than for another model suggests the first model is preferable. The results in this table show little difference among the range of stationary models (i.e. MSR and SSR), although there is a slight preference for the shortest duration change-point model ($\underline{\alpha}_\lambda = 30$) over the other three stationary models when the AR(4) mean is specified. The random walk model is the least preferred model and the long duration change-point model (with $\underline{\alpha}_\lambda = 120$) and single stationary regime model are equally ranked.

Similar results arise when estimating the bounded trend in the mean, although the differences among stationary models are less pronounced. Slightly different measures are obtained in the UC-SV case where a stationary model with longer regime durations (with $\underline{\alpha}_\lambda = 120$) appears to be the most preferable, while all others are essentially equivalent. This suggests that there is a preference in the data for stationarity but changing regimes are also very important.

To further assess the predictive abilities of these models in practice, we compare predictive likelihoods for the period: 2008Q3–2013Q2. The beginning of this period is interesting because annualized inflation dropped suddenly from 6.1% in 2008Q3 to -9.9% in 2008Q4, making it a very difficult period to forecast. In contrast, the end of the period is characterized by relatively stable inflation volatility. The results in Table 4 summarize the joint predictive likelihoods.

Evidently, there is little difference in predictive abilities between models of varying volatility specifications. Instead, the most pronounced differences are due to the conditional mean specification, where the bounded trend inflation model performs best while the simplest AR(4) model exhibits the least predictive power. These results serve to reinforce the point that for inflation forecasting purposes, either a stationary or random walk specification of the volatility process does equally well. However, our change point process does provide a useful instrument for drawing inference regarding the evolution of volatility itself.

[TABLE 4 HERE]

5 Conclusion

This paper has considered the relative advantages and disadvantages of the random walk and stationary specifications of inflation volatility and introduced a new change-point model of log inflation volatility that incorporates some desirable features. The new model ensures that inflation volatility is bounded in probability while permitting infrequent but large changes in the volatility level and persistence; both of which are frequently discussed features of volatility over the past forty years. A comparison of estimated volatility from a range of models suggests that the specification matters little for this purpose. While information criteria show some preference for the single regime stationary model over the random walk model, the stationary model will produce estimates that approximate the random walk process. On the debate over which specification is appropriate, if the objective is estimation of volatility then either would seem appropriate.

The change-point model of log volatility provides new insights on the evolution of this process. The pattern of volatility shows the often observed decline from the 1980s to the 1990s, but also indicates a rise over the 2000s. Inflation persistence generally increases over most of the sample. However, a decomposition of the shocks reveals that this increase in volatility levels and persistence is associated with transitory innovations rather than permanent ones. Likewise, the volatility spike is no longer present when core inflation data is used. We also estimate regime-switching probabilities and find that they align with periods near economic slowdowns and, interestingly, tend to follow changes of the Federal Reserve Chair.

References

Ang A., G. Bekaert and M. Wei (2007) "Do Macro Variables, Asset Markets or Surveys Forecast Inflation Better?", *Journal of Monetary Economics* 54: 1163-1212.

Berg A., R. Meyer and J. Yu (2004) "Deviance Information Criterion for Comparing Stochastic Volatility Models", *Journal of Business & Economic Statistics* 22(1): 107-120.

Billingsley, P. (1979) Probability and Measure. Wiley, New York.

Chan J. C. C. (2013) "Moving Average Stochastic Volatility Models with Application to Inflation Forecast", *Journal of Econometrics* 176(2): 162–172.

Chan J. C. C., G. Koop, and S. M. Potter (2013) "A new model of trend inflation", *Journal of Business & Economic Statistics* 31(1): 94–106.

Chan J. C. C. and I. Jeliazkov (2009) "Efficient simulation and integrated likelihood estimation in state space models", *International Journal of Mathematical modelling and Numerical Optimisation* 1: 101–120.

Clark T. E. and T. Doh (2011) "A Bayesian Evaluation of Alternative Models of Trend Inflation", Federal Reserve Bank of Cleveland, working paper no. 11-34.

Cogley T. and T. Sargent (2001) "Evolving post-World War II inflation dynamics", *NBER Macroeconomic Annual* 16: 331-373.

Cogley T. and T. Sargent (2005) "Drifts and volatilities: Monetary policies and outcomes in the post WWII U.S.", *Review of Economic Dynamics* 8, 262-302.

Croushore D. (2010) "An Evaluation of Inflation Forecasts from Surveys Using Real-Time Data", *The B.E. Journal of Macroeconomics*, De Gruyter, May 10(1): 1-32.

D'Agostino A., Gambetti, L. and Giannone, D. (2009) "Macroeconomic forecasting and structural change", ECARES working paper 2009-020.

Geweke J. (1991) “Efficient simulation from the multivariate normal and student-t distributions subject to linear constraints”, In E. M. Keramidas, editor, *Computer Science and Statistics, Proceedings of the Twenty-Third Symposium on the Interface*, pages 571–578. Fairfax: Interface Foundation of North America, Inc..

Grassi, S. and Proietti, T. (2010) “Has the volatility of U.S. inflation changed and how?”, *Journal of Time Series Econometrics*, 2(1), Article 6.

Groen, J. J. J., Paap, R. and Ravazzolo, F. (2013) “Real-time inflation forecasting in a changing world”, *Journal of Business and Economic Statistics*, 31, 29-44.

King R. and Watson, M. (1994) “The post-war U.S. Phillips curve: a revisionist econometric history”, *Carnegie-Rochester Conference Series on Public Policy*, 41, 157-219.

S. Kim, N. Shephard and S. Chib (1998) “Stochastic volatility: likelihood inference and comparison with ARCH models”, *The Review of Economic Studies*, 65(3): 361–393.

Koop G., R. Léon-González and R. W. Strachan (2009) “On the Evolution of Monetary Policy”, *Journal of Economic Dynamics and Control* 33, 997-1017.

Koop G., R. Léon- González and R. W. Strachan (2010) “Dynamic probabilities of restrictions in state space models: An application to the Phillips curve”, *Journal of Business and Economic Statistics* Vol. 28, No. 3: 370-379.

Koop G., R. Léon-González and R. W. Strachan (2011) “Bayesian Inference in the Time Varying Cointegration Model”, *The Journal of Econometrics* 165, 210-220.

Koop G. and S. M. Potter (2006) “Estimation and forecasting in models with multiple breaks”, Working Paper.

Koop G. and S. M. Potter (2007) “Estimation and forecasting in models with multiple breaks”, *Review of Economic Studies* 74: 763–789.

McCausland W. J., S. Millera, and D. Pelletier (2011) “Simulation smoothing for state-space models: A computational efficiency analysis", *Computational Statistics and Data Analysis* 55: 199–212.

Primiceri G. (2005) "Time varying structural vector autoregressions and monetary policy", *Review of Economic Studies* 72: 821-852.

Robert C. P. (2005) "Simulation of truncated normal variables", *Statistics and Computing* 5(2): 121–125.

Rue H. (2001) "Fast sampling of Gaussian Markov random fields with applications", *Journal of the Royal Statistical Society Series B*, 63: 325-338.

Sargent T., Williams, N. and Zha, T. (2006) "Shocks and government beliefs: The rise and fall of American inflation", *American Economic Review* 96: 1193-1224.

Staiger D., Stock, J. and Watson, M. (1997) "The NAIRU, unemployment and monetary policy", *Journal of Economic Perspectives* 11: 33-49.

Spiegelhalter D. J., N. G. Best, B. P. Carlin, and A. van der Linde (2002), "Bayesian Measures of Model Complexity and Fit" (with discussion), *Journal of the Royal Statistical Society, B* 64: 583-639.

Stern S. (1992) "A method for smoothing simulated moments of discrete probabilities in multinomial probit models", *Econometrica* 60(4): 943–952.

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

Stock J. H. and M. W. Watson (2009), "Phillips Curve Inflation Forecasts," in *Understanding Inflation and the Implications for Monetary Policy*, eds., J. Fuhrer, Y. Kodrzycki, J. Little, and G. Olivei, Cambridge: MIT Press, ch. 3: 99–202.

Wright J. H. (2013) "Evaluating Real-Time VAR Forecasts with an Informative Democratic Prior", *Journal of Applied Econometrics* 28: 762-776.

Tables and Figures

	AR(4)				**Bounded Trend**			
	h_t	η_{s_t}	ρ_{s_t}	σ^2_{h,s_t}	h_t	η_{s_t}	ρ_{s_t}	σ^2_{h,s_t}
min	1.28	1.89	2.59	20.71	1.30	2.23	2.37	26.95
10%	1.49	2.16	3.55	21.72	1.49	2.48	3.94	27.46
25%	1.63	2.49	3.79	23.00	1.62	2.90	4.54	28.27
50%	1.80	3.89	5.52	24.26	1.86	3.43	5.22	29.63
75%	2.08	5.51	6.08	24.85	2.24	5.06	6.83	30.51
90%	2.49	6.43	11.11	24.95	2.65	8.06	11.51	30.74
max	4.87	10.20	15.30	25.05	3.92	10.69	12.78	31.17

Table 1: Inefficiency factors (computed on *thinned* draws) for main parameters of the Multiple Stationary Regime specification under the AR(4) and Bounded Trend conditional means.

	Permanent Shocks Vol.				**Transitory Shocks Vol.**			
	h_t	η_{s_t}	ρ_{s_t}	σ^2_{h,s_t}	h_t	η_{s_t}	ρ_{s_t}	σ^2_{h,s_t}
min	1.69	3.69	0.98	7.93	0.88	1.99	1.54	7.85
10%	2.31	4.62	1.13	8.20	1.41	2.11	1.73	8.15
25%	3.57	6.02	1.61	8.67	1.77	2.68	1.85	8.56
50%	5.17	10.04	4.83	8.88	2.18	3.53	2.12	8.96
75%	9.71	12.51	7.56	8.98	3.20	4.60	3.22	9.18
90%	10.37	13.22	9.78	9.04	5.66	7.27	5.83	9.25
max	39.40	13.61	12.88	9.09	29.62	10.36	13.32	9.32

Table 2: Inefficiency factors (computed on *thinned* draws) for main parameters of the UC-SV with Multiple Stationary Regime volatilities specification.

Volatility Prior	**Conditional Mean** AR(4)	Bounded Trend	UC-SV
$\underline{\alpha}_\lambda = \underline{\xi}_1 = \underline{\xi}_2 = 30$	982	991	915
$\underline{\alpha}_\lambda = \underline{\xi}_1 = \underline{\xi}_2 = 60$	984	991	914
$\underline{\alpha}_\lambda = \underline{\xi}_1 = \underline{\xi}_2 = 120$	986	993	910
Single Stationary Regime	987	995	915
Random Walk	991	1001	914

Table 3: DICs for 15 models.

	AR(4)			**Bounded Trend**			
	MSR	**SSR**	**RW**	**MSR**	**SSR**	**RW**	**UC-SV**
2008Q3-2009Q2	-19.09	-19.81	-18.29	-18.73	-18.31	-17.68	-18.06
2009Q3-2010Q2	-11.41	-11.16	-11.31	-8.84	-8.94	-9.53	-9.78
2010Q3-2011Q2	-8.62	-8.36	-9.03	-7.37	-7.37	-8.05	-8.79
2011Q3-2012Q2	-7.92	-7.59	-7.96	-7.04	-6.80	-7.02	-7.39
2012Q3-2013Q2	-6.95	-6.68	-6.95	-6.65	-6.46	-6.59	-6.77
Joint	-53.99	-53.59	-53.55	-48.63	-47.89	-48.87	-50.80

Table 4: A comparison of log predictive likelihoods for 2008Q3–2013Q2.

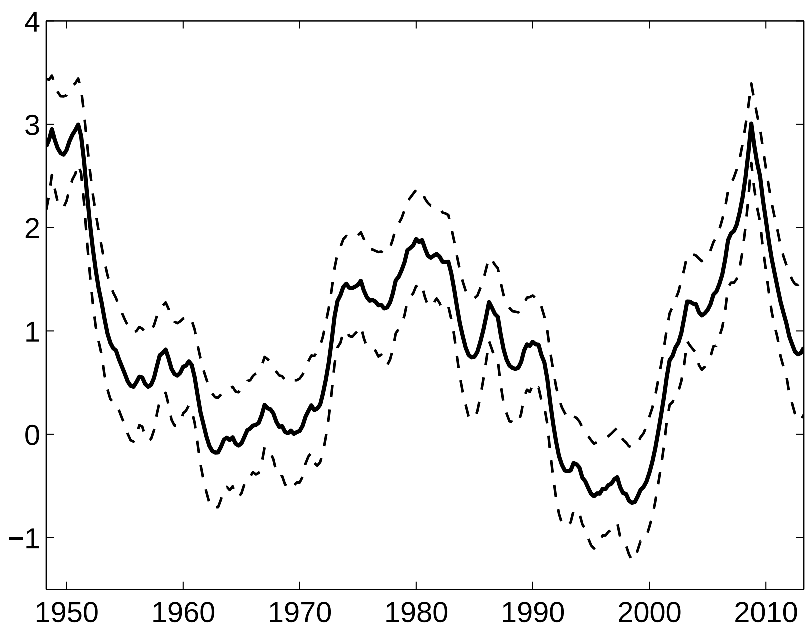

Figure 1: Plot of inflation volatility from 1948 to 2013.

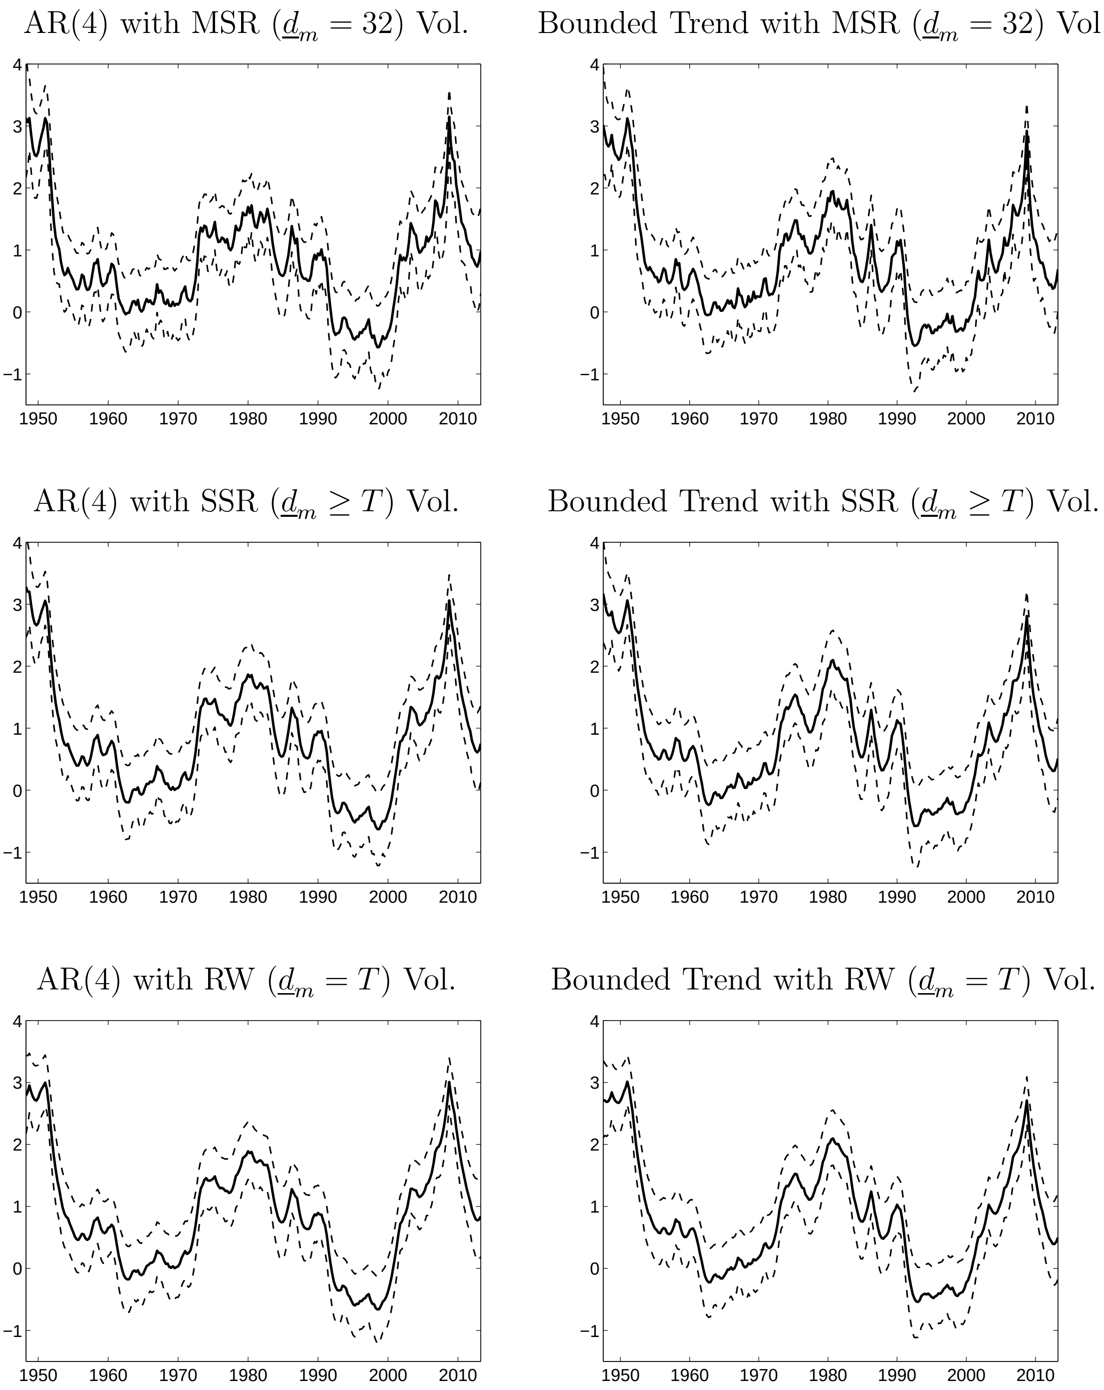

Figure 2: Posterior median and the $(16\%, 84\%)$ probability interval for the log-volatility h_t. Note: for the *single stationary regime* case (middle row), we get $\hat{\rho} \approx 0.92$ for both the AR(4) and Bounded Trend mean specifications.

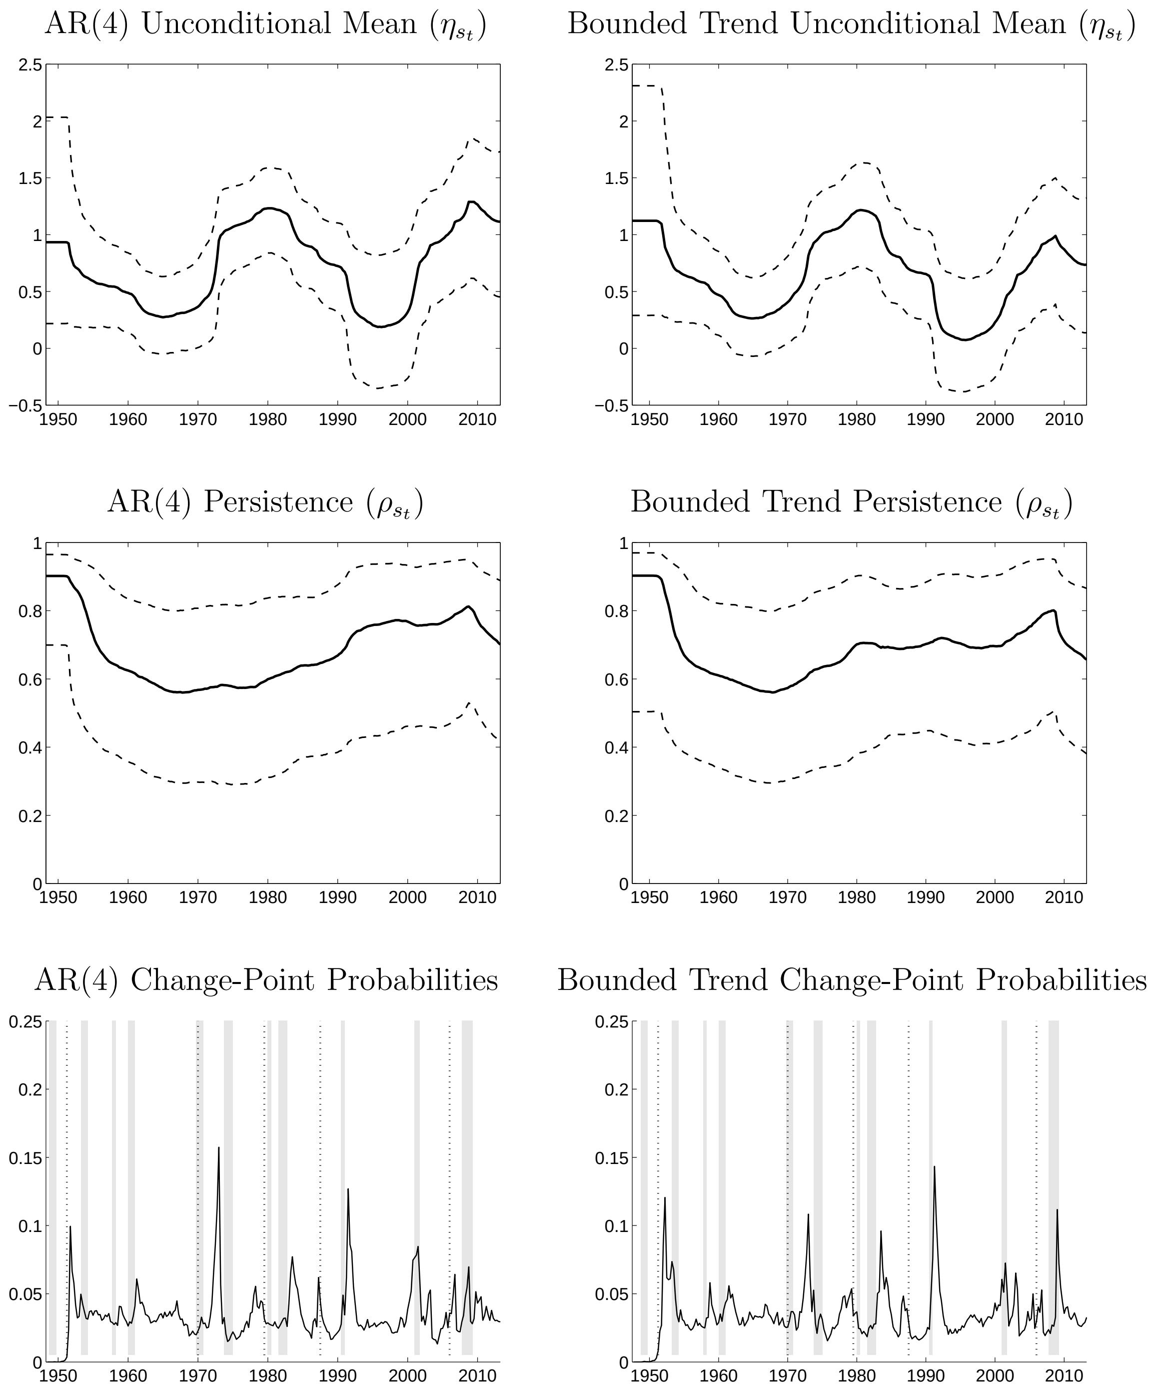

Figure 3: Posterior median and the $(16\%, 84\%)$ probability interval for parameters (η_{s_t}, ρ_{s_t}) of the Multiple Stationary Regime specification, under the AR(4) and Bounded Trend conditional means. Also shown are the regime change point probabilities with Federal Reserve Bank Chairperson changeovers (vertical dotted lines) and the NBER recessions (shaded grey bands).

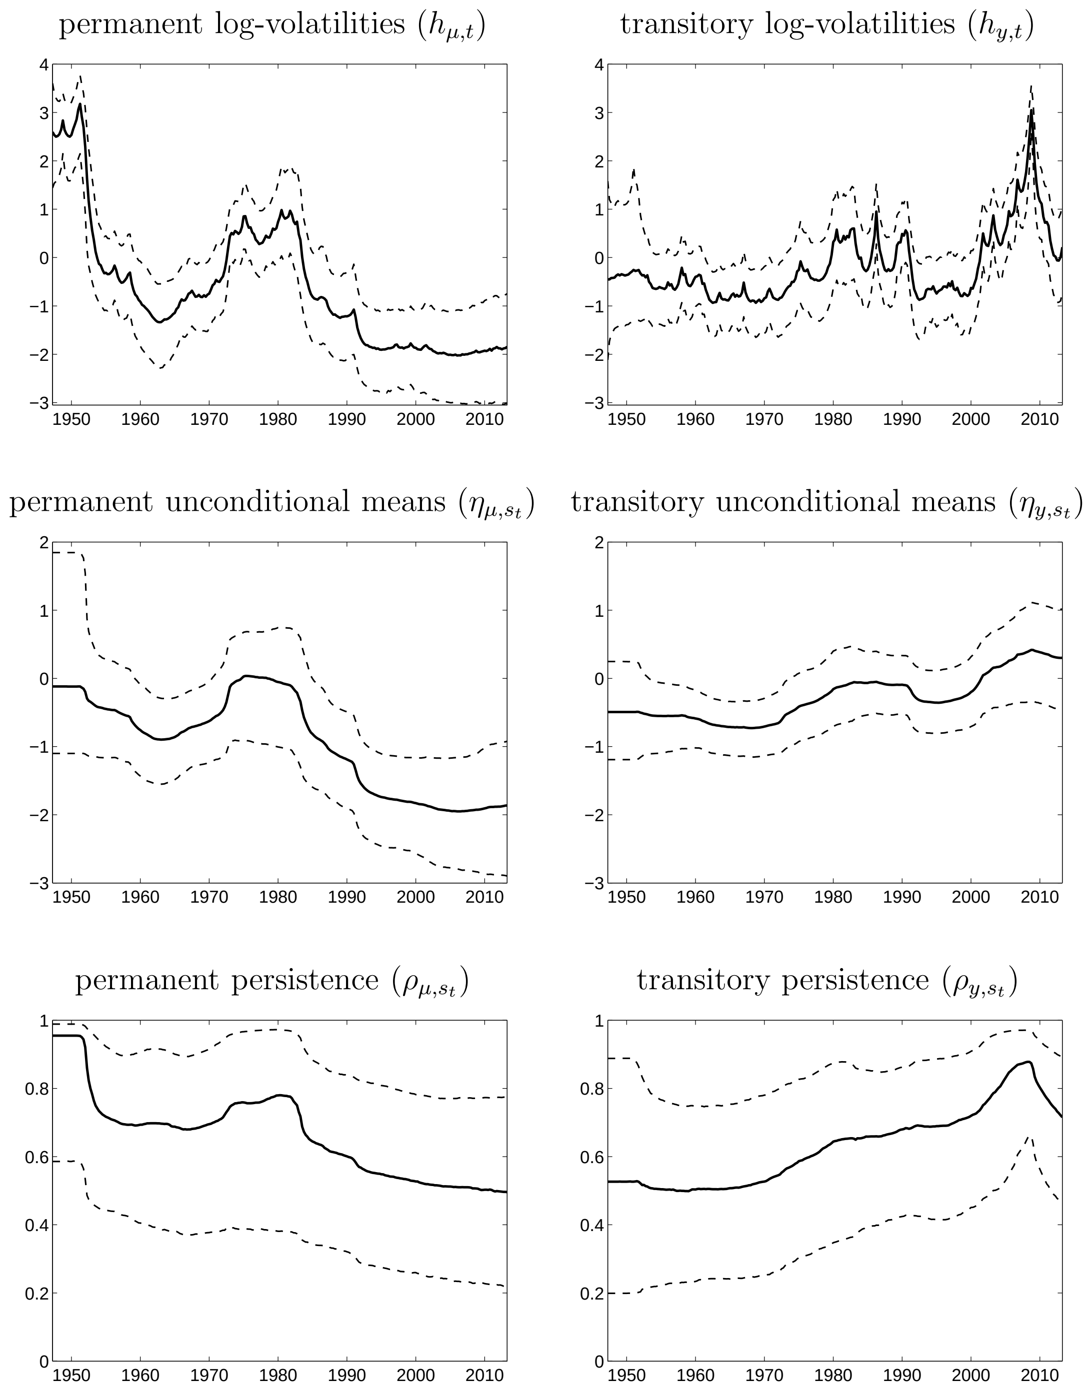

Figure 4: Estimates of the log-volatility regime processes for the permanent and transitory shocks under the UC-SV specification. The dashed lines depict the $(16\%, 84\%)$ HPD intervals.