Inflation Volatility with Regime Switching^{*}

Maksim A. Isakin[†]

Phuong V. Ngo[‡]

May 1, 2018

Abstract

This paper investigates the U.S. inflation dynamics by allowing regime switching in an unobserved components stochastic volatility model. The likelihood is constructed using particle filter with a new resampling procedure that ensures the smoothness of the likelihood function. We estimate the model with three regimes using maximum likelihood and determine prevailing regimes over the 1960:I - 2017:IV period. Our results show that since the start of the Great Recession, the U.S. inflation has entered a regime with moderate volatility where most of the volatility comes from transitory shocks. In terms of forecasting, our model produces a better outcome compared to the model without regime switching.

JEL classification: C14, C15, C32, E31, E37.

Keywords: regime switching, particle filter, unobserved components, stochastic volatility, inflation forecast.

^{*}The authors acknowledge the financial support from the Faculty Scholarship Initiative (FSI) Program of the Cleveland State University. We are also grateful for the support from the Ohio Supercomputer Center.

[†]Cleveland State University, Department of Economics, 2121 Euclid Avenue, Cleveland, OH 44115. Corresponding author. Tel. 216.687.4530, email: m.isakin@csuohio.edu.

[‡]Cleveland State University, Department of Economics, 2121 Euclid Avenue, Cleveland, OH 44115.

1 Introduction

Since the 1960s, the U.S. inflation has exhibited several episodes with distinctive dynamics. Most prominent episodes include dramatic rise of the inflation level and its volatility in 1970s, disinflation under tightening monetary policy in the early 1980s, and low and stable inflation in the 1990s. As a result of these changes, many univariate models fail to explain the inflation dynamics over a long period or reveal changes in the parameter estimates over different subsamples. See Stock and Watson (2007) for discussion. Recently, due to low but relatively volatile inflation in the wake of the Great Recession with the zero lower bound (ZLB) on nominal interest rate, a number of papers have investigated whether the U.S. economy has recently transitioned to an unintended long-lasting low inflation or deflation regime with higher volatility, for example Aruoba et al. (2018).

In this paper, we develop a model of inflation with unobserved mean and regime switching stochastic volatilities of the measurement error (transitory shock) and the mean innovation (permanent shock). We build on the random walk plus noise or local level model (see Durbin and Koopman (2012) for details) and introduce autoregressive volatilities processes with regime switching coefficients. In the model, the logarithms of the volatilities of the transitory and permanent shocks are stationary processes conditional on the regime. We present evidence that supports this specification. An important feature of our model is that it explicitly takes into account different inflationary regimes and makes it possible to estimate the transition probabilities between different regimes and forecast the probabilities of regimes.

The local level model with stochastic volatilities is used to study the predictability of inflation in Stock and Watson (2007). In their unobserved components stochastic volatility (UCSV) model both volatilities of transitory and permanent shocks are assumed to follow independent random walks with same predetermined (calibrated) variance of the error term. This variance controls the smoothness of the stochastic volatility processes, and is set equal to 0.2 in Stock and Watson (2007) for their 1953:I–2004:IV sample. As we show in Section 4, this value is very large especially for certain periods. The UCSV model requires a large variance of the error to generate enough volatility and fit the model in the whole sample. We argue that the volatilities are stationary processes with regime switching parameters.

Our paper has three main contributions. First, it introduces AR(1) volatility processes with regime switching parameters into the UCSV model of inflation, hereafter called the RS-UCSV model. This representation nests the original UCSV model of Stock and Watson (2007) as a special case with one regime and unit-root volatilities processes. In addition, the model provides predicted probabilities of regimes which are of interest. Second, the paper develops a particle filter that adopts regime switching parameters. In particular, we propose a resampling procedure that preserves the continuity of the likelihood function and enables efficient maximum likelihood estimation. The resampling procedure uses sampling from an importance density and the kernel density estimate of the important weights. Third, the model is applied to the 1960:I-2017:IV U.S. inflation series. Our estimation results show that since the start of the Great Recession, the U.S. inflation has entered a regime with moderate volatility where larger fraction of volatility is due to transitory shocks. In addition, our RS-UCSV model produces better inflation forecast compared to the model without regime switching.

The U.S. inflation series is highly persistent and suggests a unit-root or highly persistent very slow mean-reverting specification; for discussion see Barsky (1987), Ball and Cecchetti (1990), and Stock and Watson (2007). Furthermore, Engle (1983) shows that the inflation data exhibits heteroskedasticity. Recent literature emphasizes structural breaks or changing regimes in the inflation dynamics. Regime changes can stem from differences in monetary policy as discussed in Clarida et al. (2000), Taylor (1993), and Malmendier et al. (2017). Cogley and Sargent (2003) and Cogley et al. (2010) develop an empirical model that allows both coefficients and volatilities to vary over time. Since the changes in the macroeconomic and institutional environment can be significant and rapid, e.g. changes in monetary policy, discrete regime switching models can better describe the data. Garcia and Perron (1996) apply the hidden Markov regime switching approach to modeling structural changes in the ex-post real interest rate and inflation series. Clark (2006) examines the structural breaks in disaggregated inflation indexes. Zhang (2008) finds multiple structural breaks in the persistence of the U.S. inflation series. Chan et al. (2013) modify the UCSV model of Stock and Watson (2007) by imposing bounds on the volatilities of the innovations of the permanent and transitory components. They also allow for a time-varying degree of persistence of the transitory component. This modification has the same flavor of some models used in Clark and Doh (2014), where they evaluate alternative models of trend inflation.¹

The rest of the paper is organized as follows. Section 2 introduces the model. In Section 3

¹The literature on modeling inflation is too large to summarize. The reader can read Stock and Watson (2007) and Clark and Doh (2014) for a complete list of models.

we introduce the filtering procedure and describe the estimation method. Section 4 presents the estimation results and compares the models. The final section concludes.

2 The Model

We build on the unobserved components stochastic volatility model (hereafter UCSV model) of Stock and Watson (2007) by modifying the volatilities processes and allowing regime switching. In the standard UCSV model, the observed inflation π_t is equal to unobserved inflation mean τ_t which follows a random walk process plus a serially uncorrelated transitory component (measurement error). The log of volatilities of the permanent mean innovation (permanent shock) and the transitory component innovation (transitory shock) follow random walk processes.

We extend the standard UCSV model by allowing that the log volatilities of the permanent and transitory shocks evolve as independent AR(1) processes with regime switching coefficients. The regime $s(t) \in \{1, ..., S\}$ follows a Markov chain with transition matrix Λ . Our regime switching unobserved component stochastic volatility model (hereafter RS-UCSV) is summarized as follows:

$$\pi_t = \tau_t + \exp(\sigma_{\pi,t})\varepsilon_{\pi,t},\tag{1}$$

$$\tau_t = \tau_{t-1} + \exp(\sigma_{\varepsilon,t})\varepsilon_{\tau,t},\tag{2}$$

$$\sigma_{\pi,t} = \mu_{\pi,s(t)} + \rho_{\pi,s(t)} \left(\sigma_{\pi,t-1} - \mu_{\pi,s(t)} \right) + \gamma_{s(t)} \eta_{\pi,t}, \tag{3}$$

$$\sigma_{\tau,t} = \mu_{\tau,s(t)} + \rho_{\tau,s(t)} \left(\sigma_{\tau,t-1} - \mu_{\tau,s(t)} \right) + \gamma_{s(t)} \eta_{\tau,t}$$
(4)

$$\lambda_{ij} = \mathbb{P}[s_t = i | s_{t-1} = j] \tag{5}$$

where $\mu_{\pi,s}$, $\mu_{\tau,s}$, $\rho_{\pi,s}$, $\rho_{\tau,s}$, and γ_s are regime dependent coefficients, $\varepsilon_{\pi,t}$, $\varepsilon_{\tau,t}$, $\eta_{\pi,t}$, and $\eta_{\tau,t}$ are independent standard normal random variables. Parameter γ_s controls the smoothness of the stochastic volatilities processes. The parameter space includes the regime dependent coefficients and transition probabilities λ_{ij} . The total number of parameters is $S^2 + 4S$.

Our RS-UCSV model nests the standard UCSV model. When the number of regimes is one or

when the AR(1) coefficients are one across regimes, our model becomes the standard UCSV model:

$$\pi_t = \tau_t + \exp(\sigma_{\pi,t})\varepsilon_{\pi,t},\tag{6}$$

$$\tau_t = \tau_{t-1} + \exp(\sigma_{\varepsilon,t})\varepsilon_{\tau,t},\tag{7}$$

$$\sigma_{\pi,t} = \sigma_{\pi,t-1} + \gamma \eta_{\pi,t},\tag{8}$$

$$\sigma_{\tau,t} = \sigma_{\tau,t-1} + \gamma \eta_{\tau,t} \tag{9}$$

The unobserved components models with constant and stochastic volatilities have been used widely in modeling trend inflation. Chan et al. (2013) modify the UCSV model of Stock and Watson (2007) by imposing bounds on the volatilities of the innovations of the permanent and transitory components. They also allow for a time-varying degree of persistence of the transitory component. This mofification has the same flavor of some models used in Clark and Doh (2014), where they evaluate alternative models of trend inflation.

Another variant of unobserved components model is the first order integrated moving average IMA(1,1):

$$\Delta \pi_t = \varepsilon_t - \theta \varepsilon_{t-1},\tag{10}$$

where ε_t is i.i.d. with zero mean and variance σ_{ε}^2 . This model is analyzed, for example, in Nelson and Schwert (1977) and Barsky (1987). In fact, the IMA(1,1) is equivarlent to the unobserved component model with constant volatilities (UCCV) model. See Durbin and Koopman (2012) for details. We can imply the MA coefficient from the UCCV model using the equation: $\sigma_{\tau}/\sigma_{\pi} = (1+\theta)^2/\theta$. When allowing stochastic volatilities as in the standard UCSV model of Stock and Watson (2007), the MA coefficient of the IMA(1,1) model is no longer constant. Instead, it varies over time.

We are aware of the fact that there are many other forecasting models, both univariate and multivariate, that incorporate unobserved stochastic trend inflation. However, in this paper we focus on three types of models: the IMA(1,1) model, the UCSV model, and the RS-UCSV model. The reason is two fold. First, both Stock and Watson (2007) and Clark and Doh (2014) show that the UCSV model is among the best for forecasting future inflation. Second, we would like to see the role of regime switching in explaining inflation dynamics beyond what is produced by the UCSV model.

3 Particle filter and likelihood

Stock and Watson (2007) estimate the standard UCSV model using Markov Chain Monte Carlo (MCMC). Creal (2012) points out that their model has no static parameters that need to be estimated and suggests using the particle filter. In addition, the distribution of the unobserved permanent component τ_t can be calculated exactly by Kalman filter because conditional on the volatilities $\sigma_{\pi,t}$ and $\sigma_{\tau,t}$ the sub-model (1), (2) is a linear Gaussian state space model. We build on this approach and augment the particle filter with Markov regime switching filter of Hamilton (1989), which also calculates the distribution of the regime exactly.

The filter extracts the information about the regime and state variables of the model from data on the inflation rate π_t . The particle filter is implemented using the sequential importance sampling (SIS) and the Markov regime switching filter. The distribution of τ_t , $\sigma_{\pi,t}$ and $\sigma_{\tau,t}$ is approximated with the particle representation. In particular, we denote the vector of state variables $z_t \equiv (\tau_t, \sigma_{\pi,t}, \sigma_{\tau,t})$ and the vector of parameters $\theta = (\{\mu_{\pi,s}, \mu_{\tau,s}, \rho_{\pi,s}, \rho_{\tau,s}, \gamma_s\}_{s=1}^S)$. The filter calculates the posterior probability of regime $p(s_t|\pi_{1:t})$ and an approximation of the posterior density $p(z_t|\pi_{1:t}, s_t = s)$ by a sample of N particles $\{z_{s,t}^{(i)}\}_{i=1}^N$ and their corresponding weights $\{w_{s,t}^{(i)}\}_{i=1}^N$ in regime s, i.e.

$$p(z_t|\pi_{1:t}, s_t = s) \approx \sum_{i=1}^N w_{s,t}^{(i)} \delta(z_t - z_{s,t}^{(i)}),$$

where δ is the Dirac delta function and $\sum_{i} w_{s,t}^{(i)} = 1$. Thus, the moments of the distribution at time t conditional on regime s are

$$\frac{1}{N} \sum_{i=1}^{N} m(z_{s,t}^{(i)}) w_{s,t}^{(i)} \xrightarrow{a.s.} \mathbb{E}[m(z_t) | \pi_{1:t}, s_t = s].$$

where $m(\cdot)$ is a real-valued function.

The SIS method places a unit probability mass on each path of particles that has already been simulated in the previous iterations. The draws of particles are made sequentially from conditional distribution keeping existing particle paths fixed. Therefore, the joint importance distribution of the state and regime in the SIS approximation is factored as follows:

$$f(z_{0:t}, s_{0:t}|\pi_{1:t}; \theta) = f(z_t, s_t|z_{0:t-1}, s_{0:t-1}, \pi_{1:t}; \theta) f(z_{0:t-1}, s_{0:t-1}|\pi_{1:t-1}; \theta).$$
(11)

The joint conditional density the state and regime at time t is given by

$$f(z_t, s_t | z_{0:t-1}, s_{0:t-1}, \pi_{1:t}; \theta) = f(z_t | z_{0:t-1}, s_{0:t}, \pi_{1:t}; \theta) p(s_t | \pi_{1:t}; \theta)$$
(12)

A new set of particles is drawn from the importance distribution $f(z_t|z_{0:t-1}, s_{0:t-1}, s_t = s, \pi_{1:t}; \theta)$ conditional on $s_t = s$. The importance weights follow $w_{s,t} \propto w_{t-1}\tilde{w}_{s,t}$ where the incremental importance weight $\tilde{w}_{s,t}$ is given by

$$\tilde{w}_{s,t} = \frac{p(\pi_t | z_t, s_t = s; \theta) p(z_t | z_{t-1}, s_t = s; \theta)}{f(z_t | z_{0:t-1}, s_{0:t-1}, s_t = s, \pi_{1:t}; \theta)}.$$
(13)

In the numerator of (13), the first factor, $p(\pi_t|z_t, s_t; \theta)$, is the probability density of observation π_t conditional on state z_t and regime s_t and the second factor, $p(z_t|z_{t-1}, s_t; \theta)$, is the probability density of the state transition under regime s_t . Weights w_{t-1} are obtained in the resampling step by collapsing the states. This step allows us to avoid S-fold increase in the number of cases at each iteration as discussed in Kim (1994).

The posterior (updated) probability of regime is calculated by the Bayes' law:

$$p(s_t|\pi_{1:t};\theta) = \frac{p(\pi_t|s_t, \pi_{1:t-1};\theta)p(s_t|\pi_{1:t-1};\theta)}{p(\pi_t|\pi_{1:t-1})},$$
(14)

where the second factor in the numerator is the predicted probability of regime

$$p(s_t|\pi_{1:t-1};\theta) = p(s_{t-1}|\pi_{1:t-1})p(s_t|s_{t-1})$$
(15)

and

$$p(\pi_t | \pi_{1:t-1}) = \sum_{s=1}^{S} p(\pi_t | s_t = s, \pi_{1:t-1}; \theta) p(s_t = s | \pi_{1:t-1}; \theta).$$

As a result of this procedure, at each iteration we have a set of particles and their weights for each regime. In order to proceed to the next iteration, we collapse these sets into one set at a resampling step discussed later.

The likelihood of the state space model can be calculated as a by-product of the filtering. An

estimate of the contribution to the likelihood at time t is

$$p(\pi_t | \pi_{1:t-1}; \theta) = \int p(\pi_t | z_t; \theta) p(z_t | \pi_{1:t-1}; \theta) dz_t \approx \sum_{i=1}^N \hat{w}_{t-1}^{(i)} \tilde{w}_t^{(i)}$$
(16)

where $\hat{w}_{t-1}^{(i)} = w_{t-1}^{(i)} \left(\sum_{j=1}^{N} w_t^{(j)}\right)^{-1}$ is a normalized importance weight.

Resampling from importance density As it is well established in the literature, the sequential importance sampling suffers from the weight degeneracy, i.e. the probability masses of all but one particle converge to zero as t increases. Therefore, normalized importance weight of one particle converges to one while the normalized weights of the other particles converge to zero. As a result, as shown in Chopin (2004) the variance of the weights grows exponentially in time and the estimator becomes a function of a single draw. Gordon et al. (1993) show that a resampling step can prevent the weight degeneracy problem. However, traditional resampling procedures make the likelihood function discontinuous with respect to its parameters. The discontinuity problem persists even if the resampling is based on common (seeded) random numbers. The discontinuity (and non-smoothness) of the log-likelihood function can cause problems for optimizers, especially for gradient-based routines. See Hurzeler and Kunsch (2001) for discussion.

Malik and Pitt (2011) propose a resampling algorithm which ensures the smoothness of the likelihood function. The algorithm replaces a discrete CDF represented by particles and their weights with a continuous CDF based on piecewise linear approximations. Though resampling from a continuous CDF resolves the discontinuity and non-smoothness problem, the algorithm is applicable only for low-dimensional state space models and intensive computationally as it requires ordering of particles.

We propose an alternative resampling procedure which prevents weight degeneracy and preserves the continuity. The procedure makes use of kernel density estimate and importance sampling reweighing. In the resampling stage, we randomly draw new particles from a distribution chosen to be close to the distribution defined by the current particle swarm. Since the resampling distribution is not exact, we determine the particle weights as importance weights based on the kernel density of the current particles. Our particle filter with resampling from importance density is given as Algorithm 1 in Appendix A. For simplicity, we present the algorithm without the Kalman filter. We use the normal distribution as the importance distribution in the resampling step.

	1960:I-2004:IV	1960:I–2017:IV
(a) IMA(1,1): $\Delta \pi_t = \varepsilon_t - \theta \varepsilon_{t-1}$		
$\overline{ heta}_{arepsilon^2_arepsilon}$	$\begin{array}{ccc} 0.3037 & (0.0760) \\ 1.0650 \end{array}$	$\begin{array}{ccc} 0.3647 & (0.0643) \\ 1.0240 \end{array}$
(b) UCCV		
$ \sigma_{\pi} \\ \sigma_{\tau} \\ \text{Implied } \theta $	$\begin{array}{ccc} 0.4102 & (0.0002) \\ 0.9185 & (0.0004) \\ 0.2507 \end{array}$	$\begin{array}{ccc} 0.3718 & (0.0001) \\ 0.4171 & (0.0001) \\ 0.3622 \end{array}$

Table 1: Parameter estimates of IMA(1,1) and UCCV model

Notes: Panel (a) reports the parameter estimates of the integrated moving average IMA(1,1) model. Panel (b) reports estimates of the UCCV model and implied IMA(1,1) parameter θ . Standard errors are specified in the parentheses. The sample is quarterly U.S. GDP inflation from 1960:I to 2017:IV.

4 Results

4.1 Data

We compute inflation series using quarterly GDP deflator and CPI over the 1960:I to 2017:IV period. The data come from the Federal Reserve Economic Data website of the Federal Reserve Bank of St. Louis. We also use the personal consumption expenditure (PCE) index and core PCE in comparing forecasting ability among the models. Though inflation series are available from 1947:I, the earlier data is not reliable and usually excluded from the analysis.^{2,3}

4.2 The IMA(1,1) and unobserved components models

We first estimate the conventional benchmark models of inflation, the IMA(1,1), and the unobserved components model with constant volatilities, the UCCV model.⁴ Panel (a) in Table 1 reports the parameter estimates for the whole sample and 1960:I–2004:IV subsample for comparison with Stock and Watson (2007).

The coefficient θ is significant in both samples and increases from 0.30 to 0.36 when we expand the sample. As explained above, the IMA(1,1) model is equivalent to the UCCV model. If σ_{τ}^2

 $^{^{2}}$ Clark and Doh (2014) also use the data from 1960 for their analysis.

 $^{^{3}}$ Figure 7 in the appendix shows the GDP inflation series, its first difference, autocorrelation and partial autocorrelation functions for these series.

 $^{^{4}}$ The examination of the autocorrelation and partial autocorrelation functions of the inflation and the first difference in Figure 7 also suggests IMA(1,1) specification for our sample.

and σ_{π}^2 are variances of the innovations in the permanent and transitory components in the UCCV model, then $\lambda = (1+\theta)^2/\theta$, where θ is the implied IMA(1,1) parameter and $\lambda = \sigma_{\tau}/\sigma_{\pi}$ is the signal-to-noise ratio. Panel (b) in Table 1 reports the estimates of the standard errors of the error terms and implied IMA(1,1) coefficient θ from the UCCV model. From Table 1, the estimated θ in the IMA(1,1) model and the implied θ in the UCCV model are almost the same based on the 1960:I-2017:IV sample.

However, Stock and Watson (2007) argue that the MA coefficient of the IMA(1,1) should vary over time to fit structural changes in conducting monetary policy by the Fed. They actually reject the null that the MA coefficient is stable. Thus, to model inflation properly we should use the UCSV model, which is equivalent to the IMA(1,1) model with time-varying MA coefficient as variances σ_{τ}^2 and σ_{π}^2 evolve endogenously.

In this standard UCSV model, the dynamics of the volatilities processes depends on the innovation variance γ^2 and the initial conditions of the processes. Since the volatilities processes are random walks, they have no long-run steady state values that could serve as initial conditions. It is easy to see that if the initial value $\sigma_{\tau,0}$ is large relative to the initial value $\sigma_{\pi,0}$, the model attributes innovations in the observed inflation to changes in the unobserved permanent component rather than to the transitory component. As a result, the actual and filtered inflation series get closer. Moreover, if the error term volatilities in the volatilities equations (8) and (9) were allowed to be different, one could vary the speed of adjustment in the volatilities processes to control how the model recognizes future shocks. For example, if the volatility of the unobserved mean, $\sigma_{\tau,t}$, is large relative to the volatility of the measurement error $\sigma_{\pi,t}$, the model adjusts the unobserved mean τ_t . Otherwise, the model attributes the shock to the measurement error.

We follow Stock and Watson (2007) and estimate the UCSV model with one parameter (and one regime) to illustrate this point. We use the CPI inflation series and the error term variance in the volatilities equations, equal 0.2.⁵ Figure 1 plots the filtered inflation, implied IMA(1,1) parameter, and volatilities processes $\sigma_{\pi,t}$ and $\sigma_{\tau,t}$.⁶

We find that the model has relatively good fit in a subsample before 2005. After 2005, actual volatility of inflation significantly increases, and the filtered inflation process fails to accommodate the fluctuations. The reason for this dynamics is that starting from mid-1980s the volatility of

 $^{^{5}}$ We find that value of 0.1936 delivers the maximum value of likelihood function.

 $^{^{6}}$ To see the drawback of the standard UCSV model more clearly, we use the CPI inflation in this section. The result using the GDP inflation is robust. However, the volatility of inflation after 2005 is less pronounced.



Figure 1: Standard UCSV model

Notes: The figure plots actual and predicted inflation, the implied IMA(1,1) coefficient θ and filtered volatilities processes of the standard UCSV model. The sample is quarterly U.S. CPI inflation from 1960:I to 2017:IV.



Notes: The left panel plots the squared transitory shocks from the (one regime) local level model. The right panel plots the log of the squared transitory shocks from this model. The sample is quarterly U.S. GDP inflation from 1960:I to 2017:IV.

inflation had been declining substantially causing the volatility of the permanent shock to go down, see the right-bottom panel of Figure 1. At the same time, the volatility of the transitory shock had been growing. When the volatility of inflation soared after 2005, the unobserved mean inflation could not reconcile the oscillations, and the fluctuations were attributed to transitory shocks. Therefore, the unit-root volatility processes (8) and (9) cause strong dependence of the model on the previous dynamics.

Because the assumption that the volatilities processes are random walks makes the model strongly dependent on the past dynamics, we would like to test if this assumption holds with the data. Figure 2 plots squared residuals model and log of squared residuals from the local level model (with constant volatilities). The volatilities are modeled in Stock and Watson (2007) as independent random walks. While the logs of squared residuals reveal a slight low frequency trend, e.g. lower mean in the 1990s, the augmented Dickey-Fuller test with up to 9 lags rejects the unit root hypothesis at 5% significance level (p-value=0.01393 for the six-lags specification).

In addition to the limitation of the model in fitting the data after 2005, with the variance of the innovations in the volatilities equations equal 0.2 and a sample of 180 observations as in Stock and Watson (2007), ninety percent of simulated inflation series lie within the interval from approximately -405 to 405 while ten percent of the series get beyond this range. Therefore, the model could easily accommodate hyperinflation. For comparison if the variance is 0.01, then ninety percent of simulated series are contained in the interval approximately (-27, 27). Figure 3 shows



Notes: The figure plots 1,000 simulated paths of the observable variable π_t in the UCSV model over 180 periods. In the left panel the model parameter $\gamma = \sqrt{0.2}$, in the right panel $\gamma = 0.1$.

1,000 simulated inflation processes with variance equal 0.2 and 0.01. Arguably, the variance equal 0.2 is excessively large to explain the inflation series bounded between -10 and 20. As the parameter is constant, a large value is required to generate enough volatility to fit the model in the whole sample. In the next subsection, we model the volatilities using stationary and regime conditionally stationary processes.

4.3 RS-UCSV Model

In this section, we estimate our main model (1)-(5) and provide a comparison between alternative models with a focus on forecasting performance. We specify three hidden Markov regimes. Ngo and Isakin (2018) show that a conventional DSGE model has up to three steady-state regimes. In terms of the mean, one regime is associated with zero interest rate and deflation and the others – with positive interest rate. Importantly, inflation in one of these two positive-interest steady-state regimes is small and corresponds to the target inflation. The second positive-interest steady state has a much higher inflation rate, e.g. about 12 percent on the annual basis. In the inflation literature, it is well-known that different periods with different inflation means have different inflation volatilities, see Friedman (1977) and Ball (1992). For example, in the early 1980s when the inflation rate was high on average, the inflation was also very volatile. This paper sheds light on the volatility of these three regimes.

As discussed before, we assume that the coefficients in the volatilities equations (3) and (4)

(a) RS-1	UCSV model	with $AR(1)$ conditional	l volatilities		
$\mu_{\pi,1}$	-2.7520	(0.1399)	γ_1	0.0001	(0.0000)
$\mu_{ au,1}$	-2.8751	(0.0408)	γ_2	0.0001	(0.0001)
$\mu_{\pi,2}$	-0.6448	(0.0309)	γ_3	0.2622	(0.0242)
$\mu_{ au,2}$	0.1153	(0.0248)	λ_{12}	0.0001	(0.0001)
$\mu_{\pi,3}$	0.5667	(0.7503)	λ_{13}	0.0183	(0.0015)
$\mu_{ au,3}$	0.9026	(0.2065)	λ_{21}	0.0319	(0.0034)
$\rho_{\pi,1}$	0.7942	(0.0428)	λ_{23}	0.0001	(0.0001)
$\rho_{\tau,1}$	-0.0086	(0.0079)	λ_{31}	0.0098	(0.0011)
$\rho_{\pi,2}$	-0.4236	(0.1509)	λ_{32}	0.0051	(0.0022)
$\rho_{\tau,2}$	0.9990	(0.0277)			
$\rho_{\pi,3}$	0.9990	(0.0327)			
$ ho_{ au,3}$	0.7629	(0.0534)			

Table 2: Parameter estimates of models with three regime

(b) RS-UCCV model

$\gamma_{\pi,1}$	0.0400	(0.0304)	λ_{12}	0.0762	(0.0708)
$\gamma_{ au,1}$	0.0410	(0.0483)	λ_{13}	0.0000	(0.0000)
$\gamma_{\pi,2}$	0.4733	(0.0845)	λ_{21}	0.0219	(0.0233)
$\gamma_{\tau,2}$	0.0997	(0.0532)	λ_{23}	0.0075	(0.0076)
$\gamma_{\pi,3}$	0.1469	(0.3289)	λ_{31}	0.0000	(0.0000)
$\gamma_{ au,3}$	2.2157	(0.7704)	λ_{32}	0.0413	(0.0285)

Notes: Panel (a) reports the parameter estimates of the model (1)-(5) with three regimes. Panel (b) reports the parameter estimates of the three-regime UCSV model in which the volatilities are constants conditional on the regime. The diagonal elements of the transition matrix are calculated as $\lambda_{ii} = 1 - \sum_{j \neq i} \lambda_{ij}$. Standard errors are specified in the parentheses. The sample is quarterly U.S. GDP inflation from 1960:I to 2017:IV.

including coefficient γ_s are regime dependent and coefficient γ_s is the same for both volatilities processes $\sigma_{\tau,t}$ and $\sigma_{\pi,t}$. The last assumption is for parsimony only. Panel (a) of Table 2 reports the parameter estimates of our main model (1)-(5) with three regimes using the GDP inflation. The model has 21 parameters $\mu_{\pi,s}$, $\mu_{\tau,s}$, $\rho_{\pi,s}$, $\rho_{\tau,s}$, γ_s for s = 1, 2, 3 and non-diagonal elements λ_{ij} of the transition matrix Λ . The diagonal elements are calculated such that the probabilities in each row sum up to one.

Table 2 shows that the three regimes are distinct: low volatility (regime 1), moderate volatility (regime 2), and high volatility (regime 3). The average volatility of the transitory shock is smallest in regime 1 and largest in regime 3. Interestingly, the volitility of the transitory shock is quite persistent in regime 3, which is also the most persistent regime. The average volatility of the permanent shock is also smallest in regime 1 and largest in regime 3. The permanent shock is also smallest in regime 1 and largest in regime 3. The permanent shock is also smallest in regime 1 and largest in regime 3. The permanent shock volatility is very persistent in the second regime.



Figure 4: Three-regime model with AR(1) conditional volatilities – GDP inflation

Notes: The actual and predicted inflation, the implied IMA(1,1) coefficient θ , and filtered log volatilities of the transitory and permanent shocks. The volatilities are modeled as AR(1) processes conditional on the regime. The shaded areas indicate the likeliest regime in each quarter: the second regime is light grey, the third is dark grey. The sample is quarterly U.S. GDP inflation from 1960:I to 2017:IV.

Figure 4 plots actual and predicted inflation, the implied IMA(1,1) coefficient θ , and filtered log volatilities of the transitory and permanent shocks. In addition, Figure 4 shows the prevailing regimes over the sample.⁷ As follows from the estimates of the (quarterly) transition probabilities λ_{ij} and filtered regime probabilities, the regimes are quite persistent. The most volatile third regime prevails from 1972:II to 1986:IV with two short occurrences afterwards. The most stable first regime has the longest spell from 1990:III to to 2007:I. Our results mostly support findings in Stock and Watson (2007) of the changes in the inflation process. Moreover, the results validate the evidence for a structural shift in the early 1990s borne out by Clark (2006). In addition, we find that after 2007:II the inflation predominantly stays in the regime with medium permanent shock volatility and high transitory shock volatility.

For comparison, we also estimate a model where the volatilities of the transitory and permanent shocks are constants conditional on the regime. It implies that $\rho_s = 0$ and $\gamma_s = 0$ for s = 1, 2, 3in model (1)-(5). Overall, the model has twelve parameters: six error term volatilities and six free probabilities in the regime transition matrix. Panel (b) of Table (2) reports the estimates of

⁷Figure 8 in the appendix plots the filtered probabilities of the regimes for the RS-UCSV model.



Figure 5: Three-regime model with constant conditional volatilities: RS-UCCV model

Notes: Filtered and predicted value of inflation in the model where the volatilities of the transitory and permanent shocks are constants conditional on the regime. The shaded areas indicate the likeliest regime in each quarter: the second regime is light gray, the third is dark gray. The sample is quarterly U.S. GDP inflation from 1960:I to 2017:IV.

the parameters. Figure 5 plots filtered and predicted inflation.⁸ To explore the robustness of our specification, we estimate a model with regime switching MA(1) volatilities processes and find that the coefficients of the lagged error terms in all regimes are insignificant. The estimates of the filtered state variables that are close to those for the model with constant volatilities conditional on the regime.

For a robustness check, we estimate the model using the CPI inflation data. Figure 6 plots inflation and its one-step ahead forecast, the implied IMA(1,1) coefficient, and filtered log volatilities of the transitory and permanent shocks. It is worth noting that in the late 2000s and, especially, post Great Recession period CPI inflation series exhibits significantly higher volatility compared to the GDP inflation data. The model recognizes this period as the most volatile third regime. The presence of the regime switching allows the model to produce very different dynamics of the filtered inflation in this period compared to the UCSV model of Stock and Watson (2007); see Figure 1 for comparison.

We now turn our attention on the forecasting ability of the models. We use two indicators to compare the models. The first indicator is the in-sample mean squared forecasting error (MSFE)

 $^{^{8}}$ Figure 9 in the appendix plots the filtered probabilities of the regimes for the RS-UCCV model.



Figure 6: Three-regime model with AR(1) conditional volatilities (RS-UCSV model) – CPI inflation

Notes: The top four panels plot the actual and predicted inflation, the implied IMA(1,1) coefficient θ , and filtered log volatilities of the transitory and permanent shocks. The volatilities are modeled as AR(1) processes conditional on the regime. The shaded areas indicate the likeliest regime in each quarter: the second regime is light gray, the third is dark gray. The bottom three panels plot the filtered probabilities of the regimes. The sample is quarterly U.S. CPI inflation from 1960:I to 2017:IV.

calculated in the whole sample. The second criterion is the out-of-sample MSFE calculated for the last five years of the sample. To calculate this indicator, we estimate the model using the 1960:I-2012:IV period and use one-step ahead forecasts in the 2013:I-2017:IV period. We do not perform rolling estimation because multiple estimation of our RS-UCSV model is computationally cumbersome.

Table 3 summarizes the forecasting performance of the models using different inflation series: GDP inflation, CPI inflation, PCE inflation, and core CPE inflation. Based on in-sample MSFEs, our RS-UCSV model dominates the other models. It produces the smallest in-sample RMSFEs using the CPI, PCE, and core PCE inflation. For the GDP inflation, the RS-UCSV model is as good as the RS-UCCV model. The RS-UCSV model also produces better out-of-sample RMSFEs using the GDP, CPI, PCE inflation series. In sum, the models with regime switching outperform the model without regime switching both in- and out-of-sample.

	Naive	IMA(1,1)	UCSV	RS-UCSV	RS-UCCV
GDP inflation					
In-sample MSFE	1.1410	1.0236	1.0142	0.9570	0.9464
Out-of-sample MSFE	0.9606	0.7620	0.6729	0.6593	0.6658
CPI inflation					
In-sample MSFE	4.5074	3.8069	3.5028	3.4936	3.9198
Out-of-sample MSFE	3.8826	2.8466	2.6366	2.5084	2.6816
PCE inflation					
In-sample MSFE	2.1491	1.9157	1.8405	1.8277	1.9897
Out-of-sample MSFE	1.4919	1.2991	1.2611	1.1801	1.2552
PCE core inflation					
In-sample MSFE	0.7277	0.6585	0.6852	0.6537	0.6961
Out-of-sample MSFE	0.2294	0.1926	0.4628	0.4525	0.4505

Table 3: Forecasting Performance of the Models

Notes: The Table reports in-sample and out-of-sample mean squared forecast errors. In-sample MSFE's are calculated in the whole sample period. Out-of-sample MSFEs are calculated for the 2013:I–2017:IV period. The naive one-step ahead forecast equal to the previous value, i.e. $\pi_t^f = \pi_{t-1}$. The RS-UCCV model is the regime switching local level model with constant volatilities conditional on the regime. The samples contain the quarterly inflation over the 1960:I–2017:IV period.

5 Conclusion

We study the U.S. inflation dynamics using the random walk plus noise model with regime switching stochastic volatilities of the permanent and transitory shocks. The model is equivalent to the integrated moving average IMA(1,1) model with time-varying coefficient θ of the MA term. We assume that the volatility processes are stationary conditional on the regime. Since the values of the volatilities determine coefficient θ in each period, the coefficient is also stationary conditional on regime. We use maximum likelihood estimator based on the particle filter with regime switching. An important contribution is the resampling procedure that preserves the continuity of the likelihood function. We estimate the model with three regimes and find that the regime switching autoregressive representation of volatilities produces superior inflation forecast.

References

- Aruoba, B., Cuba-Borda, P., and Schorfheide, F. (2018). Macroeconomic dynamics near the zlb: A tale of two countries. *The Review of Economic Studies*, 85(1):87–118.
- Ball, L. (1992). Why does high inflation raise inflation uncertainty. Journal of Monetary Economics, 29(3):371–388.
- Ball, L. and Cecchetti, S. (1990). Inflation and uncertainty at long and short horizons. Brookings Papers on Economic Activity, 21(1):215–254.
- Barsky, R. (1987). The fisher hypothesis and the forecastability and persistence of inflation. *Journal* of Monetary Economics, 19(3):24.
- Chan, C. C. J., Koop, G., and Portter, S. M. (2013). A new model of trend inflation. Journal of Business and Economic Statistics, 31(1):94–106.
- Chopin, N. (2004). Central limit theorem for sequential monte carlo and its application to bayesian inference. *The Annals of Statistics*, 32:2385?2411.
- Clarida, R., Gali, J., and Gertler, M. (2000). Monetary policy rules and macroeconomic stability: Evidence and some theory. *The Quarterly Journal of Economics*, 115(1):147–180.

- Clark, T. (2006). Disaggregate evidence on the persistence of consumer price inflation. *Journal of* Applied Econometrics, pages 563–587.
- Clark, T. and Doh, T. (2014). Evaluating alternative models of trend inflation. *Journal of International Forecasting*, 30:426–448.
- Cogley, T., Primiceri, G. E., and Sargent, T. (2010). Inflation-gap persistence in the us. American Economic Journal: Macroeconomics, 2(1):43–69.
- Cogley, T. and Sargent, T. J. (2003). Drifts and volatilities: Monetary policies and outcomes in the post-wwii u.s. FRB Atlanta Working Paper No. 2003-25.
- Creal, D. (2012). A survey of sequential monte carlo methods for economics and finance. *Econo*metric Reviews, 31(3):245?296.
- Durbin, J. and Koopman, S. J. (2012). Time Series Analysis by State Space Methods. Oxford University Press, 2 edition.
- Engle, R. (1983). Estimates of the variance of u.s. inflation based upon the arch model. Journal of Money, Credit and Banking, 15(3):286–301.
- Friedman, M. (1977). Nobel lecture: inflation and unemployment. Journal of Political Economy, 85(3):451–472.
- Garcia, R. and Perron, P. (1996). An Analysis of the Real Interest Rate under Regime Shifts. The Review of Economics and Statistics, 78(1):111–125.
- Gordon, N., Salmond, D., and Smith, A. (1993). Novel approach to nonlinear/non-gaussian bayesian state estimation. *IEEE Proceedings F, Radar and Signal Processing*, 140(2):107–113.
- Hamilton, J. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2):357–84.
- Hurzeler, M. and Kunsch, H. (2001). Approximating and maximising the likelihood for a general state space model. In: Doucet, A., De Freitas, N., Gordon, N., eds. Sequential Monte Carlo Methods in Practice, pages 159–175.
- Kim, C.-J. (1994). Dynamic linear models with markov-switching. Journal of Econometrics, 60(1-2):1–22.

- Kim, C.-J. and Nelson, C. R. (1999). Has the u.s. economy become more stable? a bayesian approach based on a markov switching model of the business cycle. *Review of Economics and Statistics*, 81(4):608–661.
- Malik, S. and Pitt, M. K. (2011). Particle filters for continuous likelihood evaluation and maximisation. *Journal of Econometrics*, 165(2):190–209.
- Malmendier, U., Nagel, S., and Yan, Z. (2017). The making of hawks and doves: Inflation experiences on the fomc.
- Nelson, C. and Schwert, G. (1977). Short-term interest rates as predictors of inflation: On testing the hypothesis that the real rate of interest is constant. *American Economic Review*, 67(3):478–86.
- Ngo, P. V. and Isakin, M. A. (2018). Sunspot multiplicity and inflation dynamics. Working paper, Cleveland State University.
- Schwert, G. W. (1987). Effects of model specification on tests for unit roots in macroeconomic data. Journal of Monetary Economics, 20(1):73–103.
- Stock, J. H. and Watson, M. (2007). Why has u.s. inflation become harder to forecast? Journal of Money, Credit and Banking, 39(s1):3–33.
- Taylor, J. B. (1993). Discretion versus policy rules in practice. Carnegie-Rochester Conference Series on Public Policy, 39(1):195–214.
- Zhang, C. (2008). Structural instability of us inflation persistence. *Applied Economics Letters*, 15(14):1147–1151.

Appendix

5.1 Appendix A. Particle filter for the regime switching UCSV model

Algorithm 1 Particle filter for the regime switching UCSV model

I. At period t = 0:

- 1. Draw each set of particles $\{\tau_0^{(i)}\}_{i=1}^N$, $\{\sigma_{\pi,0}^{(i)}\}_{i=1}^N$, and $\{\sigma_{\tau,0}^{(i)}\}_{i=1}^N$ from their prior distributions. Set weights for each set of particles $w_{g,0}^{(i)} = 1/N$ for $g \in \{\tau, \sigma\pi, \sigma\tau\}$ and i = 1, ..., N. Set prior probabilities of regimes $\phi_0 = (\phi_0^1, ..., \phi_0^S)'$.
- II. For periods t = 1, ..., T:
 - 1. For each regime s = 1, .., S:
 - (a) Draw $\{\eta_{\pi,t}^{(i)}\}_{i=1}^{N}$ and $\{\eta_{\tau,t}^{(i)}\}_{i=1}^{N}$ from N(0,1) and calculate $\{\sigma_{\pi,t}^{(i),s}\}_{i=1}^{N}$ and $\{\sigma_{\tau,t}^{(i),s}\}_{i=1}^{N}$ according to (3) and (4) in regime s.
 - (b) Draw $\{\varepsilon_{\tau,t}^{(i)}\}_{i=1}^N$ from N(0,1) and calculate $\{\tau_t^{(i),s}\}_{i=1}^N$ according to (2) using $\{\sigma_{\tau,t}^{(i),s}\}_{i=1}^N$.
 - (c) Update the importance weights $w_{g,t}^{(i),s} \propto w_{g,t}^{(i)} N\left(\pi_t, \exp(\nu_t^{(i),s} + \sigma_{\pi,t}^{(i),s})\right)$ for $g \in \{\tau, \sigma\pi, \sigma\tau\}$ and i = 1, ..., N.
 - 2. Calculate the vector of predicted probabilities of regimes as $\hat{\phi}_t = \Lambda' \phi_{t-1}$.
 - 3. Calculate the likelihoods of regimes $\tilde{\phi}_t = \text{diag}(\bar{w}_t^1, \bar{w}_t^2, ..., \bar{w}_t^S)\hat{\phi}$ where $\bar{w}_t^s = \frac{1}{N} \sum_{i=1}^N w_t^{(i),s}$ and incremental log likelihood $\log(\tilde{\phi}'_t \mathbf{1})$.
 - 4. Obtain updated probabilities of regime as $\phi_t = \frac{\tilde{\phi}_t}{\sum_{s=1}^s \tilde{\phi}_t^s}$.
 - 5. Resample the particles $\{\{\tau_t^{(i),s}, \sigma_{\pi,t}^{(i),s}, \sigma_{\tau,t}^{(i),s}\}_{i=1}^N\}_{s=1}^S$ and form unconditional (averaged across the regimes) set of particles $\{\tau_t^{(i)}, \sigma_{\pi,t}^{(i)}, \sigma_{\tau,t}^{(i)}\}_{i=1}^N$.
 - (a) Stack each set particles $\{\{\tau_t^{(i),s}\}_{i=1}^N\}_{s=1}^S$, $\{\{\sigma_{\pi,t}^{(i),s}\}_{i=1}^N\}_{s=1}^S$, and $\{\{\sigma_{\tau,t}^{(i),s}\}_{i=1}^N\}_{s=1}^S$ for all regimes in $NS \times 1$ vectors $\tilde{\tau}_t$, $\tilde{\sigma}_{\pi,t}$, and $\tilde{\sigma}_{\tau,t}$ and calculate corresponding vectors of weights $\tilde{w}_{\tau,t}$, $\tilde{w}_{\sigma\pi,t}$, and $\tilde{w}_{\sigma\tau,t}$ such that particle *i* in regime *s* has weight $\frac{\hat{\phi}_t^s w_{g,t}^{s,(i)}}{\sum_{s=1}^S \sum_{i=1}^N \hat{\phi}_t^s w_{g,t}^{s,(i)}}$ for $g \in \{\tau, \sigma\pi, \sigma\tau\}$.
 - (b) Calculate means $\nu_{g,t}$ and variances $\varsigma_{g,t}$ for each set of particles $g \in \{\tau, \sigma\pi, \sigma\tau\}$.
 - (c) Sample N particles for each set of particles $\{\tau_t^{(i)}\}_{i=1}^N$, $\{\sigma_{\pi,t}^{(i)}\}_{i=1}^N$, and $\{\sigma_{\tau,t}^{(i)}\}_{i=1}^N$ from corresponding $N(\nu_{g,t},\varsigma_{g,t})$, where $g \in \{\tau, \sigma\pi, \sigma\tau\}$.
 - (d) Calculate the importance weights for each set of particles. For set $\{\tau_t^{(i)}\}_{i=1}^N$ the weights are

$$w_{\tau,t}^{(i)} = \frac{\sum_{j=1}^{NS} \tilde{w}_{g,t}^{(j)} K(\tau_t^{(i)} - \tilde{\tau}_t^{(j)} | h)}{\phi(\tilde{\tau}_t^{(i)} | \nu_{\tau,t}, \varsigma_{\tau,t})},$$
(17)

where $K(\cdot|H)$ is a kernel function with bandwidth h and $\phi(\cdot|\nu,\varsigma)$ is the normal PDF with mean ν and covariance matrix ς . For sets $\{\sigma_{\pi,t}^{(i)}\}_{i=1}^{N}$ and $\{\sigma_{\tau,t}^{(i)}\}_{i=1}^{N}$ the importance weights formulas are similar to (17).



Figure 7: The ACF and PACF of the GDP inflation and its first difference

Notes: The figure plots the autocorrelation and partial autocorrelation functions of the GDP inflation calculated over the 1960:I–2017:IV period.



Figure 8: Filtered probabilities of regimes for the GDP inflation: RS-UCSV model

Notes: The figure plots the filtered probabilities of the regimes. The sample is quarterly U.S. GDP inflation from 1960:I to 2017:IV.



Figure 9: Filtered probabilities of regimes for the GDP inflation: RS-UCCV model

Notes: Filtered probabilities of the regimes. The sample is quarterly U.S. GDP inflation from 1960:I to 2017:IV.