Inflation and Professional Forecast Dynamics: An Evaluation of Stickiness, Persistence, and Volatility*

Elmar Mertens† and **James M. Nason**‡

Current draft: June 14, 2015
First Draft: February 15, 2015

Abstract

This paper studies the joint dynamics of U.S. inflation and the average inflation predictions of the Survey of Professional Forecasters (SPF) on a sample running from 1968Q4 to 2014Q2. The joint data generating process (DGP) of these data consists of the unobserved components (UC) model of Stock and Watson (2007, "Why has US inflation become harder to forecast?," *Journal of Money, Credit and Banking* 39(S1), 3–33) and the sticky information (SI) forecast updating equation of Mankiw and Reis (2002, "Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve," *Quarterly Journal of Economics* 117, 1295–1328). We introduce time-varying inflation gap persistence into the Stock and Watson (SW)-UC model and a time-varying frequency of forecast updating into the SI law of motion. These models combine to produce a nonlinear state space model. This model is estimated using Bayesian tools grounded in the particle filter, which is an implementation of sequential Monte Carlo methods. The estimates reveal the data prefer the joint DGP of time-varying frequency of SI forecast updating and a SW-UC model with time-varying inflation gap persistence. The joint DGP produces estimates that indicate the inflation spike of 1974 was dominated by gap inflation, but trend inflation explains most of the inflation peak of the early 1980s. We also find the stochastic volatility (SV) of trend inflation exhibits negative comovement with the time-varying frequency of SI forecast updating while the SV and time-varying persistence of gap inflation often show positive comovement. Thus, the average SPF respondent is most sensitive to the impact of permanent shocks on the conditional mean of inflation.

JEL Classification Numbers: $E31$, $C11$, $C32$.

Key Words: inflation; professional forecasters; sticky information; particle filter; Bayesian estimation; Markov chain Monte Carlo; stochastic volatility; time-varying persistence.

†*e-mail*: elmar.mertens@frb.gov, *voice*: (202) 452–2916, *address*: Monetary Affairs Division, Board of Governors of the Federal Reserve System, 20th St. and Constitution Ave., N.W., Washington, D.C. 20551.

‡*e-mail*: jmnason@ncsu.edu, *voice*: (919) 513–2884, *address*: Department of Economics, Campus Box 8110, North Carolina State University, Raleigh, NC 27695–8110 and the Centre for Applied Macroeconomic Analysis.

*We thank Gregor Smith for several conversations that motivated this paper, Todd Clark and Monica Jain for useful comments, and participants at the 25th $(EC)^2$ Conference, Advances in Forecasting, held at the Universitat Pompeu Fabra, Barcelona, Spain, December 12 and 13, 2014 and the CIREQ Time Series and Financial Econometrics Conference, held at the Université de Montréal, Montréal, Québec, May 8 and 9, 2015. The views herein are those of the authors and do not represent the views of the Board of Governors of the Federal Reserve System or the Federal Reserve System. Updates to this paper will be posted at `www.elmarmertens.com/research/workingpapers/MertensNasonSI.pdf`.

1 Introduction

Central banks pay extraordinary attention to inflation expectations. A good reason for this preoccupation is that inflation expectations contain information about whether a central bank will be successful in stabilizing inflation. This information can be difficult for monetary policy makers to obtain because inflation expectations are not directly observed. Economists have responded by advancing methods to infer inflation expectations from realized inflation and some combination of financial market data, statistical and economic models, and inflation survey data.

This paper contributes to the literature studying the joint dynamics of realized inflation and the inflation predictions of professional forecasters. We tap a quarterly sample of inflation predictions of the Survey of Professional Forecasters (SPF) to extract the beliefs of its average respondent about the (in)stability, of the persistence, volatility, and stickiness of inflation. The SPF is an attractive source of information for studying the joint dynamics of realized and predicted inflation because, as Faust and Wright (2013) and Ang, Bekaert, and Wei (2007) observe, SPF inflation predictions often dominate model based out of sample forecasts. This superior forecasting performance suggests the average SPF participant's beliefs harbor information useful for assessing the joint data generating process (DGP) of realized and anticipated inflation.

The joint dynamics of realized inflation and average SPF inflation predictions combine two nonlinear models. These are the Stock and Watson (2007) unobserved components (UC) model of inflation and a nonlinear version of the Mankiw and Reis (2002) sticky information (SI) model. Mankiw and Reis (MR) propose a SI model in which price setting firms face incentives to update their information set infrequently rather than at every date as rational expectations (RE) maintains. A subset of firms are restricted to ground current decisions on past information because MR assume that at any moment of time only a time-invariant fraction of firms are granted access to current information.[1]

Coibion and Gorodnichenko (2015) adapt the MR-SI model to produce the h-step ahead inflation prediction of a SI forecaster, $F_t\pi_{t+h}$. This prediction is generated by Coibion and Gorodnichenko (CG) as a weighted average of the previous period's SI forecast, $F_{t-1}\pi_{t+h}$, and a RE inflation forecast, $\mathbf{E}_t\pi_{t+h}$, $F_t\pi_{t+h} = \lambda F_{t-1}\pi_{t+h} + (1-\lambda)\mathbf{E}_t\pi_{t+h}$, where λ is the constant SI parameter, $\lambda \in (0, 1)$. This scheme, which updates from $F_{t-1}\pi_{t+h}$ to $F_t\pi_{t+h}$, is a fixed coefficient-linear SI forecasting law of motion.

We innovate on the CG fixed coefficient-linear SI law of motion by investing the frequency at which

[1] In contrast, Sims (2003) constructs a dynamic optimizing model built on a primitive form of information processing in which agents react to shifts in the true DGP of the economy by smoothing their forecasts.

$F_t\pi_{t+h}$ is updated with a time-varying parameter (TVP), λ_t. The result is the nonlinear law of motion

$$F_t\pi_{t+h} \;=\; \lambda_{t-1}F_{t-1}\pi_{t+h} \;+\; (1-\lambda_{t-1})\mathbf{E}_t\pi_{t+h}, \tag{1}$$

where the SI-TVP frequency of updating is an exogenous random walk, $\lambda_t = \lambda_{t-1} + \sigma_\kappa\kappa_t$, $\kappa_t \sim \mathcal{N}(0,\,1)$, and $\lambda_t \in (0,\,1)$ for all dates t. The SI forecaster's information set includes λ_{t-1} when $F_{t-1}\pi_{t+h}$ is updated to $F_t\pi_{t+h}$ while κ_t is realized only during date t.

A motivation for including λ_{t-1} in the SI law of motion (1) is to uncover the changing beliefs the average SPF participant holds about the dynamics of realized inflation, π_t. Changes in these beliefs are embedded in movements of the average SPF participant's h-step ahead inflation prediction, $\pi^{SPF}_{t,t+h}$. We connect the observed $\pi^{SPF}_{t,t+h}$ to the latent SI forecast $F_t\pi_{t+h}$ by assuming $\pi^{SPF}_{t,t+h} = F_t\pi_{t+h} + \zeta_{t,t+h}$, where the measurement error $\zeta_{t,t+h} \sim \mathcal{N}(0,\,\sigma^2_{\zeta,h})$. This $\pi^{SPF}_{t,t+h}$-measurement error equation and the SI law of motion (1) form the mechanism by which shocks to the DGP of π_t are transmitted to $\pi^{SPF}_{t,t+h}$. This DGP produces $\mathbf{E}_t\pi_{t+h}$, which the SI law of motion (1) needs to update $F_t\pi_{t+h}$.

The true DGP of π_t is approximated with the Stock and Watson (2007) UC model. The Stock and Watson (SW-)UC model decomposes π_t into trend, τ_t, and gap, ε_t, inflation, which sets $\pi_t = \tau_t + \varepsilon_t$. Nonlinearities often appear in the SW-UC model as stochastic volatility (SV) in the innovation, η_t, to the random walk of τ_t and in ε_t. SV is induced by letting the variances of η_t and ε_t follow independent log random walks; see Stock and Watson (2010), Creal (2012), Mertens (2015), and Shephard (2013) among others. An implication is ε_t lacks own persistence. We extend the SW-UC model by endowing ε_t with a k-th order autoregresssion, AR(k) while maintaining its innovation, υ_t, is afflicted with SV. We report estimates when the AR parameters of ε_t are fixed or evolve as random walks.[2]

We estimate the joint DGP of the extended SW-UC model, the SI law of motion (1), and the SPF measurement error equation. Under this joint DGP, persistence in ε_t produces a term structure in π^{SPF}_{t+h}, $h = 1,\,\ldots,\,\mathcal{H}$, in which changes in λ_{t-1} create nonlinear state dynamics and volatility.[3] Estimates of this DGP shed light on whether movements in λ_{t-1} lead, lag, or are coincident to changes in the SV of τ_t, the SV of ε_t, or time-varying persistence in ε_t. If λ_{t-1} exhibits meaningful statistical and economic variation and this time variation moves with the state dynamics of the SW-UC model, we have evidence shifts in average SPF inflation predictions are attuned to the underlying factors driving inflation.

[2]Cogley, Primiceri, and Sargent (2010) reports estimates of time-varying inflation persistence using a TVP-VAR with SV that contains the U.S. unemployment rate, a short term nominal interest rate, and inflation.

[3]This approach to studying the joint dynamics of π_t and $\pi^{SPF}_{t,t+h}$ builds on Nason and Smith (2014). They estimate the joint dynamics of π_t and $\pi^{SPF}_{t,t+h}$ by integrating the SI-SPF regressions of Coibion and Gorodnichenko (2015) with the inflation survey term structure models of Kozicki and Tinsley (2012) and Mertens (2015).

Obtaining this evidence rests on separating the impact of changes in λ_{t-1} from the SVs of τ_t and ε_t, and time-varying persistence in ε_t on the joint dynamics of π_t and $\pi^{SPF}_{t,t+h}$. We locate information identifying these TVPs in the approximate exponentially weighted moving average (EWMA) recursions implied by the SW-UC model and the SI law of motion (1). The SW-UC model generates an approximation to a fixed coefficient EWMA recursion for $\mathbf{E}_t\pi_{t+1}$ because the discount on the history of π_t is a function of the SVs of τ_t and ε_t. Similarly, the SI law of motion (1) yields an approximate EWMA recursion for $F_t\pi_{t+h}$ in which the discount rate on the history of $\mathbf{E}_t\pi_{t+h}$ depends on the history of λ_{t-1}. Since shocks to these discount rates differ, fluctuations in $\mathbf{E}_t\pi_{t+h}$ and $F_t\pi_{t+h}$ have disparate sources. We rely on these differences to identify the mechanisms that transmit shocks to stickiness, volatility, and persistence into π_t and $F_t\pi_{t+h}$ by the joint DGP of the SW-UC model and the SI law of motion (1).

Another contribution of this paper is its estimator of the joint DGP of the SW-UC model, SI law of motion (1), the random walk of λ_t, and the SPF measurement error equation. Rao-Blackwellizing this nonlinear joint DGP creates a linear state space system on which the Kalman filter is used to generate the latent states τ_t and ε_t conditional on realizations λ_t, the SVs of τ_t and ε_t, and time-varying persistence of ε_t. We generate and evaluate the role of these realizations in this Kalman filter using sequential Monte Carlo (SMC) simulation methods (*i.e.*, the particle filter) that are reviewed by Creal (2012).

The conditionally linear likelihood of the SW-UC model, the SI law of motion (1), and the SPF measurement error equation are estimated on a sample that runs from $1968Q4$ to $2014Q2$. The sample equates π_t and $\pi^{SPF}_{t,t+h}$ with GNP or GDP deflator inflation and the average SPF predictions of this growth rate from $h = 1$- to 5-step ahead horizons, respectively.[4] The estimates show (*i*) the data prefer a joint DGP in which ε_t is driven by a TVP-AR(1), (*ii*) ε_t explains two-thirds or more of the spike in π_t of the 1973–1975 recession, but τ_t dominates the peak in π_t of the early 1980s, and (*iii*) $\pi^{SPF}_{t,t+h}$ is more sensitive to the impact of permanent shocks to the conditional mean of inflation.

The structure of the paper follows. The next section builds the nonlinear state space system of the SI law of motion (1) and the SW-UC model with SV and a persistent inflation gap. We also construct an example in this section of a state space model with a fixed coefficient-AR(k) inflation gap to motivate estimation of the joint DGP of π_t and $\pi^{SPF}_{t,t+h}$. Section 3 discusses estimation of the nonlinear state space system. Results appear in section 4. Section 5 concludes.

[4]The SPF contains average predictions of GNP or GDP inflation for a nowcast and forecasts up to four quarters ahead. However, these surveys are collected at the middle of every quarter. Since none of these predictions are based on full knowledge of current quarter inflation, we treat each survey as being conditioned only on data available through the end of the previous quarter. This identifies the average SPF nowcast, 1-step, ..., 4-step predictions with $\pi^{SPF}_{t,t+h}$, $h = 1, 2, \ldots, 5$, which implies these forecasts are made at the end of the previous quarter.

2 Statistical and Econometric Models

This section describes the statistical and economic models used to estimate the joint DGP of π_t and $\pi^{SPF}_{t,t+h}$, $h = 1, \ldots, \mathcal{H}$. The statistical model is the SW-UC model that has SV in innovations to trend and gap inflation, τ_t and ε_t, and generates persistence in ε_t with a TVP-AR(k). The economic model is the nonlinear SI law of motion (1) of $F_t\pi_{t+h}$. When combined, these models form a nonlinear state space model. We also discuss the joint DGP of π_t and $\pi^{SPF}_{t,t+h}$ when persistence in ε_t is a fixed coefficient AR(1). This example shows λ_{t-1} affects the transition dynamics and volatility of the unobserved state driving $\pi^{SPF}_{t,t+h}$ independent of whether persistence in ε_t is time-varying. In this case, changes in λ_{t-1} and the SV of innovations to τ_t and ε_t produce movements in $\pi^{SPF}_{t,t+h}$. When a TVP-AR(k) generates persistence in ε_t, it and λ_{t-1} interact to create additional movements in the transition dynamics of the state. This adds shocks to the persistence of ε_t to the list of factors that can induce fluctuations in $\pi^{SPF}_{t,t+h}$.

2.1 The SW-UC Model

The true DGP of π_t is approximated with the SW-UC model. Versions of this model are estimated by Stock and Watson (2010), Grassi and Prioietti (2010), Creal (2012), Mertens (2015), and Shephard (2013). These authors include nonlinearities in the SW-UC model by embedding SV in the innovations of τ_t, and ε_t. This paper includes an additional nonlinearity in the SW-UC model, which is time-varying persistence in the ε_t created by a TVP-AR(k).

These features of our specification of the SW-UC model are collected into its state space system

$$\pi_t = \tau_t + \varepsilon_t, \tag{2.1}$$

$$\tau_{t+1} = \tau_t + \varsigma_{\eta,t}\eta_{t+1}, \qquad \eta_{t+1} \sim \mathcal{N}(0,1), \tag{2.2}$$

$$\varepsilon_{t+1} = \sum_{j=0}^{k-1} \theta_{j,t}\varepsilon_{t-j} + \varsigma_{\upsilon,t}\upsilon_{t+1}, \qquad \upsilon_{t+1} \sim \mathcal{N}(0,1), \tag{2.3}$$

$$\ln \varsigma^2_{\ell,t+1} = \ln \varsigma^2_{\ell,t} + \sigma_\ell \xi_{\ell,t+1}, \qquad \xi_{\ell,t+1} \sim \mathcal{N}(0,1), \ \ell = \eta, \upsilon, \tag{2.4}$$

$$\theta_{j,t+1} = \theta_{j,t} + \sigma_{\phi,j}\phi_{j,t+1}, \qquad \phi_{j,t+1} \sim \mathcal{N}(0,1), \ j = 1, \ldots, k, \tag{2.5}$$

where the innovations to τ_t and ε_t, η_t and υ_t, are afflicted by SV in the form of independent log random walks $\varsigma_{\eta,t}$ and $\varsigma_{\upsilon,t}$, the TVP-AR(k) is also generated by the independent random walks $\theta_{1,t}$, ..., $\theta_{k,t}$, and the eigenvalues of the lag polynomial of these TVPs are outside the unit circle at every date t. The state space system of our modified SW-UC model consists of (2.1), which is the observation equation that decomposes π_t into τ_t and ε_t and the system of state equations, (2.2)–(2.5).

A special case of our SW-UC model gives a result traced to Muth (1960). Let $\theta_{1,t} = \ldots = \theta_{k,t} = 0$, $\sigma_\eta = \varsigma_{\eta,t}$, and $\sigma_\upsilon = \varsigma_{\upsilon,t}$ for all t. The result is a fixed coefficient UC model that has an IMA(1,1) reduced form, $(1-\mathbf{L})\pi_t = (1-\varpi\mathbf{L})\nu_t$, where $\mathbf{L}$ is the lag operator and the one-step ahead forecast error is $\nu_t = \eta_t + \varepsilon_t + \tau_t - \tau_{t-1|t-1}$.[5] The IMA(1, 1) implies the RE inflation updating equation, $\mathbf{E}_t\pi_{t+1} = \varpi\mathbf{E}_{t-1}\pi_t + (1-\varpi)\pi_t$, subsequent to some algebra. Shephard (2013) argues for replacing the fixed MA1 coefficient in the $\mathbf{E}_t\pi_{t+1}$ updating equation with the TVP ϖ_t, $\mathbf{E}_t\pi_{t+1} = \varpi_t\mathbf{E}_{t-1}\pi_t + (1-\varpi_t)\pi_t$ because of $\varsigma_{\eta,t}$ and $\varsigma_{\upsilon,t}$.[6] Backward iteration of the RE updating equation produces the approximate RE-EWMA recursion

$$\mathbf{E}_t\pi_{t+1} = \sum_{j=0}^{\infty} \mu_{\varpi,t-j}\left(\prod_{\ell=0}^{j}\varpi_{t-\ell}\right)\pi_{t-j}, \tag{3}$$

where $\mu_{\varpi,t} = (1-\varpi_t)/\varpi_t$. Thus, the SW-UC model with SV is consistent with an approximate RE-EWMA recursion (3) built from an updating equation that resembles the SI law of motion (1) of $F_t\pi_{t+h}$.

2.2 Average SPF, SI, RE Inflation Forecasts and λ_t

This section begins by reproducing the SPF measurement error equation, the nonlinear SI law of motion (1), and the random walk law of motion of λ_t. These elements form the system of equations

$$\pi^{SPF}_{t,t+h} = F_t\pi_{t+h} + \zeta_{t,t+h}, \tag{4.1}$$

$$F_t\pi_{t+h} = \lambda_{t-1}F_{t-1}\pi_{t+h} + (1-\lambda_{t-1})\mathbf{E}_t\pi_{t+h}, \tag{4.2}$$

$$\lambda_t = \lambda_{t-1} + \sigma_\kappa\kappa_t, \tag{4.3}$$

where $h = 1, \ldots, \mathcal{H}$, $\zeta_{t,t+h} \sim \mathcal{N}(0, \sigma^2_{\zeta,h})$, $\lambda_t \in (0, 1)$ for all dates t, and $\kappa_t \sim \mathcal{N}(0, 1)$. The system of equations (4.1)–(4.3) defines the mechanism through which shocks to λ_{t-1} and movements in SI and RE inflation forecasts generate fluctuations in $\pi^{SPF}_{t,t+h}$.

In part, this mechanism rests on the SI law of motion (4.2) approximating an EWMA recursion for $F_t\pi_{t+h}$. Subsequent to repeated backward iteration of and substitution into (4.2), the recursion is

$$F_t\pi_{t+h} = \sum_{j=0}^{\infty} \mu_{\lambda,t-1-j}\left(\prod_{\ell=0}^{j}\lambda_{t-1-\ell}\right)\mathbf{E}_{t-j}\pi_{t+h}, \tag{5}$$

where $\mu_{\lambda,t-1-j} = (1-\lambda_{t-1-j})/\lambda_{t-1-j}$. The SI-EWMA recursion (5) bounds $F_t\pi_{t+h}$ between the RE inflation

[5] An estimator of ϖ is built using the autocovariance functions (ACFs) of the IMA(1,1) and fixed coefficient SW-UC model; see Grassi and Proietti (2010) and Shephard (2013). At lags zero and one, the ACFs set $(1+\varpi^2)\,\sigma^2_\nu = \sigma^2_\eta + 2\sigma^2_\varepsilon$ and $-\varpi\sigma^2_\nu = -\sigma^2_\varepsilon$. Substitute for σ^2_ν to obtain the quadratic equation $\varpi^2 - \varpi\left(\sigma^2_\eta + 2\sigma^2_\varepsilon\right)/\sigma^2_\varepsilon + 1 = 0$. Its solution is $\varpi = \left[1 + 0.5\sigma^2_\eta/\sigma^2_\varepsilon\right] - (\sigma_\eta/\sigma_\varepsilon)\sqrt{1 + 0.25\sigma^2_\eta/\sigma^2_\varepsilon}$, given $\varpi \in (-1, 1)$ and $\sigma_\eta, \sigma_\varepsilon > 0$.

[6] Allowing for the TVP-AR(k) of ε_t creates higher order serial correlation in the reduced form of the SW-UC model.

forecast, $\lim_{\lambda_{t-1}\longrightarrow 0} F_t\pi_{t+h} = \mathbf{E}_t\pi_{t+h}$, and $\lim_{\lambda_{t-1}\longrightarrow 1} F_t\pi_{t+h} = \sum_{j=1}^{\infty}\mu_{\lambda,t-1-j}\left(\prod_{\ell=1}^{j}\lambda_{t-1-\ell}\right)\mathbf{E}_{t-j}\pi_{t+h}$. The former limit shows SI shutting down, as λ_t falls to zero because the history of RE inflation forecasts, $\mathbf{E}_{t-j}\pi_{t+h}$, $j \geq 1$, is discounted at a greater rate. Thus, the SI inflation forecast is updated to $\mathbf{E}_t\pi_{t+h}$ period by period. At the other extreme, less weight is placed on $\mathbf{E}_t\pi_{t+h}$ as λ_{t-1} rises to one. The SI forecaster neglects current information and instead engages the history of $\mathbf{E}_{t-j}\pi_{t+h}$, $j \geq 1$. Since $F_{t-1}\pi_{t+h}$ summarizes this history, only this information is used by the SI forecaster to set $F_t\pi_{t+h}$.[7]

Between these polar cases, movements in λ_{t-1} produce fluctuations in the rate at which the history of $\mathbf{E}_t\pi_{t+h}$ is discounted in the approximate SI-EWMA formula (5). These fluctuations offer identifying information to study the joint dynamics of π_t and $\pi_{t,t+h}^{SPF}$. The identifying information also relies on the sources of movements in the discount rate ϖ_t of the approximate RE-EWMA (3) grounded in the SV of the SW-UC model. Since shocks to $\varsigma_{\eta,t}$ and $\varsigma_{\upsilon,t}$ generate movements in the discount rate of the RE-EWMA (3) and is a potential source of variation in $\pi_{t,t+h}^{SPF}$ that is independent of shocks to λ_{t-1}, there are several sources of identifying information in the mechanism that transmits shocks into the term structure of $\pi_{t,t+h}^{SPF}$. The next two sections map the joint DGP of the SW-UC model (2.1)–(2.5) and the SI prediction mechanism of the average SPF respondent (4.1)–(4.3) into nonlinear state space systems which can be estimated on the sample $\left\{\pi_t,\ \pi_{t,t+h}^{SPF}\right\}_{t=1}^{T}$, $h = 1, \ldots, \mathcal{H}$.

2.3 An Example: The SW-UC-SI Model with a Fixed Coefficient AR(1) Inflation Gap

This section presents an example that explains and motivates our approach to evaluating the joint dynamics of π_t and $\pi_{t,t+h}^{SPF}$. The example turns the TVP-AR(k) of (2.3) into a fixed coefficient AR(1) by setting $k = 1$ and $\theta_1 = \theta_{1,t}$. Otherwise, the example leaves the rest of the SW-UC model untouched. Our motivation is to explain the impact of λ_t on the joint DGP of the SW-UC model (2.1)–(2.5) and the SI prediction mechanism of the average SPF respondent (4.1)–(4.3) separate from the additional complication of estimating a SW-UC model that includes a TVP-AR(k) inflation gap.

The state space system of the SW-UC model consists of the observation equation

$$\pi_t \ = \ \delta_x x_t, \tag{6.1}$$

which is equation (2.1), where $\delta_x = [1 \quad 1]$ and $x_t = [\tau_t \ \ \varepsilon_t]'$. Since the state vector x_t consists of the random walk (2.2) of trend inflation and the fixed coefficient AR(1) of gap inflation, $\varepsilon_{t+1} = \theta_1\varepsilon_t + \upsilon_{t+1}$,

[7] Lag the EWMA smoothing formula (5) by one period to obtain $F_{t-1}\pi_{t+h} = \sum_{j=0}^{\infty}\mu_{\lambda,t-2-j}\left(\prod_{\ell=0}^{j}\lambda_{t-2-\ell}\right)\mathbf{E}_{t-1-j}\pi_{t+h}$. Changing the indexes to i=j+1 and n=ℓ+1 yields $F_{t-1}\pi_{t+h} = \sum_{i=1}^{\infty}\mu_{\lambda,t-1-i}\left(\prod_{n=1}^{i}\lambda_{t-1-n}\right)\mathbf{E}_{t-i}\pi_{t+h}$, given λ_{t-1}=1.

the system of conditionally linear state equations evolves according to

$$x_{t+1} \;=\; \boldsymbol{\Theta}_1 x_t \;+\; \boldsymbol{\Xi}_t w_{t+1}, \tag{6.2}$$

where $\boldsymbol{\Theta}_1 = \begin{bmatrix} 1 & 0 \\ 0 & \theta_1 \end{bmatrix}$, $\boldsymbol{\Xi}_t = \begin{bmatrix} \varsigma_{\eta,t} & 0 \\ 0 & \varsigma_{\upsilon,t} \end{bmatrix}$, and $w_{t+1} = [\eta_{t+1} \;\; \upsilon_{t+1}]'$. The system of state equations (6.2) is linear conditional on the log random walk process (2.4) that generates $\varsigma^2_{\eta,t}$ and $\varsigma^2_{\upsilon,t}$. The SV shocks $\varsigma^2_{\eta,t}$ and $\varsigma^2_{\upsilon,t}$ are the source of nonlinearity in the SW-UC model when persistence in ε_t is tied to a constant AR1 coefficient, θ_1. This nonlinearity creates SV in the transition dynamics of system of state equations (6.2) because it appears in the covariance matrix of the system of state equations (6.2), $\mathbf{E}_t\left\{\boldsymbol{\Xi}_t w_{t+1} w'_{t+1} \boldsymbol{\Xi}'_t\right\} = \boldsymbol{\Xi}_t \boldsymbol{\Xi}'_t = \begin{bmatrix} \varsigma^2_{\eta,t} & 0 \\ 0 & \varsigma^2_{\upsilon,t} \end{bmatrix}$.

Given a constant θ_1, the law of iterated expectations (LIE) is available to compute h-step ahead RE inflation forecasts using the state space system (6.1) and (6.2) of the SW-UC model. This state space system gives the h-step ahead RE inflation forecast $\mathbf{E}_t \pi_{t+h} = \delta_x \boldsymbol{\Theta}_1^h x_t$. Given $h = 1, \ldots, \mathcal{H}$, these RE predictions form a term structure of inflation predictions because a AR(1) generates persistence in ε_t. Next, substitute the RE inflation forecast term structure into the SI approximate EWMA (5) to find $F_t \pi_{t+h}$ $= \delta_x \boldsymbol{\Theta}_1^h \sum_{j=0}^{\infty} \mu_{\lambda,t-1-j} \boldsymbol{\Theta}_1^j \left(\prod_{\ell=0}^{j} \lambda_{t-1-\ell}\right) x_{t-j}$. We show the infinite sum of the previous expression equals the SI forecast of x_t, which is the SI inflation term structure forecast

$$F_t \pi_{t+h} \;=\; \delta_x \boldsymbol{\Theta}_1^h F_t x_t, \quad h \;=\; 1, \ldots, \mathcal{H}, \tag{7}$$

which connects the h-step ahead SI inflation forecast to the SI nowcast of x_t.

We swap the h-step ahead state vector, x_{t+h}, for $F_t \pi_{t+h}$ in the SI law of motion (4.2), $F_t x_{t+h} =$ $(1 - \lambda_{t-1})\mathbf{E}_t x_{t+h} + \lambda_{t-1} F_{t-1} x_{t+h}$. The process of backward iteration yields the smoothing formula $F_t x_{t+h} = \boldsymbol{\Theta}_1^h \sum_{j=0}^{\infty} \mu_{\lambda,t-1-j} \boldsymbol{\Theta}_1^j \left(\prod_{\ell=0}^{j} \lambda_{t-1-\ell}\right) x_{t-j}$, where $\mathbf{E}_{t-j} x_{t+h} = \boldsymbol{\Theta}_1^{h+j} x_{t-j}$. Let $h = 0$ to obtain $F_t x_t$ $= \sum_{j=0}^{\infty} \mu_{\lambda,t-1-j} \boldsymbol{\Theta}_1^j \left(\prod_{\ell=0}^{j} \lambda_{t-1-\ell}\right) x_{t-j}$, which links $F_t \pi_{t+h}$ to $F_t x_t$. After pulling x_t from the previous infinite sum, we have $F_t x_t = (1 - \lambda_{t-1}) x_t + \sum_{j=1}^{\infty} \mu_{\lambda,t-1-j} \boldsymbol{\Theta}_1^j \left(\prod_{\ell=0}^{j} \lambda_{t-1-\ell}\right) x_{t-j}$, which by a change in index yields

$$\begin{aligned} F_t x_t \;&=\; (1 - \lambda_{t-1}) x_t + \lambda_{t-1} \boldsymbol{\Theta}_1 \sum_{i=0}^{\infty} \mu_{\lambda,t-1-i} \boldsymbol{\Theta}_1^i \left(\prod_{\ell=1}^{i+1} \lambda_{t-1-\ell}\right) x_{t-1-i} \\ &=\; (1 - \lambda_{t-1}) x_t \;+\; \lambda_{t-1} \boldsymbol{\Theta}_1 F_{t-1} x_{t-1}, \end{aligned} \tag{8}$$

where $F_{t-1} x_{t-1} = \sum_{j=0}^{\infty} \mu_{\lambda,t-1-j} \boldsymbol{\Theta}_1^j \left(\prod_{\ell=0}^{j} \lambda_{t-1-\ell}\right) x_{t-1-j}$. The system of equations (8) is a recursion that

generates the SI forecast of x_t, $F_t x_t$, given its own lag and conditions on the law of motion (6.2) of x_{t+1}, the independent log random walks of $\varsigma_{\eta,t+1}$ and $\varsigma_{\upsilon,t+1}$, and the random walk of λ_t.

We lead the recursion (8) by one period to generate SI forecasts of trend and gap inflation and substitute for x_{t+1} using the SW-UC model's state system (6.2). These actions produce the SI system of state equations

$$\begin{bmatrix} x_{t+1} \\ F_{t+1}x_{t+1} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\Theta}_1 & \mathbf{0}_{2\times 2} \\ (1-\lambda_t)\boldsymbol{\Theta}_1 x_t & \lambda_t \boldsymbol{\Theta}_1 \end{bmatrix} \begin{bmatrix} x_t \\ F_t x_t \end{bmatrix} + \begin{bmatrix} \boldsymbol{\Xi}_t \\ (1-\lambda_t)\boldsymbol{\Xi}_t \end{bmatrix} w_{t+1}, \tag{9.1}$$

which is conditionally linear on the independent bivariate SV log random walk process (16) and the random walk of λ_t, which is equation (4.3). Note that $\varsigma_{\eta,t}$, $\varsigma_{\upsilon,t}$, and λ_t are known when x_t and $F_t x_t$ are updated to x_{t+1} and $F_{t+1}x_{t+1}$, respectively. The associated system of observation equations is

$$\mathcal{Y}_t \equiv \begin{bmatrix} \pi_t \\ \pi^{SPF}_{t,t+1} \\ \vdots \\ \pi^{SPF}_{t,t+\mathcal{H}} \end{bmatrix} = \begin{bmatrix} \delta_x & \mathbf{0}_{1\times 4} \\ \mathbf{0}_{1\times 4} & \delta_x \boldsymbol{\Theta}_1 \\ \vdots & \vdots \\ \mathbf{0}_{1\times 4} & \delta_x \boldsymbol{\Theta}_1^{\mathcal{H}} \end{bmatrix} \begin{bmatrix} x_t \\ F_t x_t \end{bmatrix} + \begin{bmatrix} 0 \\ \zeta_{t,t+1} \\ \vdots \\ \zeta_{t,t+\mathcal{H}} \end{bmatrix}, \tag{9.2}$$

which draws on the observation equation (6.1) of the SW-UC model, the SPF measurement equation (4.1), and the system (7) that sets $F_t \pi_{t+h}$ equal to a term structure of h-step ahead SI forecasts of x_t.[8] Thus, the nonlinear state space system (9.1) and (9.2) eliminates $F_t \pi_{t+h}$ from the joint dynamics of π_t and $\pi^{SPF}_{t,t+h}$, $h = 1, \ldots, \mathcal{H}$.[9]

The time-varying frequency of updates to the SI inflation forecast alters the path of $F_{t+1}x_{t+1}$ in two ways. First, shocks to λ_t shift the transition dynamics of the state equations (9.1). These changes create permanent movements in the persistence of $F_{t+1}x_{t+1}$. The volatility of $F_{t+1}x_{t+1}$ is also affected by permanent shocks to λ_t. Thus, the fixed coefficient AR(1)-SW-UC model united with the nonlinear SI model (4.1)–(4.3) creates the opportunity for changes in the stickiness of $F_t \pi_{t,t+h}$ to interact with shifts in the volatility of the DGP of π_t to explain fluctuations in the term structure of π^{SPF}_{t+h}. The next section replaces the fixed coefficient inflation gap with the TVP-AR(k) of (2.3). We add this nonlinearity to the joint DGP of the SW-UC model and SI law of motion to ask whether the term structure of π^{SPF}_{t+h} responds to time-varying persistence in the inflation gap at any time during the last 45 years.

[8] Since $\zeta_{t,t+h}$ is an additive measurement error, it only effects the observation equations and not state dynamics.

[9] The state space model (9.1) and (9.2) is identified by separating time variation in $\mathbf{E}_t \pi_{t+h}$ using $\varsigma_{\eta,t}$ and $\varsigma_{\upsilon,t}$ while λ_t is responsible for these movements in $F_t \pi_{t,t+h}$. Alternatives to our identification are developed by Krane (2011), Coibion and Gorodnichenko (2012), and Jain (2013). Forecast revisions are central to their identification of the responses of professional forecasters to persistent shocks.

2.4 The State Space System with a TVP-AR(k) Inflation Gap

An inflation gap driven by the TVP-AR(k) of (2.3) complicates construction of the state space of the joint dynamics of π_t and π_{t+h}^{SPF}. In this case, the LIE cannot be engaged to calculate h-step RE inflation forecasts using the SW-UC model (2.1)–(2.5). We develop this idea starting with the observation equation of the state space system of the SW-UC model

$$\pi_t \;=\; \delta_{\mathbf{X}} \mathcal{X}_t, \tag{10.1}$$

which is equation (2.1), where the $3+2k$ row vector $\delta_{\mathbf{X}} = [1 \;\; 1 \;\; 0 \;\; \dots \;\; 0]$ is conformable with the state vector $\mathcal{X}_t = [\tau_t \;\; \varepsilon_t \;\; \varepsilon_{t-1} \;\; \dots \;\; \varepsilon_{t-k+1}]'$, which contains trend inflation, gap inflation, and its lags. The associated conditionally linear system of state equations

$$\mathcal{X}_{t+1} \;=\; \mathbf{\Theta}_t \mathcal{X}_t \;+\; \mathbf{\Upsilon}_t \mathcal{W}_{t+1}, \tag{10.2}$$

is the rest of the SW-UC state space, where $\mathbf{\Theta}_t = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 & 0 \\ 0 & \theta_{1,t} & \theta_{2,t} & \dots & \theta_{k-1,t} & \theta_{k,t} \\ \mathbf{0}_{k-1\times 1} & \mathbf{I}_{k-1} & & & & \mathbf{0}_{k-1\times 1} \end{bmatrix}$, the vector $\mathcal{D}_{\Xi} = [\varsigma_{\eta,t} \;\; \varsigma_{\upsilon,t} \;\; 0 \;\; \dots \;\; 0]$ contains the only nonzero elements of the diagonal matrix $\mathbf{\Upsilon}_t$ and $\mathcal{W}_{t+1} = [\eta_{t+1} \;\; \upsilon_{t+1} \;\; 0 \;\; \dots \;\; 0]'$. Thus, $\mathbf{E}_t\left\{\mathbf{\Upsilon}_t \mathcal{W}_{t+1} \mathcal{W}_{t+1}' \mathbf{\Upsilon}_t'\right\}$ is a diagonal matrix with its nonzero elements residing in the vector $\mathcal{D}_{\mathcal{W},t} = [\varsigma_{\eta,t}^2 \;\; \varsigma_{\eta,t}^2 \;\; 0 \;\; \dots \;\; 0]$. The first two elements of $\mathcal{D}_{\mathcal{W},t}$ represent the SV in τ_t and ε_t. This SV is the source of time-variation in the covariance matrix $\mathbf{\Upsilon}_t \mathbf{\Upsilon}_t' = \mathbf{E}_t\left\{\mathbf{\Upsilon}_t \mathcal{W}_{t+1} \mathcal{W}_{t+1}' \mathbf{\Upsilon}_t'\right\}$ and, thus, volatility in $\mathcal{X}_{t+1}$. There are also time-varying dynamics in the transition matrix, $\mathbf{\Theta}_t$, of the state system (10.2). The source of these dynamics is the TVP-AR(k) of ε_t. Nonetheless, the system of state equations (10.2) is linear conditional on $\varsigma_{\eta,t}$, $\varsigma_{\upsilon,t}$, and $\theta_{1,t}, \dots, \theta_{k,t}$, which are known at date t.

The state space system (10.1) and (10.2) is a challenge to evaluate. Part of the problem is to compute h-step ahead RE and SI forecasts of π_t and $\mathcal{X}_t$ when $\theta_{1,t}, \dots, \theta_{k,t}$ produce time-varying persistence in ε_t, which drives variation $\mathbf{\Theta}_t$. Since nonlinearities in $\mathbf{\Theta}_t$ rule out using the LIE to compute $\mathbf{E}_t \pi_{t+h}$, $F_t \pi_{t+h}$, and $F_t \mathcal{X}_t$, we appeal to two aspects of the anticipated utility model (AUM) to solve the problem. The AUM resurrects the LIE by assuming (i) agents are ignorant of the true DGP and (ii) treats the TVPs of the SW-UC model (2.1)–(2.5) and the SI law of motion (4.2) as fixed (locally) at each date t. Under the AUM, we hold the current state of these TVPs fixed when generating h-step ahead forecasts.[10]

[10] Cogley and Sbordone (2008) employ these assumptions to study the dynamics of trend and gap inflation within a TVP-new Keynesian Phillips curve. They note Kreps (1998) argues that agents engaging in these behaviors are acting rationally when seeing through to the true model is costly. These assumptions also result in decision making that is often close to Bayesian forecasting, according to Cogley and Sargent (2008).

The AUM assumptions are engaged to evaluate multi-period forecasts of $\mathbf{E}_t\pi_{t+h}$ and $F_t\pi_{t+h}$. The state space of the SW-UC model (10.1) and (10.2) produces the h-step ahead RE inflation forecast

$$\mathbf{E}_t\pi_{t+h} \;=\; \delta_{\mathbf{X}}\boldsymbol{\Theta}_{t|t}^{h}\mathcal{X}_{t|t}. \tag{11}$$

The subscripts on $\boldsymbol{\Theta}_{t|t}$ and $\mathcal{X}_{t|t}$ are held fixed at the date t state of the joint DGP of the SW-UC model (2.1)–(2.5) to be consistent with the AUM.

The next step is to find a recursion for the SI prediction of $\mathcal{X}_t$. This recursion links $F_t\pi_{t+h}$ to the state of the SW-UC models and the SI law of motion by creating a mechanism to update $F_t\mathcal{X}_t$ given $\mathcal{X}_t$ and $F_{t-1}\mathbf{X}_{t-1}$. Although the path of the backward-looking SI law of motion of $\mathcal{X}_t$ is also altered by changes in λ_{t-1}, the example in the previous section shows that shifts in λ_t are not an impediment in constructing a map from $F_t\mathcal{X}_t$ to $F_t\pi_{t+h}$. The problem is the AUM also tells us to track the history of $\boldsymbol{\Theta}_t$ when constructing the smoothing recursion that describes the path of $F_t\mathcal{X}_t$.

The map from $F_t\mathcal{X}_t$ to $F_t\pi_{t+h}$ is built starting from $F_t\mathcal{X}_{t+h} = (1-\lambda_{t-1})\mathbf{E}_t\mathcal{X}_{t+h} + \lambda_{t-1}F_{t-1}\mathcal{X}_{t+h}$, given ε_t is the TVP-AR(k) of (2.3). Subsequent to iterating backward this SI law of motion, the result is the approximate EWMA formula $F_t\mathcal{X}_{t+h} = \sum_{j=0}^{\infty}\mu_{\lambda,t-1-j}\left(\prod_{\ell=0}^{j}\lambda_{t-1-\ell}\right)\mathbf{E}_{t-j}\mathcal{X}_{t+h}$. Since $\mathbf{E}_{t-j}\mathcal{X}_{t+h} = \boldsymbol{\Theta}_{t-j|t-j}^{h+j}\mathcal{X}_{t-j|t-j}$, $F_t\mathcal{X}_{t+h} = \sum_{j=0}^{\infty}\mu_{\lambda,t-1-j}\left(\prod_{\ell=0}^{j}\lambda_{t-1-\ell}\right)\boldsymbol{\Theta}_{t-j|t-j}^{h+j}\mathcal{X}_{t-j|t-j}$. Next, set $h=0$ to obtain

$$F_t\mathcal{X}_t \;=\; \sum_{j=0}^{\infty}\mu_{\lambda,t-1-j}\left(\prod_{\ell=0}^{j}\lambda_{t-1-\ell}\right)\boldsymbol{\Theta}_{t-j|t-j}^{j}\mathcal{X}_{t-j|t-j}. \tag{12}$$

Under the hypothesis of the AUM, fastening together the EWMA formula (5) of $F_t\pi_{t+h}$, the RE inflation forecast (11), and the smoothing formula (12) of $F_t\mathcal{X}_t$ results in

$$F_t\pi_{t+h} \;=\; \delta_{\mathcal{X}}\boldsymbol{\Theta}_{t|t}^{h}F_t\mathcal{X}_t, \quad h \;=\; 1,\ldots,\mathcal{H}. \tag{13}$$

The term structure (13) joins $F_t\pi_{t,t+h}$ to the SI nowcast of $\mathcal{X}_t$, which connects $\pi_{t,t+h}^{SPF}$ to the SI forecasts of trend and gap inflation through the SPF measurement error equation (4.1). An implication is shocks to $\theta_{1,t},\ldots,\theta_{k,t}$ generate nonlinear fluctuations in $\pi_{t,t+h}^{SPF}$ as well as π_t.

We still need a recursion for the SI forecast of $\mathcal{X}_t$ when the AUM is invoked because of a TVP-AR(k) in ε_t. First, pull the first four terms out of the infinite sum of the smoothing formula (12)

$$\begin{aligned} F_t\mathcal{X}_t \;=\;& (1-\lambda_{t-1})\mathcal{X}_t + (1-\lambda_{t-2})\lambda_{t-1}\boldsymbol{\Theta}_{t-1|t-1}\mathcal{X}_{t-1|t-1} + (1-\lambda_{t-3})\lambda_{t-1}\lambda_{t-2}\boldsymbol{\Theta}_{t-2|t-2}^{2}\mathcal{X}_{t-2} \\ &+ \;\; (1-\lambda_{t-4})\lambda_{t-1}\lambda_{t-2}\lambda_{t-3}\boldsymbol{\Theta}_{t-3|t-3}^{3}\mathcal{X}_{t-3} + \sum_{j=4}^{\infty}\mu_{t-1-j}\left(\prod_{\ell=0}^{j}\lambda_{t-1-\ell}\right)\boldsymbol{\Theta}_{t-j|t-j}^{j}\mathcal{X}_{t-j|t-j}. \end{aligned}$$

By induction, the sequence of these terms point to the SI law of motion of $\mathcal{X}_t$

$$F_t\mathcal{X}_t \;=\; (1-\lambda_{t-1})\mathcal{X}_t + \lambda_{t-1}\boldsymbol{\Theta}_{t-1|t-1}F_{t-1}\mathcal{X}_{t-1}. \tag{14}$$

as agreeing with the smoothing formula (12). A key feature of the SI law of motion (14) of $F_t\mathcal{X}_t$ is that its updating relies on λ_{t-1} interacting with the lagged time-varying persistence in ε_t, which is summarized by $\boldsymbol{\Theta}_{t-1|t-1}$. The AUM assumption is responsible for this restriction on the SI law of motion (14) of $F_t\mathcal{X}_t$. This restriction is consistent with the SI forecaster holding the TVP-AR(k) fixed at date $t-1$ when updating from $F_{t-1}\mathcal{X}_{t-1}$ to $F_t\mathcal{X}_t$.

The system of state equations of the SW-UC model and the time-varying SI law of motions are the laws of motion of $\mathcal{X}_{t+1}$ and $F_{t+1}\mathcal{X}_{t+1}$. Stack the law of motion of $\mathcal{X}_{t+1}$ on top of the law of motion $F_{t+1}\mathcal{X}_{t+1}$, which are equations (10.2) and (14), to create the system of state equations

$$\mathcal{S}_{t+1} \;=\; \boldsymbol{\mathcal{A}}_t\mathcal{S}_t \;+\; \boldsymbol{\mathcal{B}}_t\mathcal{W}_{t+1}, \tag{15.1}$$

where $\mathcal{S}_t = [\mathcal{X}_t' \;\; F_t\mathcal{X}_t']'$, $\boldsymbol{\mathcal{A}}_t = \begin{bmatrix} \boldsymbol{\Theta}_t & \mathbf{0}_{m\times m} \\ \big(1-\lambda_t\big)\boldsymbol{\Theta}_t & \lambda_t\boldsymbol{\Theta}_t \end{bmatrix}$, m=1+k, and $\boldsymbol{\mathcal{B}}_t = \begin{bmatrix} \boldsymbol{\Upsilon}_t \\ (1-\lambda_t)\boldsymbol{\Upsilon}_t \end{bmatrix}$. The state equations (15.1) depict λ_t having an impact on transition dynamics and the volatility of the system similar to that seen in the state equations (9.1) of the previous section's example of a fixed coefficient AR(1) inflation gap. However, a TVP-AR(k) inflation gap introduces an additional nonlinearity into the transition dynamics of the system of state equations (15.1). The interaction of time-varying persistence in ε_t and λ_t is one more margin on which movements in $\mathcal{S}_t$ drive fluctuations in $\pi^{SPF}_{t,t+h}$.

Time-varying inflation gap persistence also appears in the system of observation equations driven by $\mathcal{S}_t$. These equations place the observation equation (10.1) of the SW-UC model on top of the SPF measurement error equations (4.1) combined with the SI nonlinear term structure (13), which is

$$\mathcal{Y}_t \;=\; \boldsymbol{\mathcal{C}}_t\mathcal{S}_t \;+\; \boldsymbol{\mathcal{D}}\mathcal{U}_t, \tag{15.2}$$

where $\boldsymbol{\mathcal{C}}_t = \begin{bmatrix} \delta_{\mathcal{X}} & \mathbf{0}_{1\times m} \\ \mathbf{0}_{1\times m} & \delta_{\mathcal{X}}\boldsymbol{\Theta}_t \\ \vdots & \vdots \\ \mathbf{0}_{1\times m} & \delta_{\mathcal{X}}\boldsymbol{\Theta}_t^{\mathcal{H}} \end{bmatrix}$, $\boldsymbol{\mathcal{D}} = \mathbf{I}_{\mathcal{H}+1}$, $\mathcal{U}_t = [0 \;\; \zeta_{t,t+1} \;\ldots\; \zeta_{t,t+\mathcal{H}}]'$, and the diagonal matrix $\boldsymbol{\Omega}_{\mathcal{U}} = \mathbf{E}\big\{\mathcal{U}_t\mathcal{U}_t'\big\}$.[11] In the system of observation equations (15.2), the response of $\pi^{SPF}_{t,t+h}$ to movements

[11] The uncorrelated measurement error could be made at the level of the output price deflator rather than to $\pi^{SPF}_{t,t+h}$

in Θ_t do not violate the AUM assumptions. Furthermore, shocks to Θ_t produce fluctuations in y_t in two ways. There is the direct impact these shocks have on $\pi^{SPF}_{t,t+h}$ in the observation system (15.2) through the term structure created by the TVP-AR(k) and the effect these shocks have by altering the transition dynamics of the system of state equations (15.1). Note, however, the state space system (15.1) and (15.2) is linear given realizations of $\varsigma_{\eta,t}$, $\varsigma_{\upsilon,t}$, $\theta_{1,t}, \ldots, \theta_{k,t}$, and λ_t.

3 Econometric Methods

We estimate the state space system (15.1) and (15.2) using Bayesian SMC methods, which involve a particle filter algorithm adapted from Creal (2012) and Herbst and Schorfheide (2014). Durbin and Koopman (2002), and Godsill, and Doucet, and West (2004) provide instructions to run a Monte Carlo smoothing simulator. These tools are also applicable to a state space system in which persistence in ε_t is generated by a fixed coefficient AR(k), as is the state space system (9.1) and (9.2) where $k = 1$.

3.1 Rao-Blackwellization of a Nonlinear State Space Model

Creal (2012) discusses a particle filter that relies on the Rao-Blackwellization process of Chen and Liu (2000). In this case, Rao-Blackwellization snaps the state vector of a nonlinear state space model in two. One vector of state variables are responsible for the nonlinearities in the state space system. Given a realization of this vector, the remaining state variables are generated by a linear state space model. Thus, the state space system (15.1) and (15.2) is consistent with applying the Rao-Blackwellization process to the joint DGP of the SW-UC model (2.1)–(2.5) and the SI prediction mechanism of the average SPF respondent (4.1)–(4.3), conditional on the latent state variables $\ln \varsigma^2_{\eta,t}$, $\ln \varsigma^2_{\upsilon,t}$, $\theta_{1,t}, \ldots, \theta_{k,t}$, and λ_t. An implication is that, conditional on these state variables, applying the Kalman filter to the state space system (15.1) and (15.2) analytically produces the distribution of s_t.

Analytic integration of the distribution of s_t endows the particle filter with greater numerical efficiency. The efficiency gains are obtained by applying the Kalman filter to the state space system (15.1) and (15.2) to generate s_t and not by simulating this state vector. However, the Kalman filter

as in the observation system (15.2). The alternative assumption implies cumulative inflation forecasts possess uncorrelated measurement errors. Cumulative inflation forecasts have correlated measurement errors when $\pi^{SPF}_{t,t+h}$ is endowed with an uncorrelated measurement error. Or results are similar using either measurement error assumption. We report estimates of the state space model (15.1) and (15.2) because log likelihoods are greater under the assumption the uncorrelated measurement error, $\xi_{t,t+h}$, is attached to $\pi^{SPF}_{t,t+h}$. However, estimates based on the alternative measurement error assumption are available on request.

needs realizations of $\ln \varsigma_{\eta,t}^2$, $\ln \varsigma_{\upsilon,t}^2$, $\theta_{1,t}, \ldots, \theta_{k,t}$, and λ_t to accomplish this task. The Rao-Blackwellized particle filter solves this problem by simulating M realizations or particles of these state variables on which the Kalman filter is run to create the analytic distribution of $\mathcal{S}_t$.[12]

The particle filter produces filtered and smoothed estimates of the variables in $\mathcal{S}_t$ using the linear state space system (15.1) and (15.2) conditional on the data $\mathcal{Y}_t$ and the state variables $\ln \varsigma_{\eta,t}^2$, $\ln \varsigma_{\upsilon,t}^2$, $\theta_{1,t}, \ldots, \theta_{k,t}$, and λ_t. Since these TVPs evolves as a independent random walks, initial conditions and volatility scale parameters are needed to generate synthetic samples of these random walks. The scale volatilities are collected in $\Psi = [\sigma_\eta \;\; \sigma_\upsilon \;\; \sigma_{\zeta,1} \;\; \ldots \;\; \sigma_{\zeta,\mathcal{H}} \;\; \sigma_{\phi,1} \;\; \ldots \;\; \sigma_{\phi,k}]'$.[13] At the moment, calibration methods are applied to obtain values for the elements of Ψ.[14]

The calibration of Ψ is summarized in table 1. Table 1 shows the scale of the volatility of innovations to the SVs of trend and gap inflation, σ_η and $\sigma_{\upsilon,t}$, equal 0.2, which matches values used by Stock and Watson (2007). There is no a prior evidence to calibrate the scale of the volatility of the innovations to λ_t, $\sigma_{\zeta,h}$, $h = 1, \ldots, 5$. We settle on $\sigma_\kappa = 0.05$ to scale the volatility of the innovation, κ_t, to λ_t subsequent to running simulation experiments on a grid of values for σ_κ. The same process is used to fix the volatilities of $\xi_{t,t+h}$, $\sigma_{\zeta,h}$, which yields $\sigma_{\zeta,h} = 1.0$ for all h. Table 1 also indicates that we estimate models with a TVP-AR(1) and a TVP-AR(2) for ε_t. The scale of the volatilities of innovations to $\theta_{1,t}$ or $\theta_{1,t}$ and $\theta_{2,t}$ are $\sigma_{\phi,1} = 0.05$ or $\sigma_{\phi,1} = 0.10$ and $\sigma_{\phi,2} = 0.05$ using a similar procedure.

We also have to calibrate initial conditions for the random walk processes of $\mathcal{V}_t$. The initial conditions of $\ln \varsigma_{\eta,0}^2$ and $\ln \varsigma_{\upsilon,0}^2$ are sampled using $\ln \varsigma_{\eta,0}^2 \sim \mathcal{N}(-[2.5 + 2 \ln 2], 10)$ and $\ln \varsigma_{\upsilon,0}^2 \sim \mathcal{N}(-2.5, 10)$. Since λ_t is restricted to the open interval between zero and one, we assume λ_0 is drawn from a truncated normal distribution with a mean of 0.5 and standard deviation of one. The initial condition $\theta_{1,0}$ rests on a similar assumption when $k = 1$ because $\theta_{1,t} \in (-1, 1)$. In this case, $\theta_{1,0} \sim \text{trunc}\mathcal{N}(0, 1)$.

The truncated normal-random walks $\vartheta_{1,t}$ and λ_t along with $\ln \varsigma_{\eta,t}^2$ and $\ln \varsigma_{\upsilon,t}^2$ are the TVPs that enter our particle filter algorithm when $k = 1$. In this case, the state space system (15.1) and (15.2) is linear conditional on $\mathcal{Y}_t$ and $\mathcal{V}_t = [\ln \varsigma_{\eta,t}^2 \;\; \ln \varsigma_{\upsilon,t}^2 \;\; \vartheta_{1,t} \;\; \lambda_t]'$. We simulate this state vector using

$$\mathcal{V}_{t+1} \quad = \quad \mathcal{V}_t \; + \; \mathbf{\Omega}_{\mathcal{E}}^{0.5} \mathcal{E}_{t+1}, \tag{16}$$

where the vector $\mathcal{D}_{\mathcal{E}} = [\sigma_\eta^2 \;\; \sigma_\upsilon^2 \;\; \sigma_{\phi,1}^2 \;\; \sigma_\kappa^2]$ contains the non-zero elements of the diagonal matrix $\mathbf{\Omega}_{\mathcal{E}}$,

[12] Chen and Liu (2000) refer to this particle filter algorithm as a mixture of Kalman filters.

[13] When ε_t is a fixed coefficient-AR(k), $\theta_1, \ldots, \theta_k$ are added to Ψ and $\sigma_{\phi,1} \ldots, \sigma_{\phi,k}$ are deleted.

[14] In a future draft, we plan to estimate $\mathcal{S}_t$ and Ψ jointly by embedding a particle filtering inside a Metropolis-Hastings Markov chain Monte Carlo simulator using methods developed by Andrieu, Doucet, and Holenstein (2010), Flury and Shephard (2011), Herbst and Schorfheide (2014) and Schorfheide, Song, and Yaron (2014).

the elements of $\mathcal{E}_{t+1} = [\xi_{\eta,t+1}\ \xi_{\upsilon,t+1}\ \phi_{1,t+1}\ \kappa_{t+1}]'$ are innovations distributed *IID* standard normal, and $\theta_{1,t+1}$ and λ_{t+1} are truncated to remain in on the open intervals $(-1,\ 1)$ and $(0,\ 1)$, respectively. This system describes the evolution of $\mathcal{V}_t$ and shows that given knowledge of it the state space system (15.1) and (15.2) can be evaluated using standard Kalman filter methods.

Guaranteeing the TVP-AR(2) for ε_t is stationary at every date t presents another estimation problem. Morley, Nelson, and Zivot (2003) provide an example in which a pair of AR(2) coefficients have roots outside the unit circle; also see Morley (1999) and Xu (2013). We adapt their approach to restrict $\theta_{1,t}$ and $\theta_{2,t}$ to guarantee ε_t is stationary. The first step revises equation (2.5) of the SW-UC model, which are the random walks of $\theta_{j,t}$, $j = 1, \ldots, k$, by imposing the restrictions

$$\theta_{1,t} = \frac{2\vartheta_{1,t}}{1 + \left|\vartheta_{1,t}\right|}, \tag{17.1}$$

$$\theta_{2,t} = \left(1 - \left|\vartheta_{1,t}\right|\right)\vartheta_{2,t} + \left|\vartheta_{1,t}\right|, \tag{17.2}$$

where $\vartheta_{1,t}$ is unrestricted, but $\vartheta_{2,t} \in (-1,\ 1)$. The "raw" TVP-AR coefficients $\vartheta_{1,t}$ and $\vartheta_{2,t}$ evolve as independent random walks

$$\vartheta_{j,t} = \vartheta_{j,t-1} + \sigma_{\phi,j}\phi_{j,t}, \tag{17.3}$$

where $j = 1, 2$, and $\phi_{j,t}$ remains a standard normal random variable. The initial conditions for $\vartheta_{1,t}$ and $\vartheta_{2,t}$ are $\vartheta_{1,0} \sim \mathcal{N}(0,\ 100)$ and $\vartheta_{2,0}$ is sampled from a truncated standard normal distribution because $\vartheta_{2,t} \in (-1,\ 1)$. The particle filter algorithm described below draws $\vartheta_{1,0}$ and $\vartheta_{2,0}$ to initialize the two random walks of (17.3) and samples from $\phi_{1,t} \sim \mathcal{N}(0, 1)$ and $\phi_{2,t} \sim \mathcal{N}(0, 1)$ to generate synthetic samples of $\vartheta_{1,t}$ and $\vartheta_{2,t}$. These artificial samples are transformed into samples of $\theta_{1,t}$ and $\theta_{2,t}$ using equations (17.1) and (17.2) to ensure the TVP-AR(2) of ε_t has roots outside the unit circle.

The "raw" TVP-AR coefficient $\vartheta_{1,t}$ replaces the TVP-AR1 coefficient $\theta_{1,t}$ in our particle filter algorithm and $\vartheta_{2,t}$ is added to it when $k = 2$. Thus, the state space system (15.1) and (15.2) is linear conditional on $\mathcal{Y}_t$ and the multivariate random walk (16), where $\vartheta_{1,t}$ and $\vartheta_{2,t}$, which evolve according to the random walk (17.3), become part of $\mathcal{V}_t = [\ln\varsigma^2_{\eta,t}\ \ \ln\varsigma^2_{\upsilon,t}\ \ \vartheta_{1,t}\ \ \vartheta_{2,t}\ \ \lambda_t]'$, and equations (17.1) and (17.2) produce $\theta_{1,t}$ and $\theta_{2,t}$.

3.2 The Particle and Kalman Filters

The Rao-Blackwellization process yields a state space model (15.1) and (15.2) that is linear in $\mathcal{S}_t$ given a history of $\mathcal{V}_{1:i}$, $i = 1, \ldots, t$, $t \leq T$. Thus, Kalman filtering techniques are used to generate samples

of the unobserved states, $\mathcal{S}_t$. Nonetheless, we resort to the particle filter to approximate the density of the latent state $\mathcal{V}_t$ using a sequential importance sampler with resampling (SISR). A SIS builds the density of $\mathcal{S}_{1:T}$ by iterating through the sample from $t=1$ given initial conditions, to $t=2$ conditional on $\mathcal{S}_{1:1}$, $\mathcal{Y}_{1:1}$, $\mathcal{V}_{1:1}$, and Ψ, and ending at $t=T$ conditional on $\mathcal{S}_{1:T-1}$, $\mathcal{Y}_{1:T-1}$, $\mathcal{V}_{1:T-1}$, and Ψ. The resampling step shuffles the M particles of $\left\{\mathcal{S}_{1:T}^{(i)}\right\}_{i=1}^{M}$ using weights that are estimates of a particle's share of the state space model's likelihood, which guarantees no single particle will receive all the probability mass in constructing the density of $\left\{\mathcal{S}_{1:T}^{(i)}\right\}_{i=1}^{M}$ as M becomes large. The weights are calculated by running each particle through the Kalman filter to obtain an estimate of the likelihood.

We adapt algorithm 3 of Creal (2012, section 2.5.7) and an algorithm of Herbst and Schorfheide (2014, section 7.4.3) to estimate the state space model (15.1) and (15.2) conditional on $\mathcal{Y}_t$ and the multivariate random walk (16). This particle filter consists of the following steps.

1. The filter is initialized by drawing M particles $\mathcal{V}_0^{(i)}$, $i = 1, \ldots, M$, from the prior distributions (16), where conditional on $\mathcal{V}_0^{(i)}$, $\mathcal{S}_0$ has a normal prior $\mathcal{S}_0 \sim \mathcal{N}(\mathcal{S}_{0|0}^{(i)}, \boldsymbol{\Sigma}_{0|0}^{(i)})$.

2. Repeat the following steps for $t = 1, \ldots, T$, where each step uses the particles $\mathcal{V}_{t-1}^{(i)}$, $\mathcal{S}_{t-1|t-1}^{(i)}$, and $\boldsymbol{\Sigma}_{t-1|t-1}^{(i)}$, $i = 1, \ldots, M$, which are the outcomes of applying the Kalman filter in the previous step.

   (a) For $i = 1, 2, \ldots, M$, draw new particles $\mathcal{V}_t^{(i)}$ conditional on $\mathcal{V}_{t-1}^{(i)}$, given the prior (16).

   (b) At date t, engage the Kalman filter to compute

$$
\begin{aligned}
\boldsymbol{\Sigma}_{t|t-1}^{(i)} &= \boldsymbol{\mathcal{A}}_t^{(i)} \boldsymbol{\Sigma}_{t-1|t-1}^{(i)} \left(\boldsymbol{\mathcal{A}}_t^{(i)}\right)' + \boldsymbol{\mathcal{B}}_t^{(i)} \left(\boldsymbol{\mathcal{B}}_t^{(i)}\right)', \\
\boldsymbol{\Omega}_{t|t-1}^{(i)} &= \mathcal{C}_t^{(i)} \boldsymbol{\Sigma}_{t|t-1}^{(i)} \left(\mathcal{C}_t^{(i)}\right)' + \mathcal{D} \boldsymbol{\Omega}_{\mathcal{U}}^{(i)} \mathcal{D}', \\
\tilde{\mathcal{Y}}_t^{(i)} &= \mathcal{Y}_t - \boldsymbol{\mathcal{A}}_t^{(i)} \mathcal{S}_{t|t-1}^{(i)}, \\
\mathcal{L}_t^{(i)} &= -\frac{1}{2}\left[\ln\left|\boldsymbol{\Omega}_{t|t-1}^{(i)}\right| + \left(\tilde{\mathcal{Y}}_t^{(i)}\right)' \left(\boldsymbol{\Omega}_{t|t-1}^{(i)}\right)^{-1} \tilde{\mathcal{Y}}_t^{(i)}\right], \\
\mathcal{K}_t^{(i)} &= \boldsymbol{\Sigma}_{t|t-1}^{(i)} \left(\mathcal{C}_t^{(i)}\right)' \left(\boldsymbol{\Omega}_{t|t-1}^{(i)}\right)^{-1}, \\
\mathcal{S}_{t|t}^{(i)} &= \boldsymbol{\mathcal{A}}_t^{(i)} \mathcal{S}_{t|t-1}^{(i)} + \mathcal{K}_t^{(i)} \tilde{\mathcal{Y}}_t^{(i)}, \\
\boldsymbol{\Sigma}_{t|t}^{(i)} &= \boldsymbol{\Sigma}_{t|t-1}^{(i)} - \boldsymbol{\Sigma}_{t|t-1}^{(i)} \left(\mathcal{C}_t^{(i)}\right)' \left(\boldsymbol{\Omega}_{t|t-1}^{(i)}\right)^{-1} \mathcal{C}_t^{(i)} \boldsymbol{\Sigma}_{t|t-1}^{(i)},
\end{aligned}
$$

   across the M particles, $(i) = 1, 2, \ldots, M$.[15]

   (c) Store the conditional moments $\mathcal{S}_{t|t}^{(i)}$ and $\boldsymbol{\Sigma}_{t|t}^{(i)}$ and particle draws $\mathcal{V}_t^{(i)}$ to report estimates of the joint DGP of the SW-UC model and the average SPF respondent's SI prediction mechanism.

[15] There are a missing observations in the SPF inflation data that the Kalman filter handles using standard methods.

(d) Compute particle weights $\omega_t^{(i)} = \dfrac{\exp\left\{\ell_t^{(i)}\right\}}{\sum_i^M \exp\left\{\ell_t^{(i)}\right\}}$.

(e) For $t < T$, prepare the next iteration by resampling the particles $\mathcal{V}_t^{(i)}$, $\mathcal{S}_{t|t}^{(i)}$, and $\boldsymbol{\Sigma}_{t|t}^{(i)}$, which corresponds to drawing from a multinominal distribution for i using the *pdf* of $\omega_t^{(i)}$.[16]

3. The filtered distribution of $\mathcal{V}_t$ conditional on $\mathcal{Y}_{1:t}$ and Ψ is approximated by the discrete distribution of particles $\mathcal{V}_t^{(i)}$ using the *pdf* of $\omega_t^{(i)}$ and the associated filtered distribution of $\mathcal{S}_t$ is approximated by a mixture of normals $\mathcal{N}\left(\mathcal{S}_{t|t}^{(i)}, \boldsymbol{\Sigma}_{t|t}^{(i)}\right)$ with the weights $\omega_t^{(i)}$. Thus, the filtered means of $\mathcal{S}_t$ and $\mathcal{V}_t$ are approximated by $\mathcal{S}_{t|t} = \sum_{i=1}^M \omega_t^{(i)} \mathcal{S}_{t|t}^{(i)}$ and $\mathcal{V}_{t|t} = \sum_{i=1}^M \omega_t^{(i)} \mathcal{V}_t^{(i)}$.

4. Since the M particles have been reweighted at every step, the date t data density is estimated by calculating the average of $\ell_t^{(i)}$ over the M particles

$$\mathcal{P}\left(\mathcal{Y}_t \,\middle|\, \mathcal{Y}_{1:t-1}; \Psi\right) \propto \frac{1}{M} \sum_{i=1}^{M} \exp\left\{\ell_t^{(i)}\right\}, \quad t = 1, \ldots, T. \tag{18}$$

The data density $\mathcal{P}(\mathcal{Y}_t \,|\, \mathcal{Y}_{1:t-1}; \Psi)$ is a vehicle for computing the log likelihood of the state space system (15.1) and (15.2). The density of (18) is summed across the $t = 1, \ldots, T$ observations to produce

$$\mathcal{L}\left(\Psi \,\middle|\, \mathcal{Y}_{1:T}\right) = \sum_{t=1}^{T} \log\left(\mathcal{P}\left(\mathcal{Y}_t \,\middle|\, \mathcal{Y}_{1:t-1}; \Psi\right)\right). \tag{19}$$

Below we report the log likelihood (19) of versions of the conditionally linear state space system (15.1) and (15.2) that differ by restrictions on the lag length of the TVP-AR(k) of ε_t and whether the AR coefficients are fixed or time-varying. Thus, we evaluate competing versions of the state space model (15.1) and (15.2) using the log likelihood (19) when Ψ is calibrated.[17]

3.3 The Mixture Smoother

We compute smoothed draws of $\mathcal{V}_t$ conditional on the entire sample of observations $\mathcal{Y}_{1:T}$ and the calibrated vector Ψ using the algorithm of Godsill, Doucet, and West (2004). Given a distribution of M

[16] Creal (2012) advises reweighting the particles $\mathcal{V}_t^{(i)}$, $\mathcal{S}_{t|t}^{(i)}$, and $\boldsymbol{\Sigma}_{t|t}^{(i)}$ only to prepare the $t+1$ step of the particle filter. However, the unweighted particles are retained to report results and to compute smoothed estimates.

[17] When Ψ is estimated, a result in Andrieu, Doucet, and Holenstein (2010) is relevant. They show the distribution of a Markov chain Monte Carlo (MCMC) simulator is independent of the error generated by a particle in a SMC algorithm. An implication is the density (18) provides an unbiased estimate of the likelihood at date t. Thus, a Metropolis-Hasting MCMC simulator wrapped around a particle filter produces estimates of the elements Ψ across different models that can be evaluated using unbiased estimates of the log likelihood (19).

particles $\mathcal{V}_t^{(i)}$ with weights $\omega_t^{(i)}$, $i = 1, \ldots, M$, generate J sequences of smoothed draws, denoted $\widetilde{\mathcal{V}}_t^{(j)}$ for $j = 1, \ldots, J$ by iterating the next algorithm. For each j, draw $\widetilde{\mathcal{V}}_T^{(j)}$ from the filtered particle draws $\mathcal{V}_T^{(i)}$ weighted by $\omega_T^{(i)}$, and then iterate backwards from $t = T-1$ to $t = 1$ in the following steps

1. For each particle $i = 1, \ldots, M$, calculate $\widetilde{\omega}_t^{(i)} = w_t^{(i)} f\left(\widetilde{\mathcal{V}}_{t+1}^{(j)} \middle| \mathcal{V}_t^{(i)}\right)$, where $f\left(\widetilde{\mathcal{V}}_{t+1}^{(j)} \middle| \mathcal{V}_t^{(i)}\right)$ is computed from the system of transition equations (16).[18]

2. Draw $\widetilde{\mathcal{V}}_t^{(j)}$ from the set of filtered particles $\mathcal{V}_t^{(i)}$ with weight $\widetilde{\omega}_t^{(i)}$.

3. Approximate the posterior distribution of $\mathcal{V}_t$ conditional on the entire sample of observations $\mathcal{Y}_{1:T}$ and the vector of parameters Ψ using the J sequences of $\widetilde{\mathcal{V}}_{1:T}^{(j)}$.

4. The smoothed distribution of $\mathcal{S}_t$ is calculated given the previous approximation

$$\mathcal{P}\left(\mathcal{S}_{1:T} \middle| \mathcal{Y}_{1:T}; \Psi\right) = \int_{\mathcal{V}_{1:T}} \mathcal{P}\left(\mathcal{S}_{1:T} \middle| \mathcal{V}_{1:T}, \mathcal{Y}_{1:T}; \Psi\right) dF(\mathcal{V}_{1:T}) \approx \sum_{j=1}^{J} \mathcal{P}\left(\mathcal{S}_{1:T} \middle| \widetilde{\mathcal{V}}_{1:T}^{(j)}, \mathcal{Y}_{1:T}; \Psi\right). \quad (20)$$

The density $\mathcal{P}\left(\mathcal{S}_{1:T} \middle| \widetilde{\mathcal{V}}_{1:T}^{(j)}, \mathcal{Y}_{1:T}; \Psi\right)$ is the multivariate normal distribution produced by the Kalman smoother for the conditionally linear state space (15.1) and (15.2) because the approximation on the right of (20) conditions on a single trajectory of $\mathcal{V}_{1:T}$. For each sequence of $\mathcal{V}_t$ generated by the particle smoother, we generate draws for $\mathcal{S}_{1:T}$ from the Kalman smoother's distribution using the disturbance smoothing algorithm of Durbin and Koopman (2002).[19]

4 The Data and Estimates

Our motivation for opening up the SW-UC model SI law of motion along these dimensions rests on the existing literature. As surveyed by Creal (2012) and Shephard (2013), there is persuasive evidence that the SVs of trend inflation and the inflation gap have varied substantially in recent US data; also see Stock and Watson (2010), Grassi and Proietti (2010) and Mertens (2015). Further, Nason and Smith (2014) report that in the last 45 years of U.S. data stickiness in SPF inflation predictions, as described by λ, has changed. This section contributes new evidence about the comovement of $\sigma_{\bar{\epsilon},t}$, $\sigma_{\tilde{\epsilon},t}$, and λ_t over the business cycle along with evidence of time variation in inflation gap persistence.

[18] For λ_t, $\theta_{1,t}$ ($k = 1$) and $\vartheta_{2,t}$ ($k = 2$), the conditional *pdf*s are truncated normals.

[19] An alternative smoothing algorithm generating jointly draws $\widetilde{\mathcal{V}}_t^{(j)}$ and conditional moments $\mathcal{S}_{t|T}^{(j)}$ and $\boldsymbol{\Sigma}_{1,t|T}^{(j)}$ is sketched by Prado and West (2010).

4.1 The Data

Our estimates are conditional on a sample of real time realized inflation, π_t, and average SPF inflation prediction, $\pi^{SPF}_{t,t+h}$. We obtain this sample from the Real-Time Data Set for Macroeconomists (RTDSM), which is compiled by the Federal Reserve Bank (FRB) of Philadelphia.[20] The data consist of observations from $1968Q4$ through $2014Q2$ for π_t as well as $\pi^{SPF}_{t,t+h}$ at horizons $h = 1, \ldots, 5$.

Realized inflation is computed from the RTDSM's quarterly real-time vintages of the GNP and GDP deflator.[21] These vintages reflect data releases that were publicly available around the middle of quarter t and most often the publicly available information contains observations through quarter $t-1$. Using these vintages of real time realized inflation, we compute the quarterly difference in the log levels of real time observations on the implicit GNP or GDP deflator. These quarterly price level data are transformed into inflation measured at an annualized rate using $\pi_t = 400[\ln P_t - \ln P_{t-1}]$, where P_t denotes the date t level of the GNP or GDP deflator.

The average SPF predictions include a nowcast of the GNP or GDP deflator's level and forecasts of these price levels for the next four quarters. These surveys are most often collected at the middle of date t (*i.e.,* the quarter), which suggests $\pi^{SPF}_{t,t+h}$ is reported without complete knowledge of π_t. We comply with this timing protocol by assuming date t SPF inflation predictions are conditional on data available through the end of date $t-1$. Thus, the average SPF nowcast, 1-step, ..., 4-step predictions, which are denoted $\pi^{SPF}_{t,t+1}, \pi^{SPF}_{t,t+2}, \ldots, \pi^{SPF}_{t,t+5}$, are made at the end of date $t-1$. These inflation predictions are computed as described previously using the annualized log differences between the SPF prediction of a deflator's level and the quarterly lagged real time realized price level supplied by the RTDSM.

4.2 The Fit of the Joint DGPs

Table 2 lists log marginal data densities, which are generated using the particle filter algorithm described in section 3.2, especially its step 2(b), and equations (19) and (20). The log marginal data densities are employed to gauge the fit of the competing joint DGPs of the SW-UC model and SI prediction mechanism of the average SPF respondent. The competing DGPs are distinct along two dimensions. Two DGPs are defined by the lag length of the TVP-AR(k) of ε_t, which is either $k = 1$ or $k = 2$. Setting $\theta_1 = \theta_{1,t}$ or $\theta_1 = \theta_{1,t}$ and $\theta_2 = \theta_{2,t}$ creates two more DGPs. The four joint DGPs are labeled in table 2 the

[20] The data are available at `http://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/`.

[21] Before $1992Q1$, the SPF and RTDSM measured the U.S. output price level with the implicit GNP deflator. From $1992Q1$ to $1996Q4$, the implicit GDP deflator pays this role, but this deflator is replaced by the chain weighted GDP deflator from $1997Q1$ to the end of the sample.

SI-λ_t Law of Motion plus SW-UC-SV-AR(1), SW-UC-SV-AR(2), SW-UC-SV-TVP-AR(1), or SW-UC-TVP-AR(2) model that denote the SI prediction mechanism of the average SPF respondent plus the SW-UC model with SV and a fixed coefficient AR(1), fixed coefficient AR(2), TVP-AR(1), or TVP-AR(2), respectively.

This gives us four versions of the joint DGP of the SW-UC model and SI prediction mechanism of the average SPF respondent to estimate. We evaluate the fit of the SI-λ_t-law-of-motion-SW-UC-SV-AR(1), -SW-UC-SV-AR(2), -SW-UC-SV-TVP-AR(1), and -SW-UC-TVP-AR(2) models by comparing the associated log marginal data densities. Table 2 contains these log marginal data densities, which show the models with TVP-AR inflation gaps dominate models with inflation gap persistence tied to fixed coefficient ARs. Across the SI-λ_t-law-of-motion-SW-UC-SV-TVP-AR(1) and SI-λ_t-law-of motion-SW-UC-SV-TVP-AR(2) models, the relevant log marginal data densities indicates the former model dominates the model with the higher-order TVP-AR(2) driving time-varying persistence in ε_t. Thus, the data favor the SI-λ_t-law-of-motion-SW-UC-SV-TVP-AR(1) model. The rest of the paper focuses almost exclusively on estimates of the SI-λ_t-law-of-motion-SW-UC-SV-TVP-AR(1) model.

4.3 Realized Inflation, Trend Inflation, and Gap Inflation

Figures 1 to 5 plot filtered RE trend inflation, $\tau_{t|t}$, filtered SI trend inflation $F_{t|t}\tau_t$, the h-step ahead average SPF inflation prediction, $\pi^{SPF}_{t,t+h}$, and realized real time inflation, π_t on the $1968Q4$ to $2014Q2$ sample. The average SPF inflation nowcast, 1-step ahead inflation prediction, 2-step ahead inflation prediction, 3-step ahead inflation prediction, and 4-step ahead inflation prediction appear in figures 1 to 5, as solid (green) lines with diamonds. However, remember we link $\pi^{SPF}_{t,t+1}$, $\pi^{SPF}_{t,t+2}$, ..., $\pi^{SPF}_{t,t+5}$ to the average SPF inflation nowcast, 1-step ahead inflation prediction, ..., 4-step ahead inflation prediction to respect the timing of when the SPF is collected by the FRB of Philadelphia. Otherwise, figures 1 to 5 contain the same solid (black) plots of $\tau_{t|t}$, dashed (red) plots $F_{t|t}\tau_t$, and dotted (blue) plots π_t. Also, the shaded areas in figures 1 to 5 are NBER dated recessions.

Differences across $\pi^{SPF}_{t,t+1}$, ..., $\pi^{SPF}_{t,t+5}$ distinguish figures 1 to 5. Figures 1 to 5 agree the inflation spike that coincides with the first oil price shock of 1973-1974 is dominated by gap inflation, ε_t, while trend inflation, τ_t, is responsible for the bulk of the peak in inflation during the early 1980s. Under the SI-law-of-motion-SW-UC-SV-TVP-AR(1) model, the average SPF respondent leans more on the "signal" of τ_t and less on the "noise" of ε_t between the first oil price shock and the Volcker disinflation.

As is well known, $\pi^{SPF}_{t,t+h}$ become smoother (*i.e.*, less volatile) moving from h=1 to h=5. This information and π_t is employed by the SI-λ_t law of motion plus SW-UC-SV-TVP-AR(1) model to generate

estimates of $\tau_{t|t}$ and $F_{t|t}\tau_t$. The estimates $\tau_{t|t}$ and $F_{t|t}\tau_t$ are close to identical for the entire sample period except from late 1968 to early 1970 and from early 1995 to late 1999 as seen in figure 1 to figure 5. During the latter episode, $\tau_{t|t}$ is lower than $F_{t|t}\tau_t$ while the opposite is true at the start of the sample period. Although $F_{t|t}\tau_t$ and $\tau_{t|t}$ differ in late 1968 and 1969, by the early 1970s the SI forecaster updates $F_{t|t}\tau_t$ to $\tau_{t|t}$, which suggests λ_t is closer to zero than one at the start of the sample. The second half of the 1990s sees $F_{t|t}\tau_t$ is greater than $\tau_{t|t}$. The persistence of this gap indicates the average SPF participant is slow to update $F_{t|t}\tau_t$ in response to movements in $\tau_{t|t}$, which implies λ_t is greater than a half in the late 1990s. We confirm these observations below when discussing estimates of λ_t, which are plotted in figures 10 and 11.

Figures 1 to 5 also differ by the distance between $\pi^{SPF}_{t,t+1}, \ldots, \pi^{SPF}_{t,t+5}$ and $\tau_{t|t}$ and $F_{t|t}\tau_t$ from early 1995 to late 1999. During this period, $\pi^{SPF}_{t,t+1}$ is nearer to $\tau_{t|t}$ than to $F_{t|t}\tau_t$ as displayed in figure 1, but for the same years figure 5 shows $\pi^{SPF}_{t,t+5}$ and $F_t\tau_{t|t}$ are nearly identical. Across $h = 2$, 3, and 4, $\pi^{SPF}_{t,t+h}$ moves to $F_{t|t}\tau_t$ and away from $\tau_{t|t}$ in the second half of the 1990s as depicted in figures 2, 3, and 4. These shifts in the average SPF respondent's inflation predictions are further confirmation that the frequency of SI updating is low from early 1995 to late 1999.[22]

Filtered, $\varepsilon_{t|t}$, and smoothed, $\varepsilon_{t|T}$, inflation gaps are plotted in figure 6. Figure 6 shows $\varepsilon_{t|t}$ and $\varepsilon_{t|T}$ rise between 7 to 9% from $1973Q4$ to $1974Q4$. At the same time, $\tau_{t|t}$ and $F_{t|t}\tau_t$ are less than 4%. The inflation spike of the early 1980s provides a striking contract because $\tau_{t|t}$ and $F_{t|t}\tau_t$ rise to almost 8%. The upshot is the average SPF participant gives $\varepsilon_{t|t}$ (or $\varepsilon_{t|T}$) less responsibility for generating the inflation spike of the early 1980s. Thus, figure 6 provides evidence the average SPF respondent attributed the inflation spike tied to the first oil price shock to movements in ε_t, but by the early 1980s the beliefs of the average SPF respondent about the DGP of inflation have changed. However, this evidence differs from Cogley and Sbordone (2008) whose estimates of a TVP-new Keynesian Phillips curve model ascribe a large role to gap inflation in generating π_t between 1975 and 1982.

Another interesting feature of figure 6 is $\varepsilon_{t|t}$ and $\varepsilon_{t|T}$ are most often negative from 1983 to 2000. This behavior in gap inflation is consistent with π_t below trend inflation, which is observed in figures 1

[22] During this period, $\tau_{t|t}$, $F_{t|t}\tau_t$, $\pi^{SPF}_{t,t+1}, \ldots,$ and $\pi^{SPF}_{t,t+5}$ move in ways consistent with the Fed engaging in a monetary policy of "opportunistic disinflation" as described by Meyer (1996) and Orphanides and Wilcox (2002). For example, Orphanides and Wilcox quote Vice Chairman Blinder and President Boehne of the FRB-Philadelphia as advocating the 1990s Fed wait for a state of the world in which there is little cost to a monetary policy that lowers $F_{t|t}\tau_t$, $\pi^{SPF}_{t,t+3}$, $\pi^{SPF}_{t,t+4}$, and $\pi^{SPF}_{t,t+5}$ instead of taking actions during periods when the potential for a costly disinflation are large. Although the SW-UC-SV-TVP-AR(1) model recovers the average SPF respondent's beliefs about changes in the inflation regime, this evidence cannot be used to evaluate regime shifts in monetary policy. We need information about monetary policy interventions to conduct this evaluation as studied, for example, by Leeper and Zha (2003).

to figures 5. The intuition is inflation below its trend predicts the average SPF respondent anticipates inflation to rise in the future; see Nelson (2008). However, (ex post) inflation is falling from 1983 to 2000, which suggests that during these years the average SPF respondent is updating less frequently to changes in $\mathbf{E}_t\pi_{t+h}$ driven by negative realizations of ε_t. Thus, shocks to ε_t that push π_t below τ_t in the SW-UC-SV-TVP-AR(1) model accords with a rise in λ_t. This prediction of the SI-λ_t law of motion plus SW-UC-SC-TVP-AR(1) model is supported by figures 10 and 11, which are reviewed below.

Fluctuations in $\varepsilon_{t|t}$ and $\varepsilon_{t|T}$ also display less variation after the double dip recessions of the early 1980s. These measures of gap inflation are less than 3% from $1980Q4$ to $1982Q4$. During the rest of the sample, figure 6 shows continued moderation in gap inflation because $\varepsilon_{t|t}$ and $\varepsilon_{t|T}$ are larger than 2% (in absolute value) only in $1987Q1$, $1991Q1$, $2001Q4$, and $2008Q3$. Except for $1987Q1$, these dates fall within NBER dated recessions. Our results indicate the volatility of ε_t is subdued from 1983 to the early 2000s compared with its volatility during the 1970s, which is similar to estimates presented by Nason (2006), Stock and Watson (2007, 2010), Grassi and Prioietti (2010), Creal (2012), Shephard (2013), Nason and Smith (2014), and Mertens (2015). The next section reports estimates of the SVs of τ_t and ε_t, which are $\varsigma_{\eta,t}$ and $\varsigma_{\upsilon,t}$, respectively.

4.4 Trend and Gap Inflation Volatilities

Figures 7 and 8 contain estimates of the filtered and smooth SVs $\varsigma_{\eta,t|t}$, $\varsigma_{\eta,t|T}$, $\varsigma_{\upsilon,t|t}$ and $\varsigma_{\upsilon,t|T}$. The filtered (smoothed) SVs appear in the top (bottom) panel of figures 7 and 8 as solid blue (red) lines. These panels include the interquantile range confidence (*i.e.*, running from the 25th to 75th quantile) bands, which are the thinner lines in these figures.

Several aspects of the filtered and smoothed SVs of τ_t and ε_t stand out in figures 7 and 8. First, the plots of $\varsigma_{\eta,t|t}$ and $\varsigma_{\upsilon,t|t}$ exhibit multiple peaks. Figure 8 displays the largest spike in $\varsigma_{\upsilon,t|t}$ at $1974Q4$. A smaller peak in $\varsigma_{\upsilon,t|T}$ is also observed after the 1973–1975 recession in the bottom panel of figure 8. The largest peak in $\varsigma_{\eta,t|t}$ occurs in $1982Q1$ as seen in figure 7. This is additional evidence that the first oil price is dominated by shocks to ε_t while around the Volcker disinflation the underlying sources of realized inflation were tied to permanent factors influencing τ_t. These estimates of the SV of τ_t and ε_t differ from estimates reported by Nason (2006), Grassi and Prioietti (2010), Stock and Watson (2010), Creal (2012), Shephard (2013), and Nason and Smith (2014). They obtain estimates of the SV of ε_t that do not begin to decline until the mid-1980s. However, Mertens (2015) reports estimates of the SV of τ_t

and ε_t that are similar to figures 7 and 8.[23] Nonetheless, $\varsigma_{\eta,t|t}$ ($\varsigma_{\upsilon,t|t}$) is always less than $\varsigma_{\upsilon,t|t}$ ($\varsigma_{\upsilon,t|T}$) as displayed in figures 7 and 8, which matches estimates beginning with Stock and Watson (2007).

Figures 7 and 8 also depict long run information about the SVs of τ_t and ε_t. The information is the low frequency movements in $\varsigma_{\eta,t|T}$ and $\varsigma_{\upsilon,t|T}$. For example, $\varsigma_{\eta,t|T}$ exhibits a steady decline that begins in 1983 and lasts to the end of the sample. Clark and Davig (2011) also report a drop in the SV of shocks driving $\pi^{SPF}_{t,t+h}$ in the long run using a vector AR with SV.[24] Similar behavior is displayed by $\varsigma_{\upsilon,t|T}$, but the decline starts in 1975. Also, $\varsigma_{\upsilon,t|T}$ rises from 2002 to 2008 and then falls to $2014Q2$. An increase in the SV of ε_t between the recessions of 2001 and 2007–2009 is consistent with results in Creal (2012), Mertens (2015), and Nason and Smith (2014).

Another revealing feature of figures 7 and 8 is the behavior of the SVs of trend and gap inflation around NBER dated recessions. The filtered SVs, $\varsigma_{\eta,t|t}$ and $\varsigma_{\upsilon,t|t}$, often rise during or after a NBER recessions as depicted by figures 7 and 8. These movements are not observed in the two-sided estimates $\varsigma_{\eta,t|T}$ and $\varsigma_{\upsilon,t|T}$ of figures 7 and 8.

4.5 Time-Varying Inflation Gap Persistence and SI Updating

Estimates of $\theta_{1,t}$ and λ_t are found in figures 9, 10, and 11. Figure 9 (10) presents estimates of the inflation gap's TVP-AR1 (time-varying SI updating frequencies) in filtered, $\theta_{1,t|t}$ ($\lambda_{t|t}$), and smoothed, $\theta_{1,t|T}$ ($\lambda_{t|T}$), formats. The filtered estimates appear in the top panel of figures 9 and 10 as solid (blue) plots surrounded by thinner lines, which are inter-quantile confidence bands. The bottom panel of these figures plot $\theta_{1,t|T}$ and $\lambda_{t|T}$ with a solid (red) lines and its confidence bands with thinner lines. The plots of $\theta_{1,t|t}$, $\lambda_{t|t}$, $\theta_{1,t|T}$ and $\lambda_{t|T}$ are the posterior means of the particle draws moving quarter by quarter through the $1968Q4$ to $2014Q2$ sample.

Figure 9, which plots $\theta_{1,t|t}$ and $\theta_{t|T}$, contains meaningful statistical and economic evidence about time-varying persistence in ε_t. First, filtered persistence in ε_t ranges from 0.25 to 0.8 in the 1970s. Next, $\theta_{1,t|t}$ drops from about 0.75 in 1980 to near 0.25 by 1991 before rising to almost 0.6 before the 2001 recession. Third, $\theta_{1,t|t}$ falls from 0.5 to about -0.1 between 2001 and 2009 and stays there for the rest of the sample. Fourth, confidence bands of $\theta_{1,t|t}$ do not contain zero between the 1969–1970 recession and the recent financial crisis, but do since the 2007–2009 recession. A similar narrative holds for $\theta_{t|T}$, except it shows less quarter to quarter variation and a smaller peak in the mid-1990s.

[23] Only Nason and Smith (2014) and Mertens (2015) estimate models conditional on π_t and $\pi^{SPF}_{t,t+h}$, h=1, 2, ..., 5.
[24] Clark and Davig estimate a VAR with SV on professional inflation forecasts of various horizons, CPI inflation, a real activity variable, and the fed funds rate from $1981Q3$ to $2008Q2$.

Plots of $\lambda_{t|t}$ and $\lambda_{t|T}$ display a low frequency swing from more frequent to less frequent updating around the late 1980s and early 1990s in figure 10. The average SPF inflation respondent is estimated to update frequently to news about $\mathbf{E}_t\pi_{t+h}$ from the $1969Q4$ to the end of the 1981–1982 recession because the bottom panel of figure 10 shows $\lambda_{t|T} \approx 0.3$ during this period. Also, confidence bands run from about 0.15 to around 0.5 during this period, which is evidence $\lambda_{t|T} \neq 0$. From $1983Q1$ to 1995, $\lambda_{t|T}$ rises from close to 0.4 to about 0.7 before beginning to decline to 0.6 by the end of the sample. Although confidence bands are wider for $\lambda_{t|T}$ during the second half of the sample, these never cover zero while ranging from about 0.4 to more than 0.8 by $2014Q2$. The top panel of figure 10 tells a similar story, but in the first several years of the sample $\lambda_{t|t}$ varies between 0.4 and 0.6. This signals $\lambda_{t|t}$ is more volatile early in the sample than from 1975 to the end of the sample.

Our estimates of $\lambda_{t|t}$ and $\lambda_{t|T}$ show the frequency of SI updating by the average SPF respondent has shifted during the last 45 years. The evidence indicates the average SPF respondent never relies exclusively on RE forecasts of inflation when forming $\pi^{SPF}_{t,t+h}$. These results are in line with Coibion and Gorodnichenko (2015). However, $\lambda_{t|t}$ and $\lambda_{t|T}$ are estimated within the SI law of motion plus SW-UC-SV-TVP-AR(1) model rather than inferred from reduced form predictive regressions.

We conclude this section by summarizing several of our results in figure 11. Figure 11 plots the same information found in figures 9 and 10. This information is the $\lambda_{t|T}$s estimated by the SI law of motion plus SW-UC-SV-AR(1) and SI law of motion plus SW-UC-SV-TVP-AR(1) models along with the associated measures of persistence in ε_t. The latter (former) model's estimates of $\lambda_{t|T}$ are displayed as a solid red (dotted black) plot. The absolute values of the posterior means of θ_1 and $\theta_{1,t|T}$ are found in figure 11 as solid green and dotted blue lines, respectively.

Figure 11 explains the preferences the data have for the SI-λ_t law of motion plus SW-UC-SV-TVP-AR(1) model over the SI-λ_t law of motion plus SW-UC-SV-AR(1) model. One reason is the former model lacks time-varying persistence in ε_t. This model yields an estimate of $|\theta_1| \approx 0.1$, which leaves ε_t with no persistence and in effect ties its movements to shocks to its SV. Thus, the SI-λ_t law of motion plus SW-UC-SV-AR(1) model attributes transitory fluctuations in $\pi^{SPF}_{t,t+h}$ to SV in ε_t. However, changes in λ_t shift the transition matrix of the system of state equations (9.1) creating fluctuations in $F_{t+1}x_t$ that in turn generate movements in π^{SPF}_{t+h} through the observation equations (9.2). In contrast, the SI law of motion plus SW-UC-SV-TVP-AR(1) model alters the transition matrix of the state equation system (15.1) by changes in $\theta_{1,t}$ and λ_t. Since figure 11 shows the dotted black line is always greater than the solid red line, the SI-λ_t law-of-motion plus SW-UC-SV-AR(1) model estimates less frequent updating compared with the SI-λ_t law-of-motion plus SW-UC-SV-TVP-AR(1) model. This indicates time-varying persistence

in ε_t lessens the average SPF participant's dependence on changing beliefs about the dynamics of π_t that are embedded in λ_t to generate fluctuations in $\pi_{t,t+h}^{SPF}$ using the observation equations (15.2).

The other reason is the SI law of motion plus SW-UC-SV-TVP-AR(1) model estimates meaningful statistical and economic comovement in $\theta_{1,t}$ and λ_t. Figure 11 plots the smoothed estimates of $|\theta_{1,t}|$ and λ_t produced by this model with the solid red and green lines. There is negative comovement in these estimates for most of the sample, which indicates time-varying SI updating is more frequent when ε_t exhibits more time-varying persistence. An implication is there are persistent movements in ε_t that play a large role in generating fluctuations in π_t. Observing these movements leads the average SPF participant to anticipate transitory shocks are important for movements in $\mathbf{E}_t\pi_t$. These beliefs are consistent with the average SPF participant updating more frequently in response to shocks to the time-varying persistence of gap inflation that produce transitory fluctuations in the conditional mean of inflation. Thus, time-varying inflation gap persistence mitigates the stickiness the average SPF respondent needs to produce inflation predictions in the context of the SW-UC model.

5 Conclusions

This paper contributes to the literature that evaluates the dynamics of realized inflation and SPF inflation predictions. We combine the Stock and Watson (2007) unobserved-components (SW-UC) model with the Mankiw and Reis (2002) sticky information (SI) law of motion. The SW-UC model with stochastic volatility is extended to include time-varying inflation gap persistence while we endow the SI law of motion with a time-varying updating parameter. The joint SW-UC model and SI law of motion yield different predictions of trend inflation. A source of the different trend inflation predictions is the SW-UC model and SI law of motion rely on different time-varying rates to discount the histories which generate these forecasts.

We employ this identifying information to estimate the joint dynamics of real time realized inflation and average SPF inflation predictions using different combinations of the SW-UC model and SI law of motion. A particle filter algorithm generates estimates of these joint models. Comparing the estimated likelihoods of these models indicate the data prefer models with time-varying persistence in gap inflation. This persistence lowers the estimated time-varying stickiness of the average SPF inflation prediction. Less sticky average SPF inflation predictions are updated more often given new information about the conditional mean of inflation, which is provided by the SW-UC model with stochastic volatility and time varying inflation gap persistence. This information consists of shocks to these elements of

the SW-UC model. On a quarterly sample from $1969Q1$ to $2014Q2$, our estimates show less frequent SI updating occurs at the same time stochastic volatility and inflation gap persistence are declining. The drop in the frequency of SI updating is observed as realized inflation begins to fall in the mid 1980s. Thus, time-variation in SI updating occurs less frequently at the same moment in time the average SPF participant believes there is less potential for a permanent or persistent shift in the inflation regime.

Results presented in this paper fit into an emerging literature represented by Krane (2011), Nason and Smith (2014), and Mertens (2015). This literature finds professional forecasters are more sensitive to permanent shocks than to transitory shocks when revising their predictions of inflation and other aggregate time series. We hope this paper stimulates more research into the ways in which professional forecasters and other economic agents process information to form beliefs and predictions about future outcomes and events.

References

Andrieu, C., A. Doucet, R. Holenstein (2010). Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society, Series B* 72, 269–342.

Ang, A., G. Bekaert, M. Wei (2007). Do macro variables, asset markets, or surveys forecast inflation better? *Journal of Monetary Economics* 54, 1163–1212.

Chen, R., J.S. Liu (2000). Mixture Kalman filters. *Journal of the Royal Statistical Society, Series B* 62, 493–508.

Clark, T.E., T. Davig (2011). Decomposing the declining volatility of long-term inflation expectations. *Journal of Economic Dynamics and Control* 35, 981–999.

Cogley, T., G. Primiceri, T.J. Sargent (2010). Inflation-gap persistence in the US. *American Economic Journal: Macroeconomics* 2, 43–69.

Cogley, T., T.J. Sargent (2015). Measuring price-pevel uncertainty and instability in the U.S., 1850–2012. *Review of Economics and Statistics* forthcoming.

Cogley, T., T.J. Sargent (2008). Anticipated utility and rational expectations as approximations of Bayesian decision making. *International Economic Review* 49, 185–221.

Cogley, T., T.J. Sargent (2005). Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics* 8, 262–305.

Cogley, T., A. Sbordone (2008). Trend inflation, indexation, and inflation persistence in the new Keynesian Phillips curve inflation-gap persistence in the US. *American Economic Review* 98, 2101–2126.

Coibion, O., Y. Gorodnichenko (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, forthcoming.

Coibion, O., Y. Gorodnichenko (2012) What can survey forecasts tell us about informational rigidities? *Journal of Political Economy* 120, 116–159.

Creal, D. (2012). A survey of sequential Monte Carlo methods for economics and finance. *Econometric Reviews* 31, 245–296.

Doucet, A., S. Godsill, C. Andrieu. (2000) On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing* 10, 197–208.

Durbin, J., S.J. Koopman 2002. A simple and efficient simulation smoother for state space time series analysis. *Biometrika* 89, 603–615.

Faust, J., J.H. Wright (2013). Forecasting inflation. In Elliot, G., A. Timmermann (eds.), HANDBOOK OF ECONOMIC FORECASTING, VOL. 2, pp. 2–56. New York, NY: Elsevier.

Flury, T., N. Shephard (2011). Bayesian inference based only on simulated likelihood: Particle filter analysis of dynamic economic models. *Econometric Theory*, 27, 933–956.

Godsill, S.J., A. Doucet, M. West 2004. Monte Carlo smoothing for nonlinear time series. *Journal of the American Statistical Association* 99, 156–168.

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

Herbst, E., F. Schorfheide (2014). Bayesian inference for DSGE models. Manuscript. Board of Governors of the Federal Reserve System, Washington, D.C.

Jain, M. (2013). Perceived inflation persistence. Working Paper 2013–43, Bank of Canada.

Kozicki, S., P.A. Tinsley (2012). Effective use of survey information in estimating the evolution of expected inflation. *Journal of Money, Credit and Banking* 44, 145–169.

Krane, S.D. (2011). Professional forecasters' views of permanent and transitory shocks to GDP. *American Economic Journal: Macroeconomics* 3, 184–211

Kreps, D.M. (1998). Anticipated utility and dynamic choice. In FRONTIERS OF RESEARCH IN ECONOMIC THEORY: THE NANCY L. SCHWARTZ MEMORIAL LECTURES, 1983-1997, Jacobs, D. P., E. Kalai, M. I. Kamien (eds.), 242–274. Cambridge, MA: Cambridge University Press.

Leeper, E.M., T. Zha (2003). Modest policy interventions. *Journal of Monetary Economics* 50, 1673–1700.

Mankiw, N.G., R. Reis (2002). Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve. *Quarterly Journal of Economics* 117, 1295–1328.

Mertens, E. (2015). Measuring the level and uncertainty of trend inflation. Manuscript. Board of Governors of the Federal Reserve System, Washington, D.C.

Meyer, L. (1996). Monetary policy objectives and strategy. Remarks given at the National Association of Business Economists 38th Annual Meeting, Boston, Massachusetts (September 8). Available at `http://www.federalreserve.gov/boarddocs/speeches/1996/19960908.htm`.

Morley, J. (1999). A note on constraining AR(2) parameters in estimation. Manuscript. Department of Economics, University of New South Wales, Sydney, Australia.

Morley, J.C., C.R. Nelson, E. Zivot (2003). Why are the Beveridge-Nelson and unobserved-components decompositions of GDP so different? *Review of Economics and Statistics* 85, 235–243.

Muth, J.F. (1960). Optimal properties of exponentially weighted forecasts. *Journal of the American Statistical Association* 55, 299–306.

Nason, J.M. (2006). Instability in US inflation: 1967–2005. Federal Reserve Bank of Atlanta *Quarterly Review* 91 (Q2) 39–59.

Nason, J.M., G.W. Smith (2014). Measuring the slowly evolving trend in US inflation with professional forecasts. CAMA working paper 7/2014, Australian National University.

Nelson, C.R. (2008). The Beveridge-Nelson decomposition in retrospect and prospect. *Journal of Econometrics* 146, 202–206.

Orphanides, A., D. Wilcox (2002). The opportunistic approach to disinflation. *International Finance* 5, 47–71.

Prado, R., M. West (2010). TIME SERIES: MODELLING, COMPUTATION AND INFERENCE. Boca Raton, FL: CRC Press, The Taylor Francis Group.

Schorfheide, F., D. Song, A. Yaron (2014). Identifying long-run risks: A Bayesian mixed-frequency approach. Working Paper 20303, NBER, Cambridge, MA.

Shephard, N. (2013). Martingale unobserved component models. Economics Series Working Papers 644, Department of Economics, University of Oxford.

Sims, C.A. (2003). Implications of rational inattention. *Journal of Monetary Economics* 50, 665–690.

Stock, J.H., M.W. Watson (2007). Why has US inflation become harder to forecast? *Journal of Money, Credit and Banking* 39(S1), 3–33.

Stock, J.H., M.W. Watson (2010). Modeling inflation after the crisis. In MACROECONOMIC CHALLENGES: THE DECADE AHEAD, Kansas City, MO: Federal Reserve Bank of Kansas City.

Xu, X.C. (2013). A note on "A Note on constraining AR(2) parameters in estimation." Manuscript. Department of Economics, University of New South Wales, Sydney, Australia.

Table 1. Calibration of Exogenous Random Walk Process $\mathcal{E}_t$

$$\Psi = \left[\sigma_\eta \quad \sigma_\upsilon \quad \sigma_{\zeta,1} \quad \ldots \quad \sigma_{\zeta,\mathcal{H}} \quad \sigma_{\phi,1} \quad \ldots \quad \sigma_{\phi,k}\right]'$$

Ψ	Source
a) Volatility of Innovations to τ_t's and ε_t's SVs.	Stock and Watson (2007)
$\ln \varsigma^2_{\eta,t+1} = \ln \varsigma^2_{\eta,t} + \sigma_\eta \xi_{\eta,t}$	$\sigma_\eta = 0.20$
$\ln \varsigma^2_{\upsilon,t+1} = \ln \varsigma^2_{\upsilon,t} + \sigma_\upsilon \xi_{\upsilon,t}$	$\sigma_\upsilon = 0.20$
b) Volatility of Innovation to λ_t, TVP of SI Forecast Updating.	calibrated
$\lambda_{t+1} = \lambda_t + \sigma_\kappa \kappa_t$	$\sigma_\kappa = 0.05$
c) Volatility of Innovations to SPF Measurement Errors.	calibrated
$\zeta_{t,t+h} \sim \mathcal{N}\left(0, \sigma^2_{\zeta,h}\right)$, $h = 1, \ldots, \mathcal{H}$, $\mathcal{H} = 5$	$\sigma_{\zeta,h} = 1.00$
d) Volatility of Innovations to TVP-AR(k) slope coefficients of the inflation gap, ε_t.	calibrated
$\theta_{j,t+1} = \theta_{j,t} + \sigma_{\phi,j} \phi_{j,t}$, $j = 1, \ldots, k$	
$k = 1$	$\sigma_{\phi,1} = 0.05$
$k = 2$	$\sigma_{\phi,1} = 0.10$ $\sigma_{\phi,2} = 0.05$
e) Fixed Coefficient AR(k) of the inflation gap, ε_t.	MCMC estimates of the SW-UC Model
$k = 1$	$\sigma_{\phi,1} = 0.01$
$k = 2$	$\sigma_{\phi,1} = -0.16$ $\sigma_{\phi,2} = -0.28$

The calibrated parameters listed in panels a)–c) are applied to all the estimated joint DGPs of the SW-UC model and the SI prediction mechanism of the average SPF respondent. The volatilities of the TVP-AR slope coefficients of ε_t listed in panel d) apply to the SW-UC model (2.1)–(2.5). The innovations $\xi_{\eta,t}$, $\xi_{\upsilon,t}$, κ_t and $\phi_{j,t}$ are random variables that are assumed to have standard normal distributions, where $j = 1$ or $j = 1, 2$. Panel e) applies only to the SW-UC model that has persistence in ε_t produced by a fixed coefficient AR(1) or AR(2).

Table 2 Marginal Log Likelihoods of Estimated Models

SI-λ_t Law of Motion plus	$\mathcal{L}\left(\mathcal{Y}_{1:T} \middle\vert \Psi\right)$
SW-UC-SV-AR(1)	-397.01
SW-UC-SV-AR(2)	-406.66
SW-UC-SV-TVP-AR(1)	-378.67
SW-UC-SV-TVP-AR(2)	-387.71

Figure 1: Filtered RE, $\tau_{t|t}$, and SI, $F_{t|t}\tau_t$, Trend Inflation, $\pi^{SPF}_{t,t+1}$, and Realized Real Time Inflation, π_t, 1968Q4 to 2014Q2

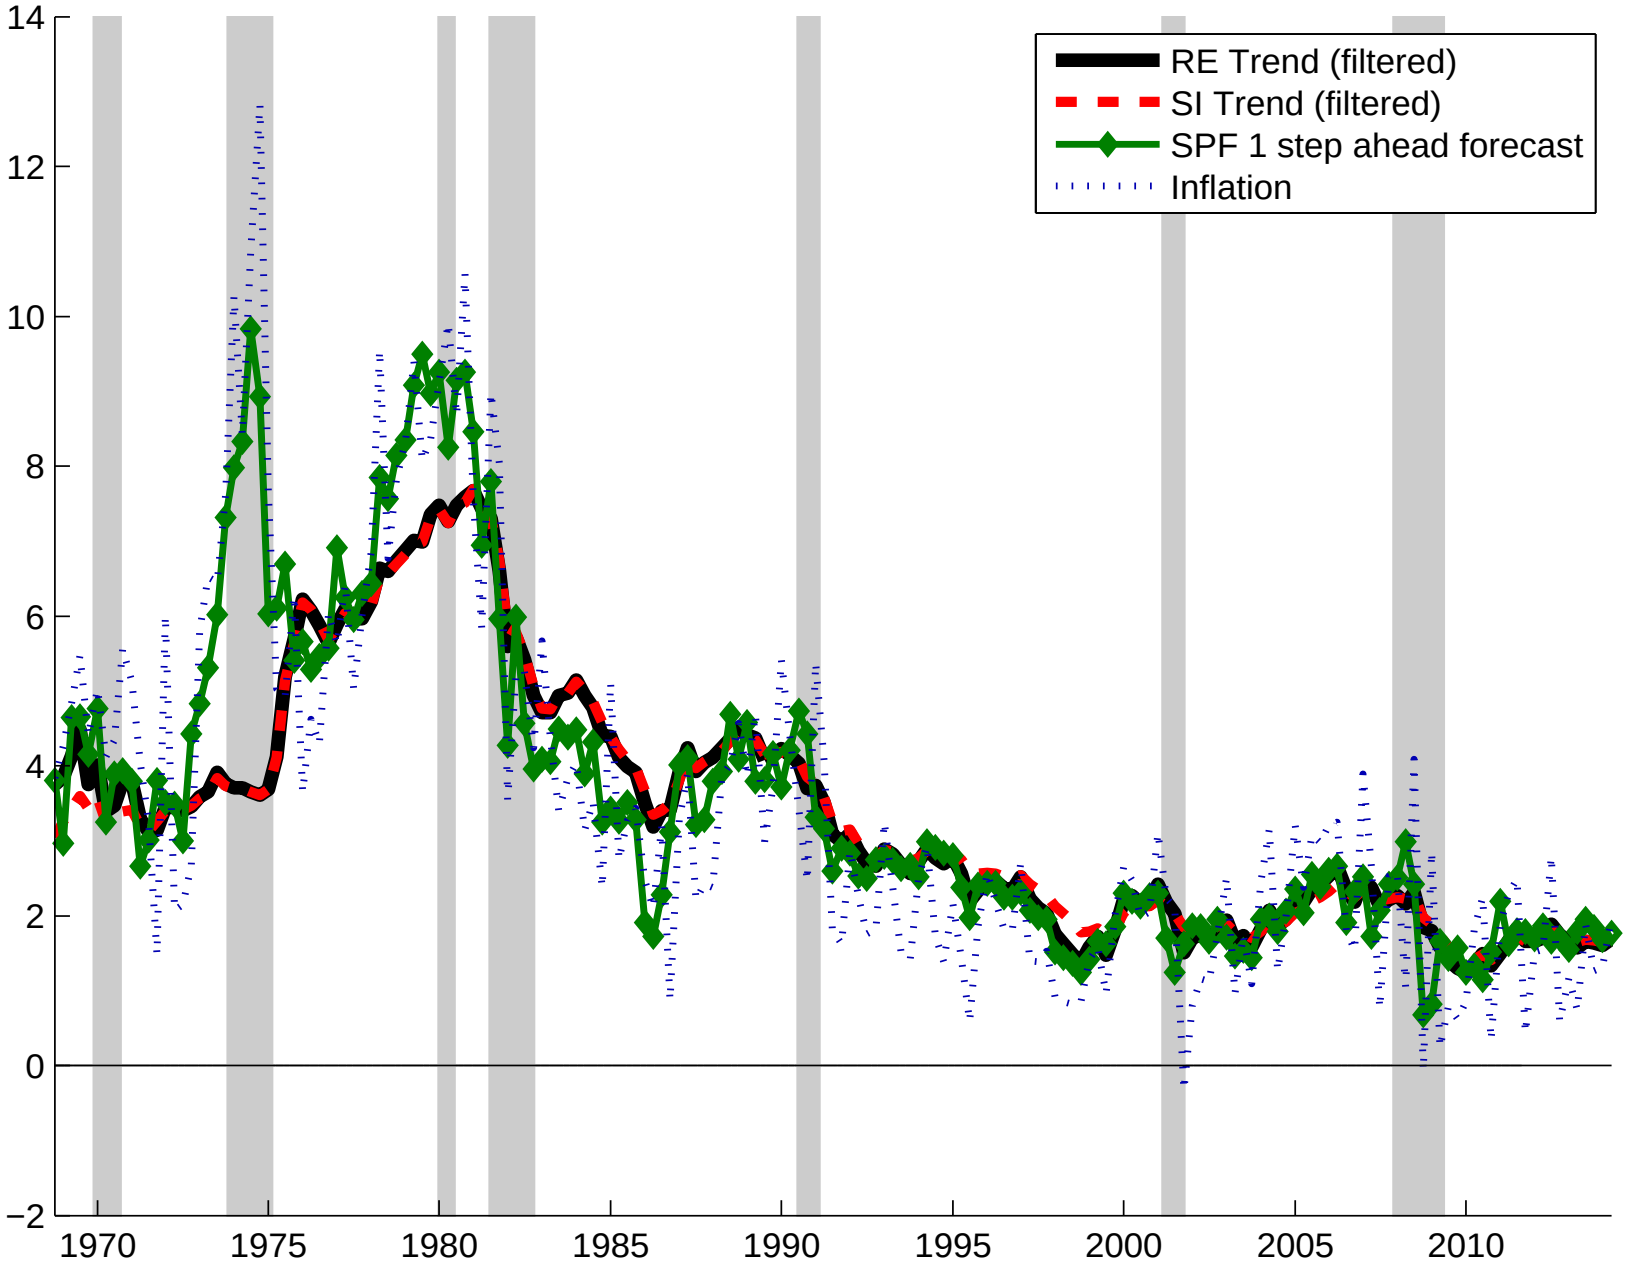

The plots of filtered RE trend inflation, $\tau_{t|t}$, and filtered SI trend inflation, $F_{t|t}\tau_t$, are generated using the particle filter estimates of the TVP-SI law of motion and the SW-UC-SV-TVP-AR(1) model. The average SPF inflation nowcast prediction corresponds to $\pi^{SPF}_{t,t+1}$.

Figure 2: Filtered RE, $\tau_{t|t}$, and SI, $F_{t|t}\tau_t$, Trend Inflation, $\pi^{SPF}_{t,t+2}$, and Realized Real Time Inflation, π_t, $1968Q4$ to $2014Q2$

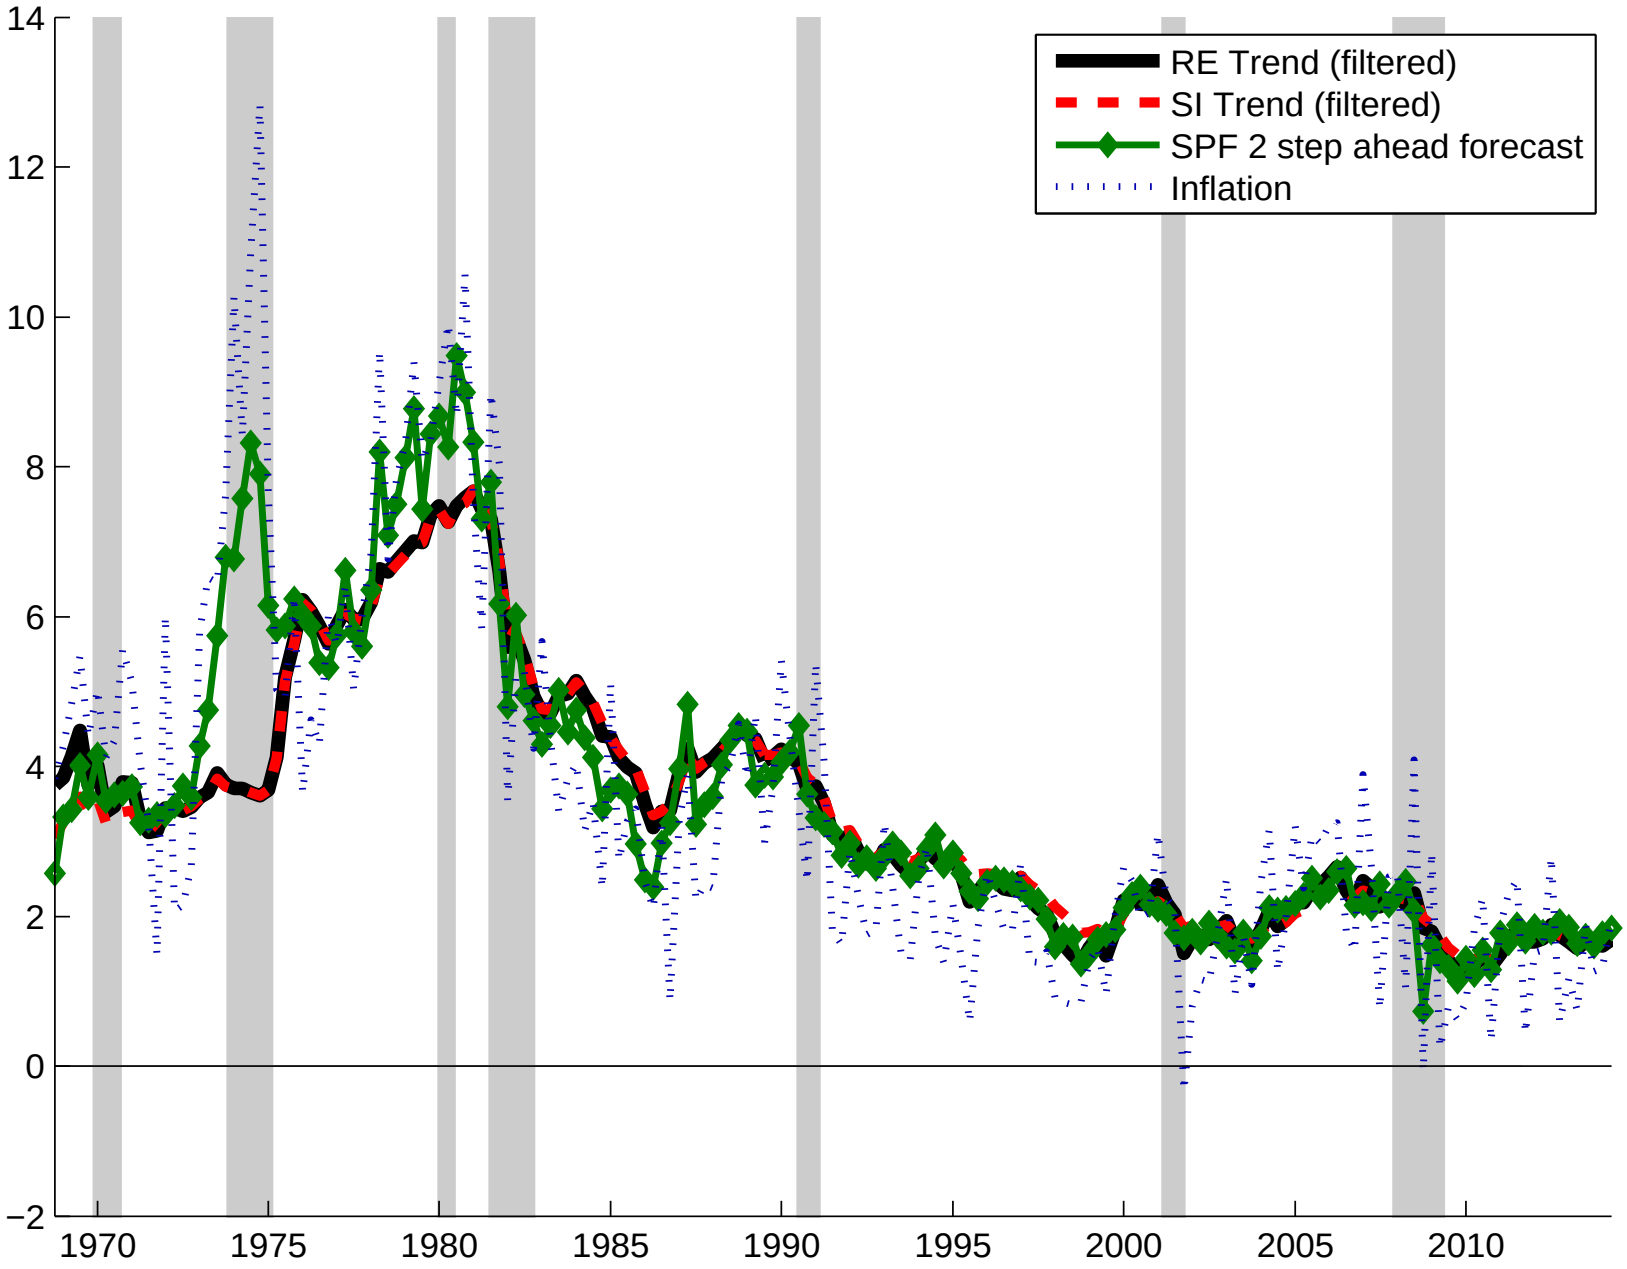

See the notes to figure 1, but the 1-step ahead average SPF inflation prediction corresponds to $\pi^{SPF}_{t,t+2}$.

FIGURE 3: FILTERED RE, $\tau_{t|t}$, AND SI, $F_{t|t}\tau_t$, TREND INFLATION, $\pi^{SPF}_{t,t+3}$, AND REALIZED REAL TIME INFLATION, π_t, 1968$Q4$ TO 2014$Q2$

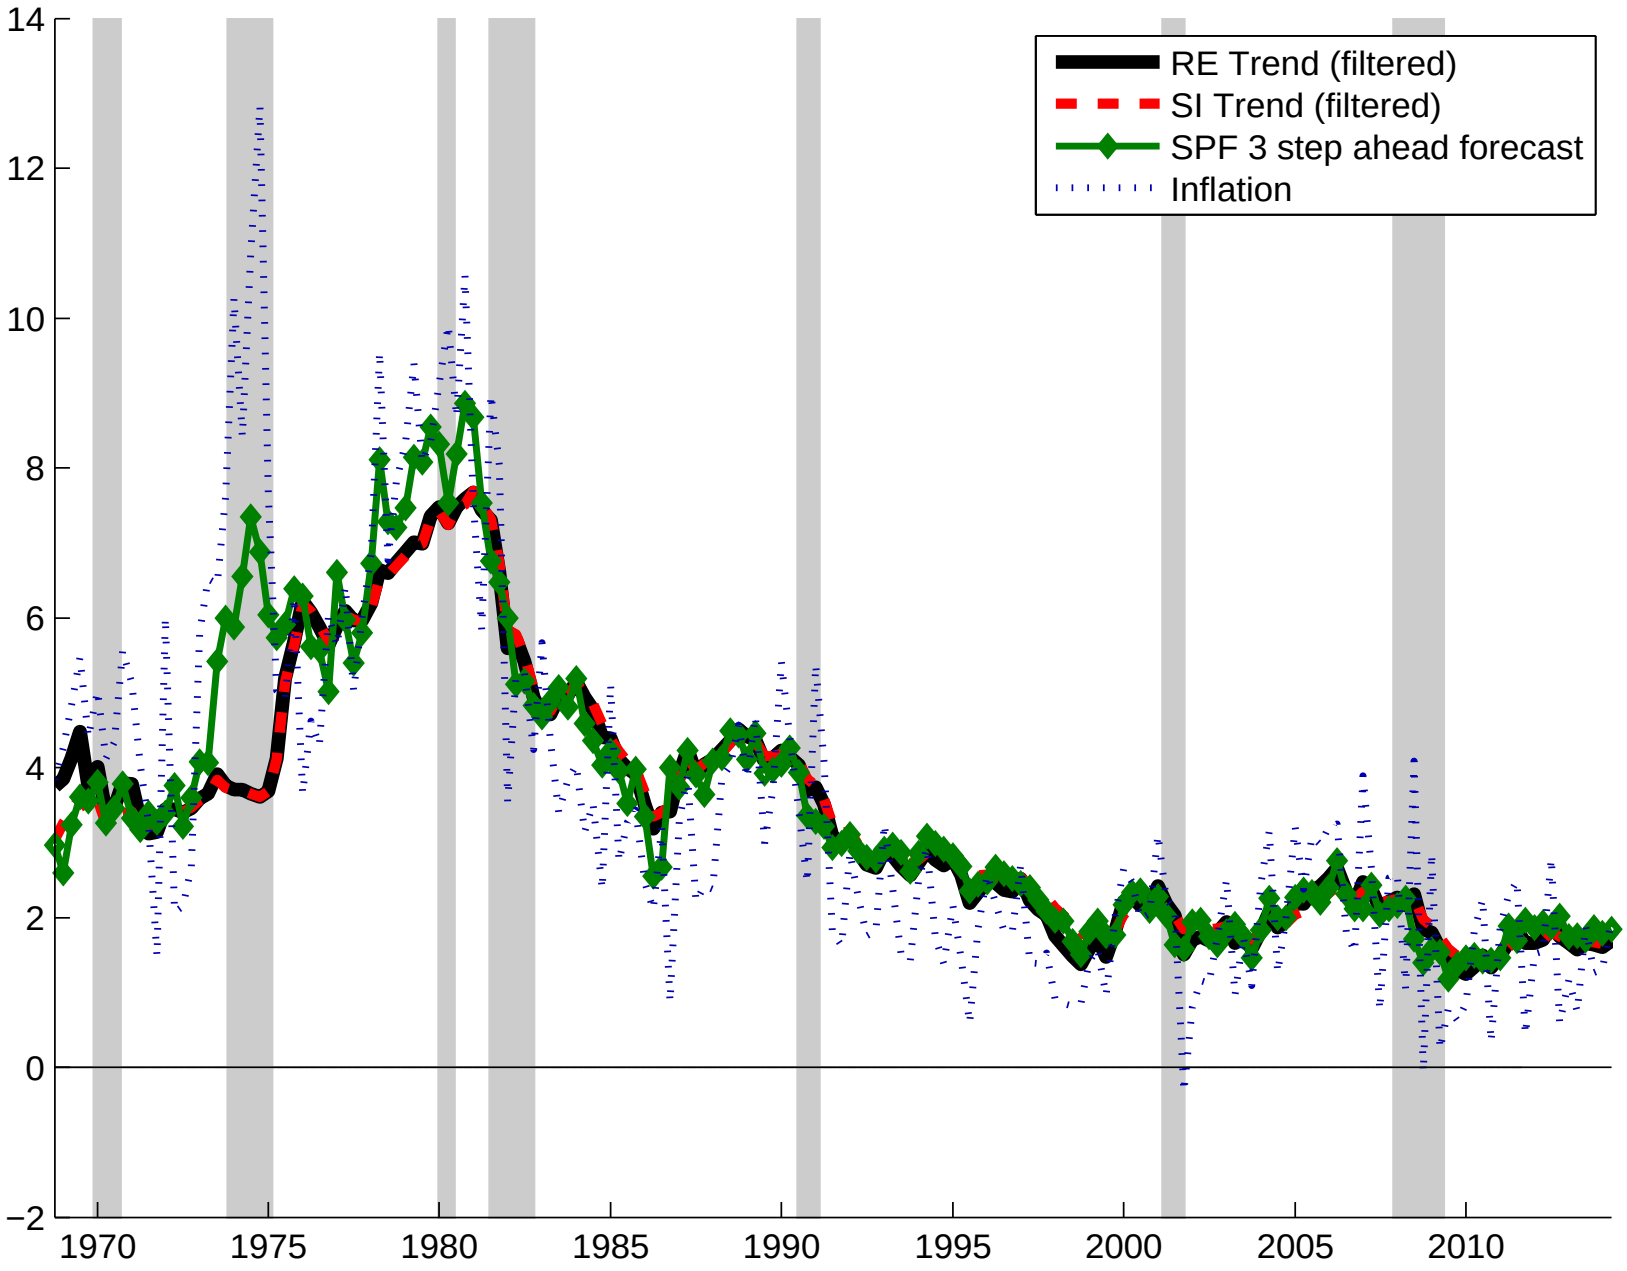

See the notes to figure 1, but the 2-step ahead average SPF inflation prediction corresponds to $\pi^{SPF}_{t,t+3}$.

Figure 4: Filtered RE, $\tau_{t|t}$, and SI, $F_{t|t}\tau_t$, Trend Inflation, $\pi^{SPF}_{t,t+4}$, and Realized Real Time Inflation, π_t, $1968Q4$ to $2014Q2$

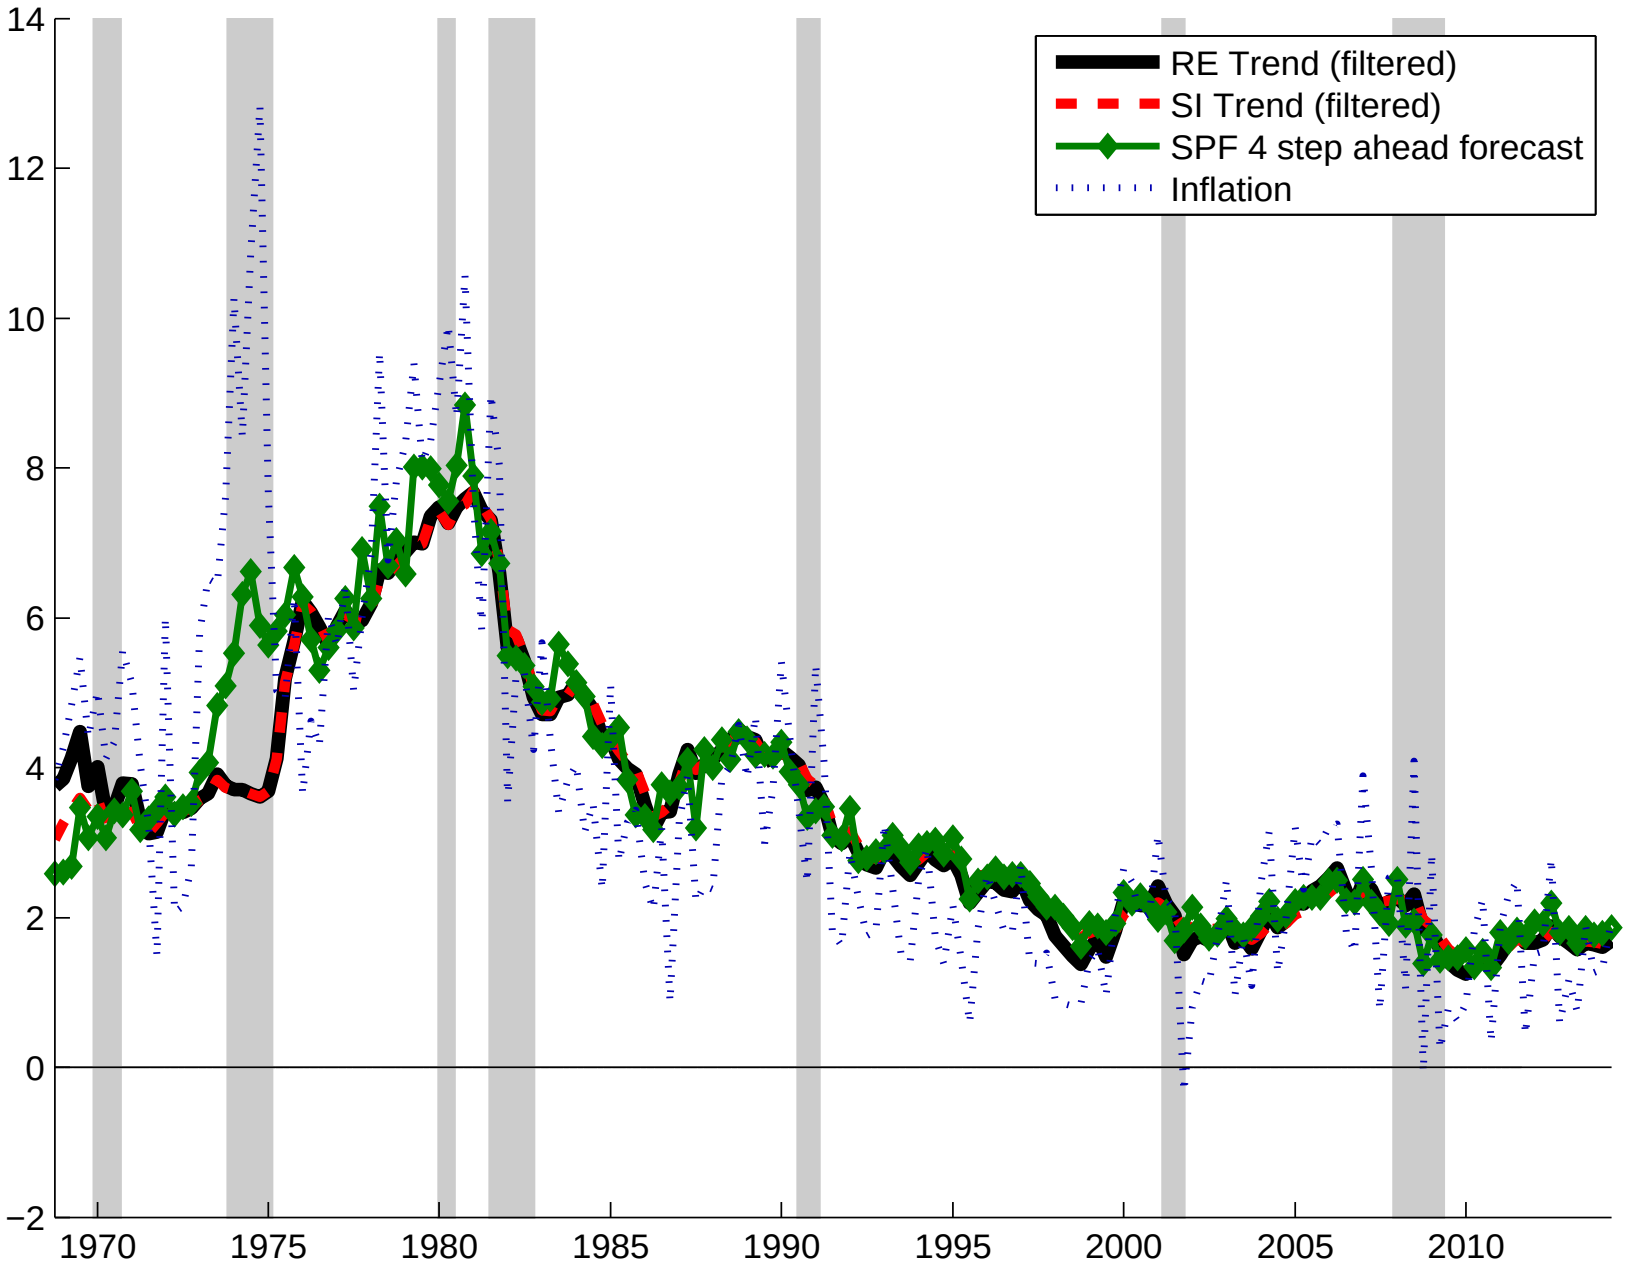

See the notes to figure 1, but the 3-step ahead average SPF inflation prediction corresponds to $\pi^{SPF}_{t,t+4}$.

FIGURE 5: FILTERED RE, $\tau_{t|t}$, AND SI, $F_{t|t}\tau_t$, TREND INFLATION, $\pi^{SPF}_{t,t+5}$, AND REALIZED REAL TIME INFLATION, π_t, 1968$Q4$ TO 2014$Q2$

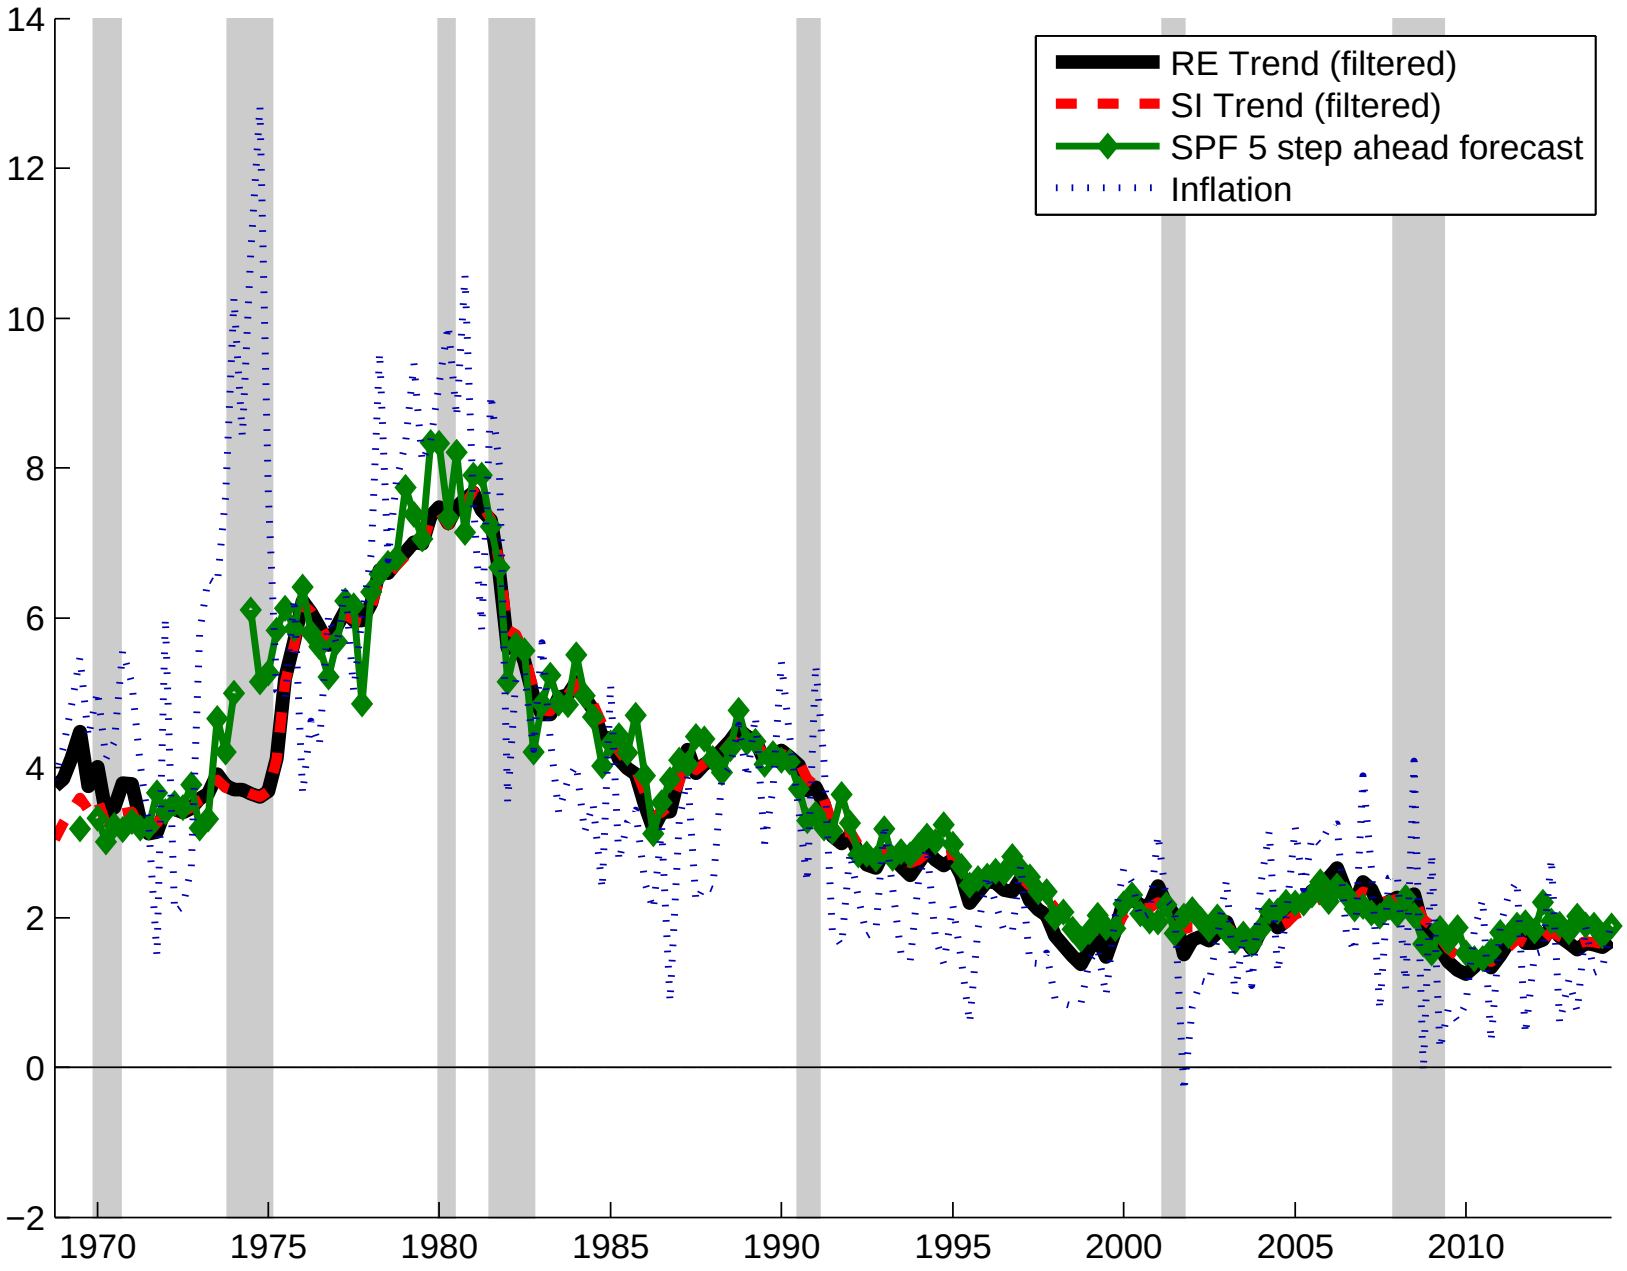

See the notes to figure 1, but the 4-step ahead average SPF inflation prediction corresponds to $\pi^{SPF}_{t,t+5}$.

Figure 6: Filtered, $\varepsilon_{t|t}$, and Smoothed, $\varepsilon_{t|T}$, Gap Inflation 1968Q4 to 2014Q2

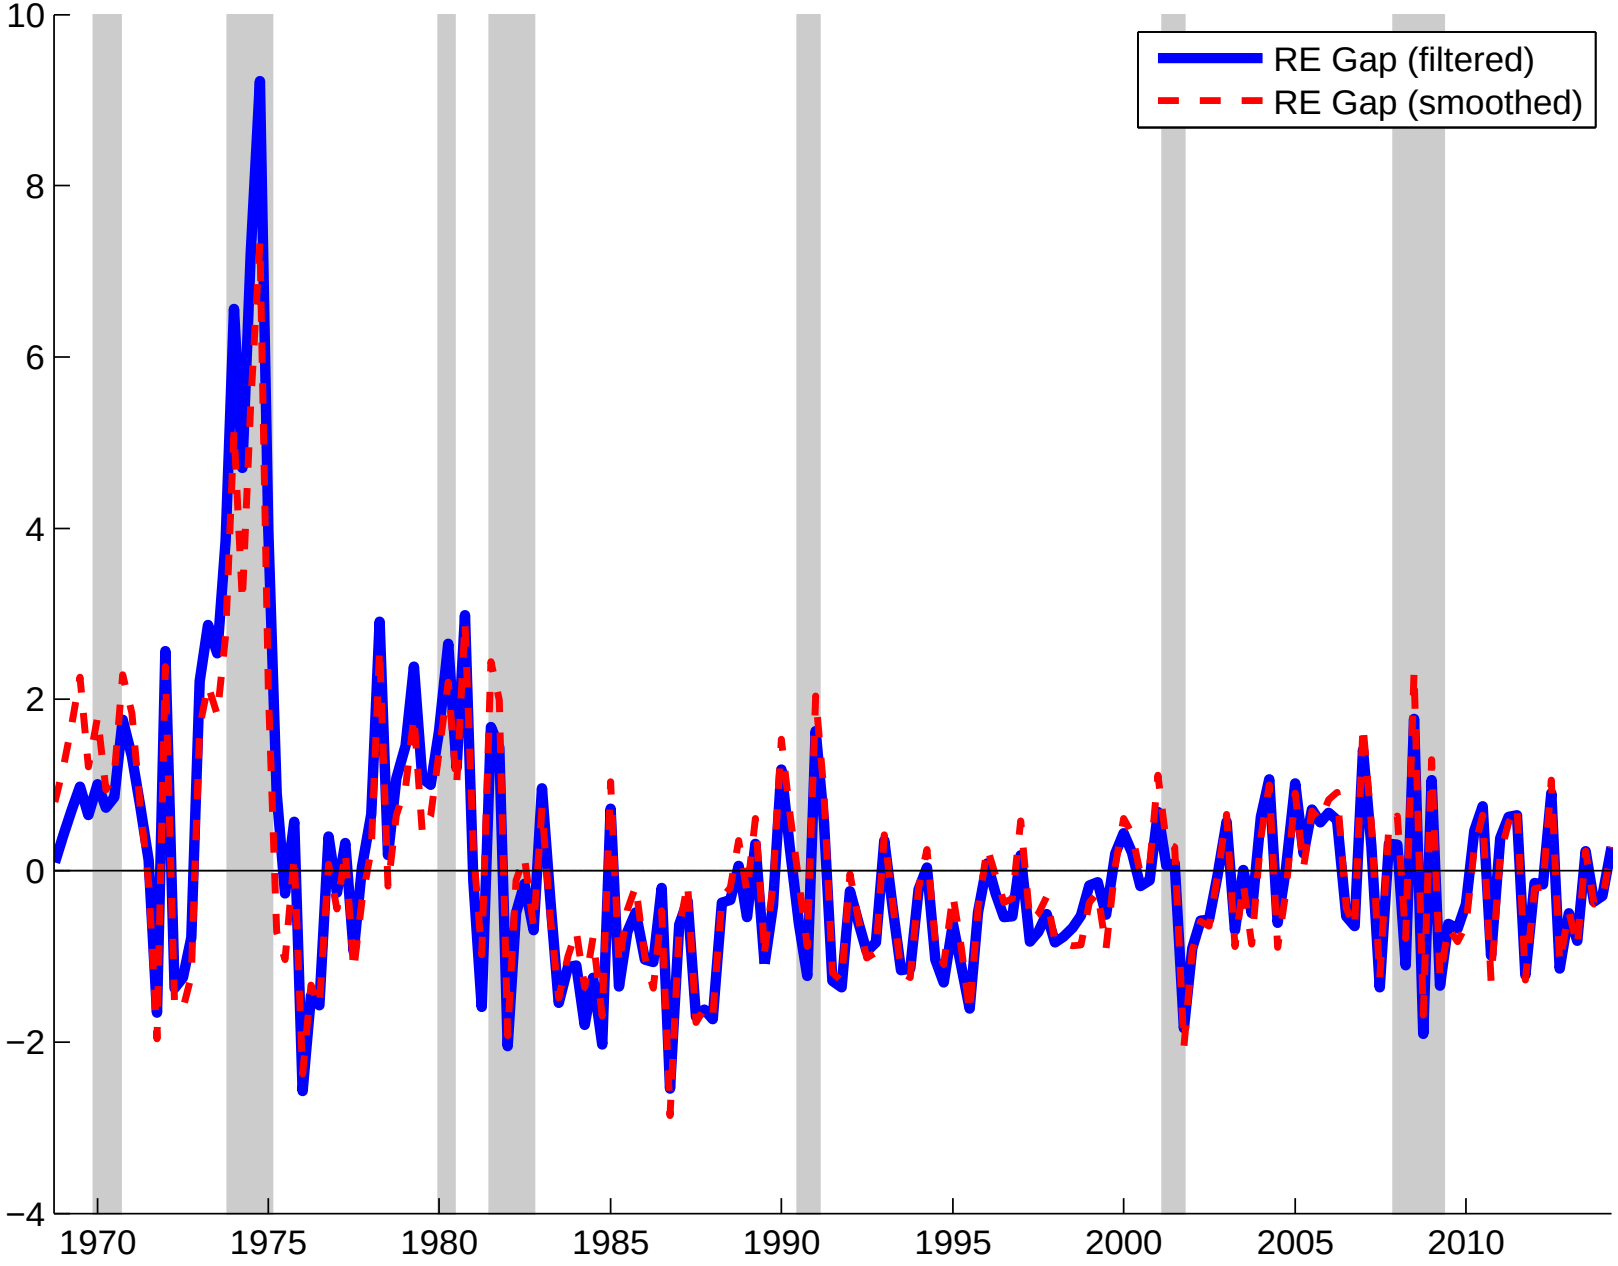

The plots of filtered, $\varepsilon_{t|t}$, and smoothed, $\varepsilon_{t|T}$, inflation gaps are generated using the particle filter estimates of the TVP-SI law of motion and SW-UC-SV-TVP-AR(1) model.

Figure 7: Filtered, $\varsigma_{\eta,t|t}$, and Smoothed, $\varsigma_{\eta,t|T}$, Trend Inflation SV, 1968Q4 to 2014Q2

Filtered = Blue and Smoothed = Red

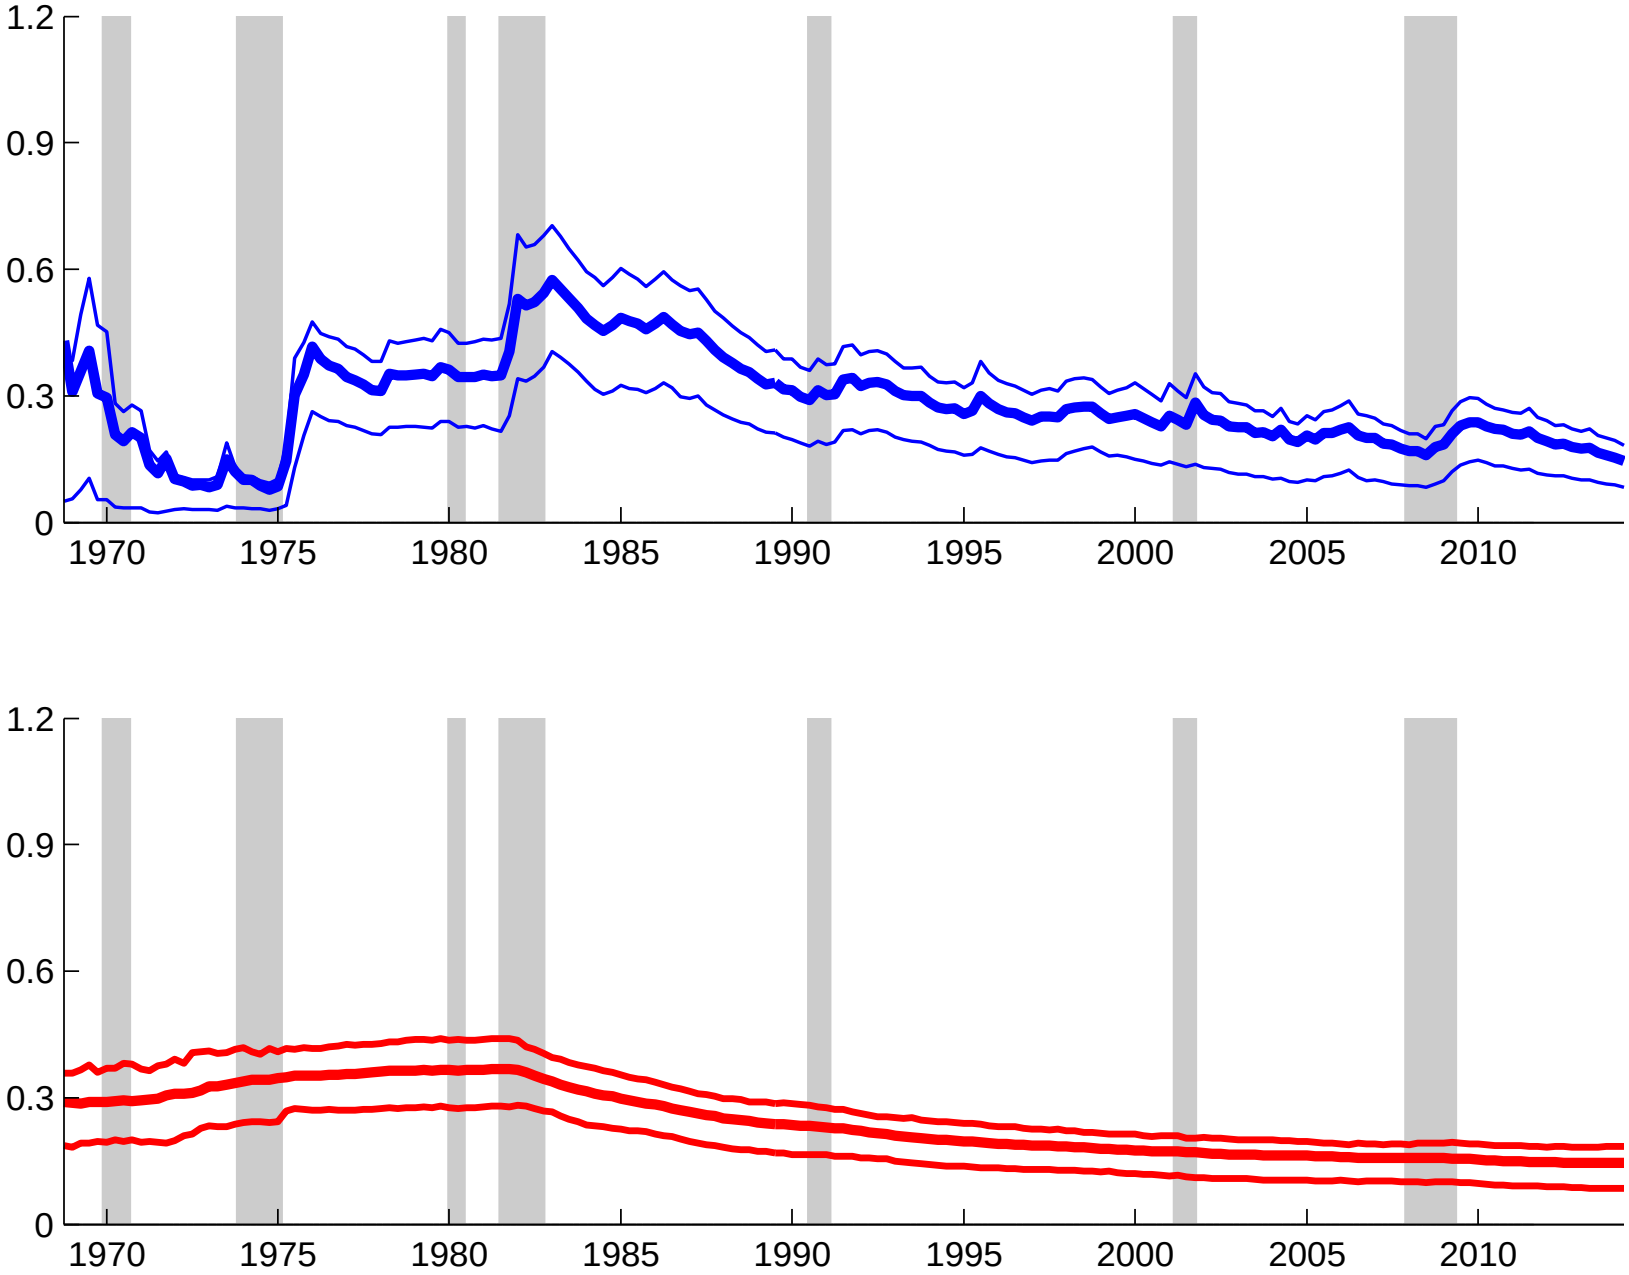

The plots of the filtered, $\varsigma_{\eta,t|t}$, and smoothed, $\varsigma_{\eta,t|T}$, SV of trend inflation are generated using the particle filter estimates of the TVP-SI law of motion and SW-UC-SV-TVP-AR(1) model. The thick and thin blue (red) plots are the filtered (smoothed) estimates and interquartile range coverage bands of $\varsigma_{\eta,t|t}$ $\left(\varsigma_{\eta,t|T}\right)$, respectively.

FIGURE 8: FILTERED, $\varsigma_{v,t|t}$, AND SMOOTHED, $\varsigma_{v,t|T}$, INFLATION GAP SV, 1968Q4 TO 2014Q2

FILTERED = BLUE AND SMOOTHED = RED

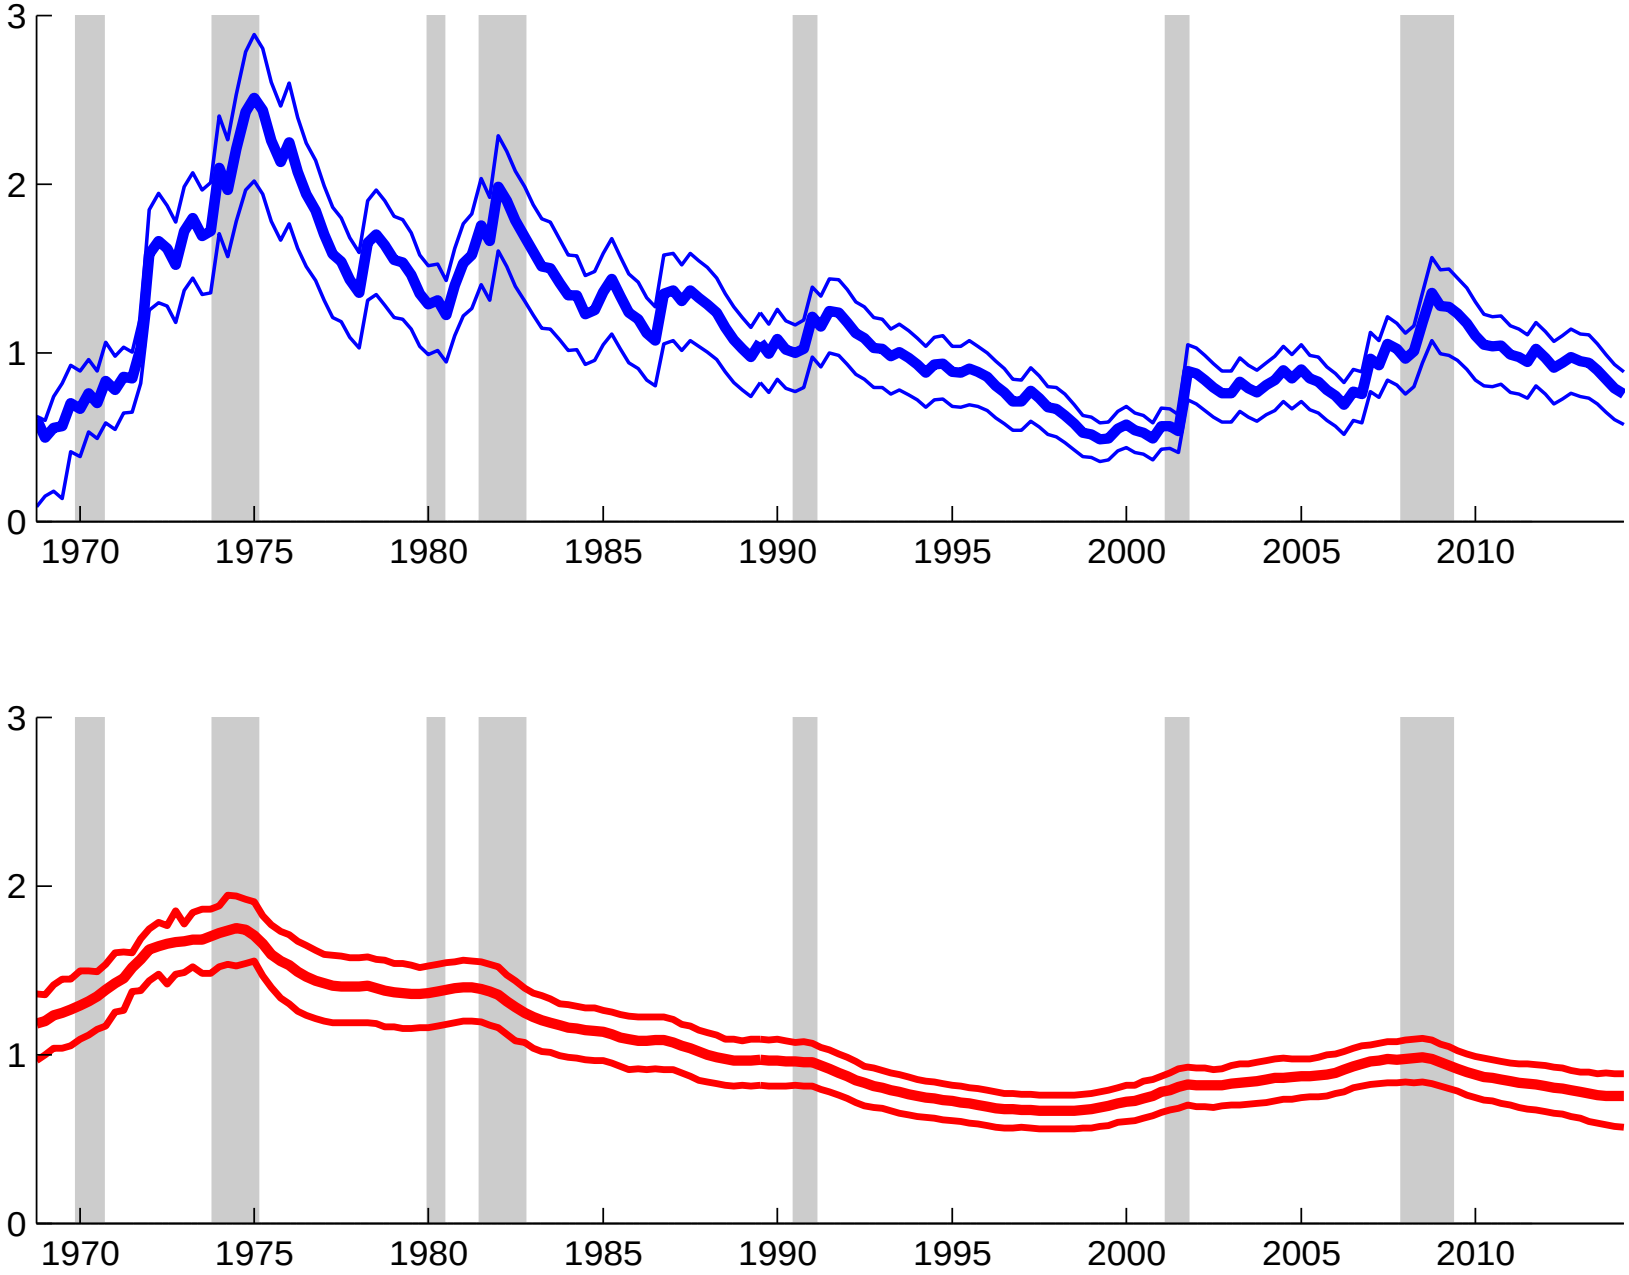

The plots of the filtered, $\varsigma_{v,t|t}$, and smoothed, $\varsigma_{v,t|T}$, SV of gap inflation are generated using the particle filter estimates of the TVP-SI law of motion and SW-UC-SV-TVP-AR(1) model. The thick and thin blue (red) plots are the filtered (smoothed) estimates and interquartile range coverage bands of $\varsigma_{v,t|t}$ $(\varsigma_{v,t|T})$, respectively.

Figure 9: Filtered, $\theta_{1,t|t}$, and Smoothed, $\theta_{1,t|T}$, Time-Varying Inflation Gap Persistence, 1968Q4 to 2014Q2

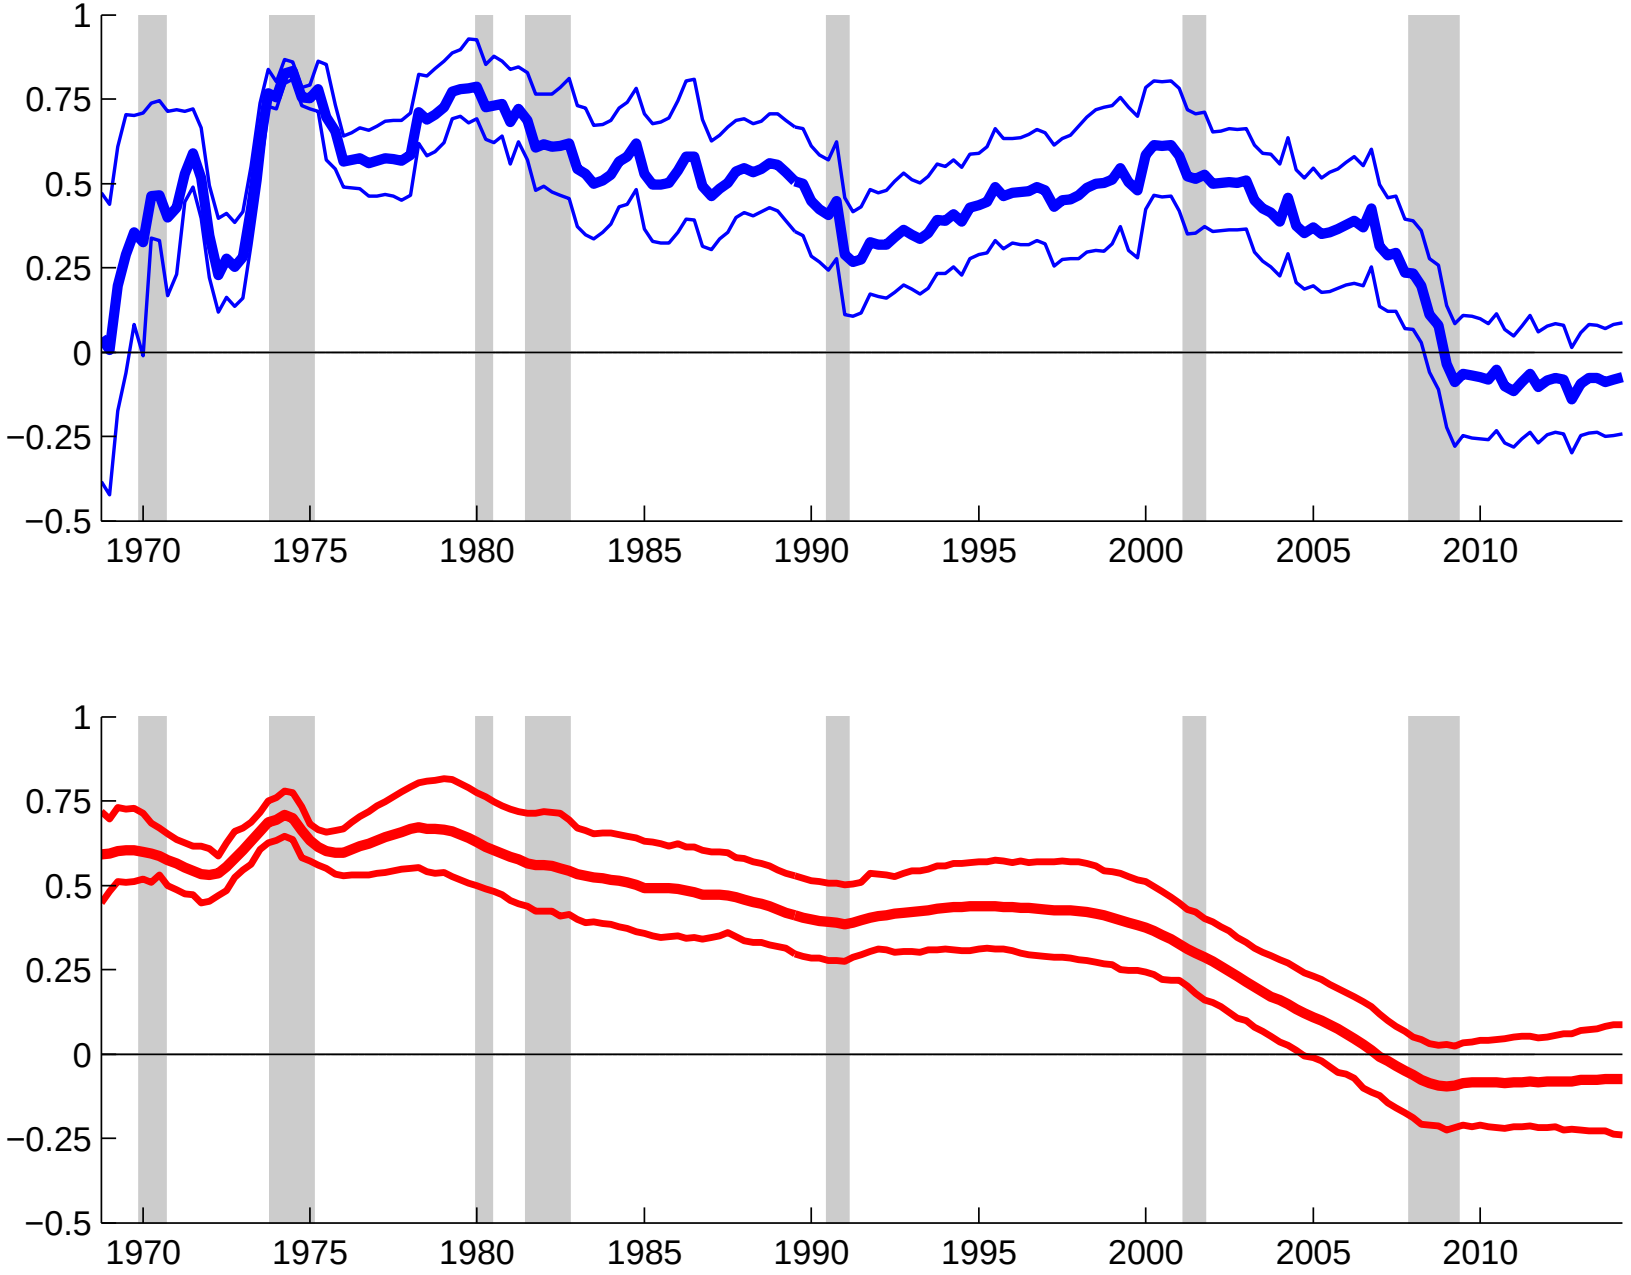

The plots of the filtered, $\theta_{1,t|t}$, and smoothed, $\theta_{1,t|T}$, TVP-AR1 coefficient are generated using the particle filter estimates of the TVP-SI law of motion and SW-UC-SV-TVP-AR(1) model. The thick and thin blue (red) plots are the filtered (smoothed) estimates and interquartile range coverage bands of $\theta_{1,t|t}$ ($\theta_{1,t|T}$), respectively.

Figure 10: Filtered, $\lambda_{t|t}$, and Smoothed, $\lambda_{t|T}$, Time-Varying SI Updating, 1968Q4 to 2014Q2

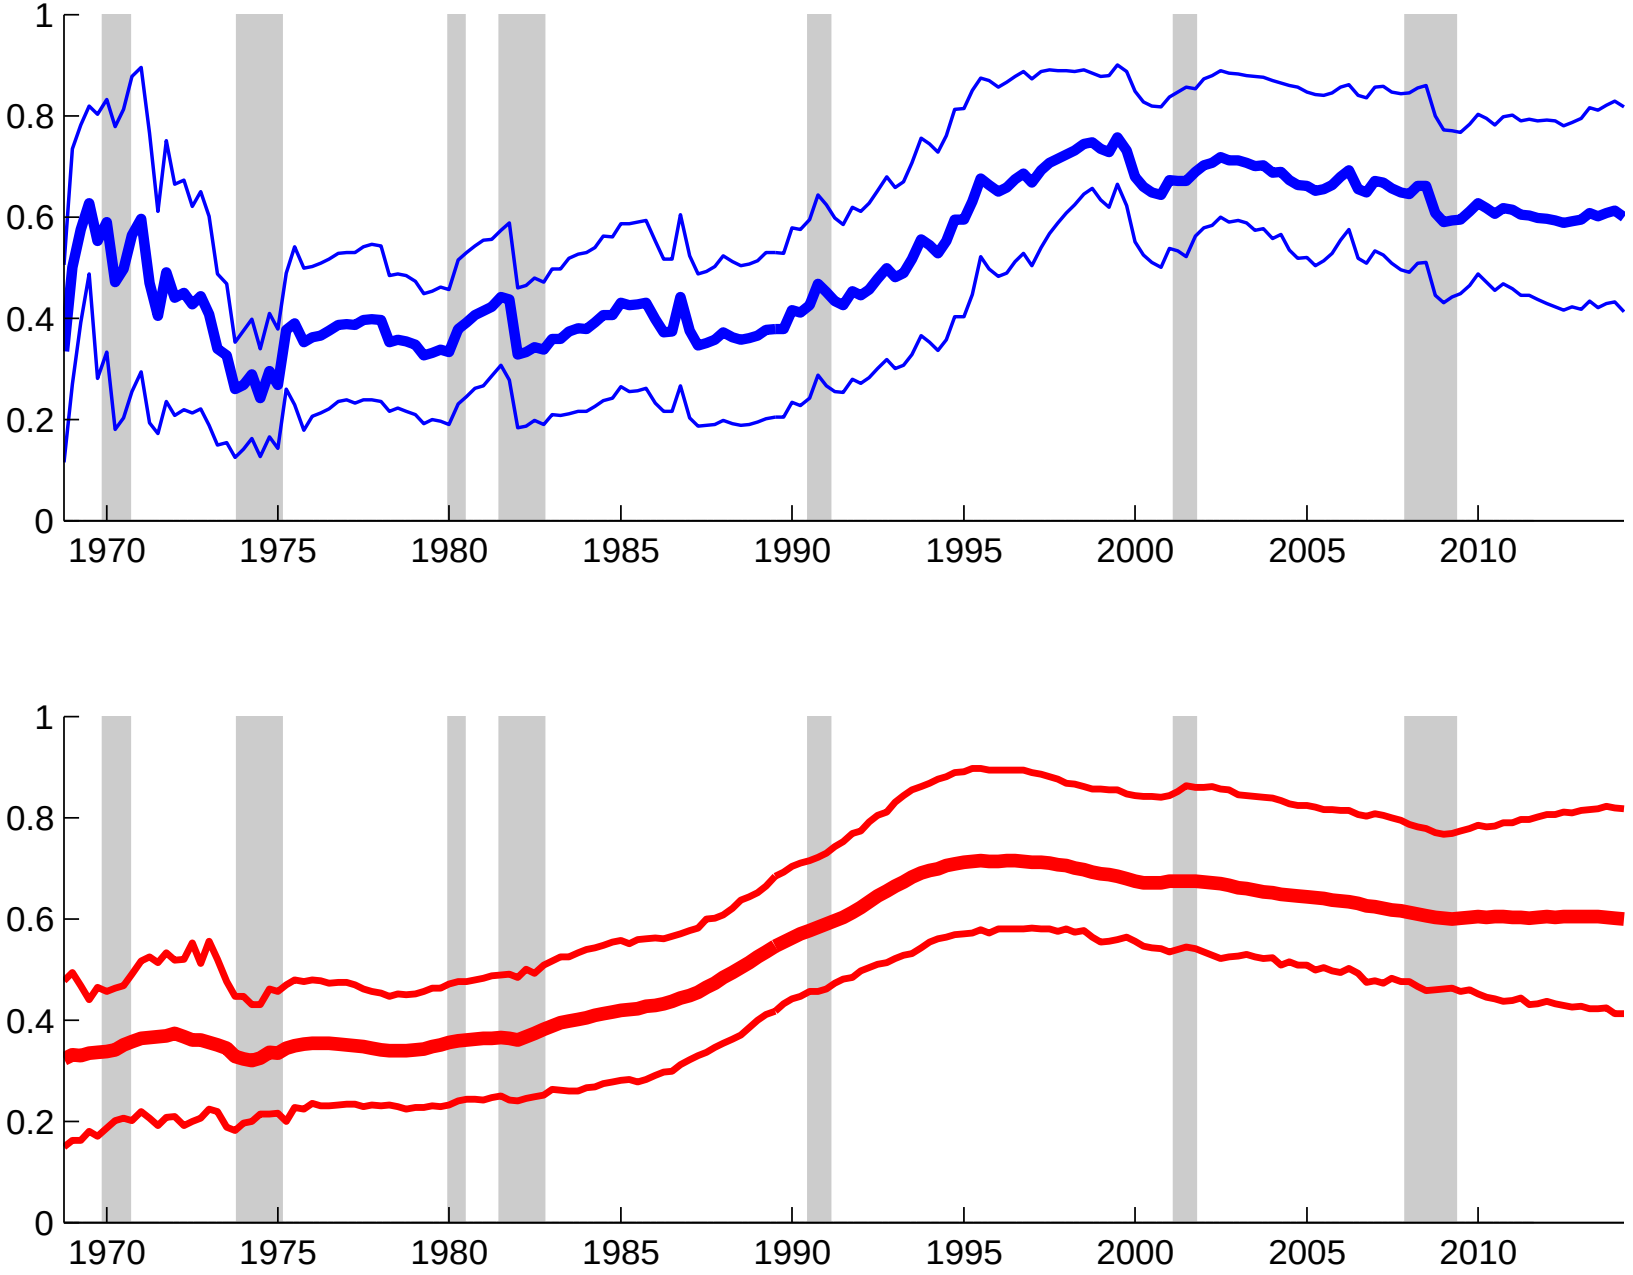

The plots of the filtered, $\lambda_{t|t}$, and smoothed, $\lambda_{t|T}$, TVP-AR1 coefficient are generated using the particle filter estimates of the TVP-SI law of motion and SW-UC-SV-TVP-AR(1) model. The thick and thin blue (red) plots are the filtered (smoothed) estimates and interquartile range coverage bands of $\lambda_{t|t}$ ($\lambda_{t|T}$), respectively.

FIGURE 11: SMOOTHED ESTIMATES OF $\lambda_{t|T}$, $\left|\theta_1\right|$, AND $\left|\theta_{1,t|T}\right|$, 1968$Q$4 TO 2014$Q$2

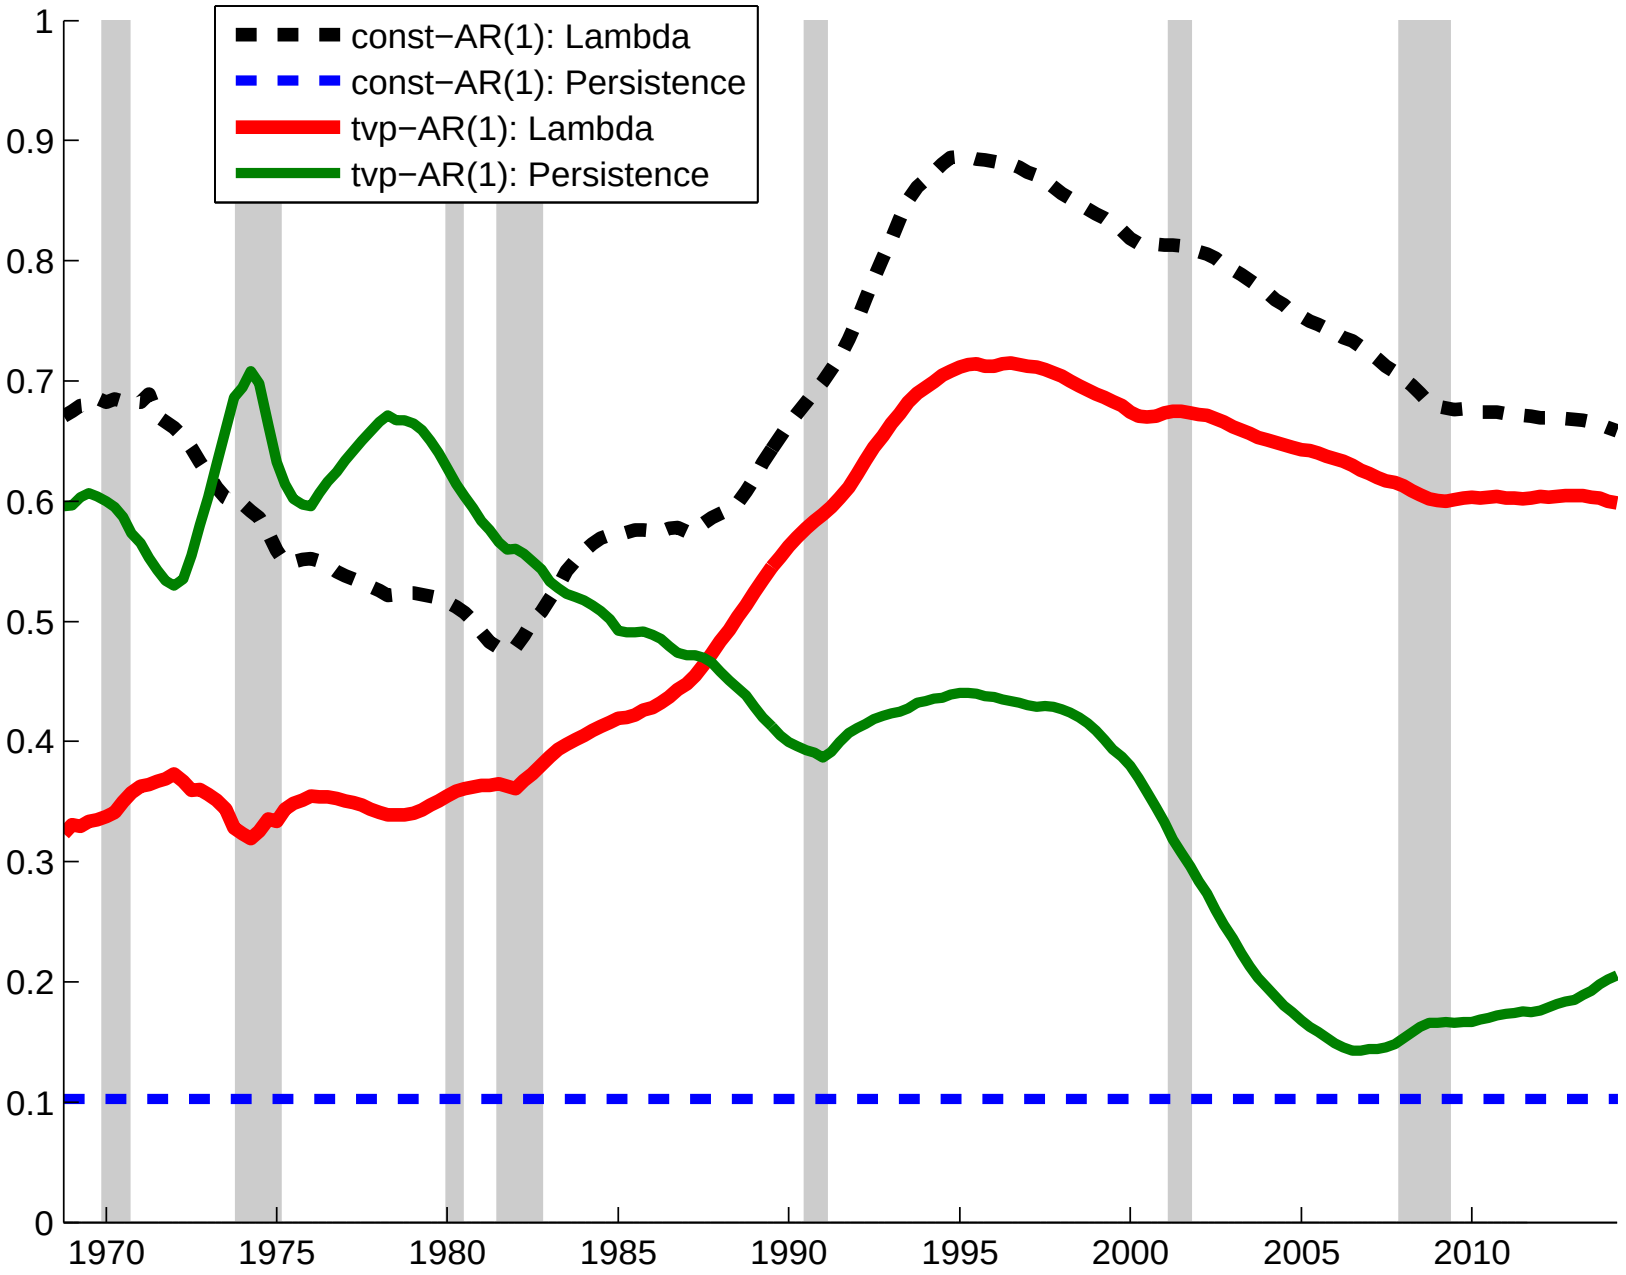

The absolute value of the fixed coefficient $\left|\theta_1\right|$ is the blue dotted line, which is the posterior mean of a MCMC simulator applied to the SW-UC-SV-AR(1) model. This model combined with the TVP-SI law of motion produces $\lambda_{t|T}$, which is the black dotted line plots. The solid green line is $\left|\theta_{1,t|T}\right|$, which equals the absolute values of the posterior means of the particle draws of this TVP moving through the sample quarter by quarter. This TVP is estimated within the SW-UC-SV-TVP-AR(1) model, which also is the source of the solid red plot of $\lambda_{t|T}$. Estimates of $\lambda_{t|T}$ are the posterior mean of the particle draws.