

INFLATION AND PROFESSIONAL FORECAST DYNAMICS: AN EVALUATION OF STICKINESS, PERSISTENCE, AND VOLATILITY*

ELMAR MERTENS[†] AND JAMES M. NASON[‡]

Current draft: February 16, 2018

First Draft: February 15, 2015

Abstract

This paper studies the joint dynamics of real-time U.S. inflation and average inflation predictions of the Survey of Professional Forecasters (SPF) based on sample ranging from 1968Q4 to 2017Q2. The joint data generating process (DGP) comprises an unobserved components (UC) model of inflation and a sticky information (SI) prediction mechanism for the SPF predictions. We add drifting gap inflation persistence to a UC model in which stochastic volatility (SV) affects trend and gap inflation. Another innovation puts a time-varying frequency of inflation forecast updating into the SI prediction mechanism. The joint DGP is a nonlinear state space model (SSM). We estimate the SSM using Bayesian tools grounded in a Rao-Blackwellized auxiliary particle filter, particle learning, and a particle smoother. The estimates show that (i) longer horizon average SPF inflation predictions inform estimates of trend inflation; (ii) gap inflation persistence is procyclical and SI inflation updating is frequent before the Volcker disinflation; and (iii) subsequently, gap inflation persistence turns countercyclical and SI inflation updating becomes infrequent.

JEL Classification Numbers: E31, C11, C32.

Key Words: inflation; unobserved components; professional forecasts; sticky information; stochastic volatility; time-varying parameters; Bayesian; particle filter.

[†]*email:* elmar.mertens@bis.org, *address:* Bank for International Settlements, Centralbahnplatz 2, CH 4051 Basel, Switzerland.

[‡]*email:* jmnason@ncsu.edu, *address:* Department of Economics, Campus Box 8110, NC State University, Raleigh, NC 27695-8110 and the Centre for Applied Macroeconomic Analysis.

*We thank Gregor Smith for several conversations that motivated this paper. We also received valuable comments from Todd Clark, Patrick Conway, Drew Creal, Bill Dupor, Andrew Filardo, Monica Jain, Alejandro Justiniano, and Wolfgang Lemke and suggestions from colleagues and participants at numerous seminars and conferences. Jim Nason thanks the Jenkins Family Economics Fund at North Carolina State University for financial support. The views herein are those of the authors and do not represent the views of the Bank for International Settlements.

1 Introduction

Central banks pay particular attention to inflation expectations. A good reason for this focus is that inflation expectations contain information about private agents' beliefs about the underlying factors driving observed inflation dynamics. We label these factors the inflation regime. For example, Bernanke (2007) argues that well anchored inflation expectations are necessary for a central bank to be able to stabilize inflation. However, since monetary policy makers lack direct knowledge of inflation expectations, they must infer such expectations from estimates of the inflation regime. These estimates often rely on realized inflation and combinations of financial market data, statistical and economic models, and forecast surveys.

This paper estimates inflation regimes from the joint data generating process (DGP) of realized inflation and the inflation predictions of professional forecasters grounded in a non-linear state space model (SSM). We tap a sample of inflation predictions from the Survey of Professional Forecasters (SPF) to extract the “beliefs” held by the average respondent about the (in)stability of the persistence, volatility, and stickiness of inflation. Average SPF inflation predictions are attractive for evaluating the SSM 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 forecasting performance suggests that average SPF inflation predictions coupled with realized inflation harbor useful information to measure inflation expectations.

We study the joint DGP of realized inflation, π_t , and average SPF inflation predictions by linking a Stock and Watson (2007) unobserved components (SW-UC) model of inflation to a version of the Mankiw and Reis (2002) sticky information (SI) model. The SW-UC model is useful for evaluating the impact of different types of shocks on inflation and inflation expectations. First, it decomposes π_t into trend inflation, τ_t , and gap inflation, ε_t , which restricts the impact of permanent and transitory shocks on π_t . When permanent shocks dominate movements in π_t , the inference is that inflationary expectations are not well anchored. The SW-UC model also imposes stochastic volatility (SV) on the innovations of τ_t and ε_t . Trend and gap SV creates nonlinearities in inflation dynamics, which produce bursts of volatility in π_t . Persistence is not

often imposed on ε_t when estimating the SW-UC-SV model. We depart from this assumption by giving ε_t drifting persistence in the form of a time-varying parameter first-order autoregression, or a TVP-AR(1). Drifting gap persistence is another source of nonlinearity in a SW-UC model, which can exhibit pro- or countercyclical changes. We label the extended version of the DGP of π_t as the SW-UC-SV-TVP-AR(1) model.

Coibion and Gorodnichenko (2015) adapt a SI model to a setup in which forecasters update their rational expectations (RE) information set with a fixed probability $1-\lambda$. Averaging across forecasters defines the h -step ahead SI inflation prediction, $F_t\pi_{t+h}$, $h = 1, \dots, \mathcal{H}$. The result is that the SI inflation prediction evolves as a weighted average of the lagged SI forecast, $F_{t-1}\pi_{t+h}$, and a RE inflation forecast, $\mathbf{E}_t\pi_{t+h}$, where the weights are λ and $1-\lambda$. The result is the SI law of motion $F_t\pi_{t+h} = \lambda F_{t-1}\pi_{t+h} + (1-\lambda)\mathbf{E}_t\pi_{t+h}$, where $F_t\pi_{t+h}$ updates at the frequency $1/(1-\lambda)$ on average. In this reading, λ reflects the average forecaster's beliefs about the persistence or stickiness of the inflation regime.¹

We innovate on the Coibion-Gorodnichenko static coefficient SI-law of motion by investing λ with drift. The result is a nonlinear SI-law of motion $F_t\pi_{t+h} = \lambda_t F_{t-1}\pi_{t+h} + (1-\lambda_t)\mathbf{E}_t\pi_{t+h}$, where the TVP-SI parameter, λ_t , evolves as an exogenous and bounded random walk (RW), $\lambda_{t+1} = \lambda_t + \sigma_\kappa \kappa_{t+1}$, and its innovation is drawn from a truncated normal distribution (\mathcal{TN}), $\kappa_{t+1} \sim \mathcal{TN}(0, 1; \lambda_{t+1} \in (0, 1))$. The SI forecaster's information set includes the innovation κ_t when $F_{t-1}\pi_{t+h}$ is updated to $F_t\pi_{t+h}$, which implies that λ_t is also part of this information set.

A motivation for placing λ_t in the SI-law of motion is to uncover evidence about changes in the beliefs that the average SPF participant holds about the inflation regime. Changes in these beliefs are embedded in observed movements of the average SPF participant's h -step ahead inflation prediction, $\pi_{t,t+h}^{SPF}$. We relate $\pi_{t,t+h}^{SPF}$ to $F_t\pi_{t,t+h}$ by adding a classical measurement error, $\zeta_{t,h}$, to set $\pi_{t,t+h}^{SPF} = F_t\pi_{t,t+h} + \sigma_{\zeta_h}\zeta_{h,t}$, where $\zeta_{h,t} \sim \mathcal{N}(0, 1)$, $h = 1, \dots, \mathcal{H}$. The $\pi_{t,t+h}^{SPF}$ observation equation, SI-law of motion, and RW of λ_t form the SI-prediction mechanism.

The joint DGP of the SI-prediction mechanism and SW-UC-SV-TVP-AR(1) model maps shocks

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

to τ_t , ε_t , and SI state variables into movements in π_t and $\pi_{t,t+h}^{SPF}$.² Estimates of the joint DGP provide evidence about drift in λ_t and its co-movement with the SVs of τ_t and ε_t and drifting persistence in ε_t . If λ_t exhibits meaningful statistical and economic time variation and it moves with the SVs or drifting inflation gap persistence, we find evidence that shifts in SI inflation updating are attuned to the hidden factors driving the inflation regime.

Another contribution is the sequential Monte Carlo (SMC) methods that we use to estimate the joint DGP of the SI-prediction mechanism and SW-UC-SV-TVP-AR(1) model. These methods consist of the particle learning estimator (PLE) of Storvik (2002) and the particle smoother (PS) of Lindsten, Bunch, Särkkä, Schön, and Godsill (2016). The PLE and PS rely on a Rao-Blackwellized auxiliary particle filter (RB-APF). Our joint DGP is susceptible to Rao-Blackwellization because τ_t , ε_t , and the SI state variables form a linear SSM for given realizations of the nonlinear state variables — which are trend and gap inflation SVs, drifting inflation persistence, and λ_{t+1} , and estimates of the static coefficients of the SI-prediction mechanism and SW-UC-SV-TVP-AR(1) model. Applying the Kalman filter (KF) produces estimates of the distribution of the conditionally linear states that are integrated analytically, which increases the efficiency of the RB-APF. The RB-APF estimates the nonlinear states by simulation.

We estimate the joint DGP of the SI-prediction mechanism and SW-UC-SV-TVP-AR(1) model on a quarterly sample from 1968Q4 to 2017Q2. The sample matches π_t with the GNP or GDP deflator inflation available to the SPF in real time at date t . The average SPF inflation prediction is denoted $\pi_{t,t+h}^{SPF}$, where $\mathcal{H} = 5$ or $h = 1, \dots, 1$ - to 5-quarter ahead forecast horizons.³

Given only a sample of $\{\pi_t, \pi_{t,t+1}^{SPF}, \dots, \pi_{t,t+5}^{SPF}\}_{t=1}^T$ and our priors, the SSM yields posterior estimates of the beliefs that the average SPF participant has about the hidden factors underlying the inflation regime. Our estimates of trend inflation are aligned with average SPF inflation

²Our approach to studying the joint dynamics of π_t and $\pi_{t,t+h}^{SPF}$ builds on Kozicki and Tinsley (2012), Mertens (2016), and Nason and Smith (2016a, b).

³The SPF contains average predictions of the GNP or GDP deflator for a nowcast and forecasts up to 4-quarters ahead. The surveys are collected at the middle of the quarter, which suggests $\pi_{t,t+h}^{SPF}$ is not based on full knowledge of π_t . We treat $\pi_{t,t+h}^{SPF}$ as reflecting only information available through the end of the previous quarter. This identifies the average SPF nowcast, 1-quarter, \dots , 4-quarter ahead predictions with $\pi_{t,t+h}^{SPF}$, $h = 1, 2, \dots, 5$. We discuss these timing issues in section 4.1.

predictions, especially at longer horizons. Gap inflation is more volatile before the Volcker disinflation than afterwards. There is a spike in gap inflation SV during the 1973–75 recession while trend inflation SV displays peaks during the 1981–82 and 2007–09 recessions. The drift in gap inflation persistence is procyclical before the Volcker disinflation, turns countercyclical afterwards, disappears by the 2007–09 recession, and returns to pre-2000 rates by 2014. The average SPF participant updates SI inflation forecasts frequently from the late 1960s to 1988. The frequency of SI inflation updating falls from 1990 to 1995 and then remains steady until 2007. During the 2007–09 recession, SI inflation updating occurs more frequently and drops slowly afterwards. Thus, movements in the frequency of SI inflation updating displays co-movement with trend inflation, its SV, and drifting inflation persistence. We conclude that the beliefs of the average SPF respondent are sensitive to the impact of permanent shocks on the conditional mean of inflation and that the Volcker disinflation marks the moment at which the behavior of trend inflation, its SV, and the cyclicity of the drift in inflation gap persistence changed.

The structure of the paper is as follows. In section 2, we build a SSM of the joint DGP of π_t and $\pi_{t,t+h}^{SPF}$, $h = 1, \dots, \mathcal{H}$. Section 3 discusses the SMC methods used to estimate the SSM. Results appear in section 4. Section 5 offers our conclusions.

2 Statistical and Econometric Models

This section describes the statistical and economic models used to estimate the joint dynamics of π_t and $\pi_{t,t+h}^{SPF}$, $h = 1, \dots, \mathcal{H}$. Stock and Watson (2007) is the source of the statistical model to which we add drifting persistence to ε_t . The economic model is a SI-prediction mechanism that has a SI-TVP parameter. Drift in inflation persistence and the frequency of SI inflation updating create nonlinearities in the state transition dynamics of the SSM. The SI-TVP also interacts with trend and gap inflation SVs to produce nonlinearities in the impulse structure of the SSM.⁴

⁴We relegate to an online appendix construction of a SSM in which persistence in ε_t is a AR(1) with a static slope coefficient. The online appendix is available at <http://www.e1marmertens.com/>.

2.1 The SW-UC Model

The SW-UC model generates π_t . Stock and Watson (2010), Creal (2012), Shephard (2013), Cogley and Sargent (2015), and Mertens (2016) have estimated versions of the model in which SV in innovations to τ_t and ε_t is the source of nonlinearity in π_t . We add an additional nonlinearity to the SW-UC-SV model in the form of drift in the persistence of ε_t created by a TVP-AR(1). We collect these features into the SW-UC-SV-TVP-AR(1) model

$$\pi_t = \tau_t + \varepsilon_t + \sigma_{\zeta,\pi}\zeta_{\pi,t}, \quad \zeta_{\pi,t} \sim \mathcal{N}(0, 1), \quad (1.1)$$

$$\tau_{t+1} = \tau_t + \zeta_{\eta,t+1}\eta_t, \quad \eta_t \sim \mathcal{N}(0, 1), \quad (1.2)$$

$$\varepsilon_{t+1} = \theta_{t+1}\varepsilon_t + \zeta_{\nu,t+1}\nu_t, \quad \nu_t \sim \mathcal{N}(0, 1), \quad (1.3)$$

$$\ln \zeta_{\ell,t+1}^2 = \ln \zeta_{\ell,t}^2 + \sigma_{\ell}\xi_{\ell,t+1}, \quad \xi_{\ell,t+1} \sim \mathcal{N}(0, 1), \quad \ell = \eta, \nu, \quad (1.4)$$

$$\theta_{t+1} = \theta_t + \sigma_{\phi}\phi_{t+1}, \quad \phi_{t+1} \sim \mathcal{N}(0, 1), \quad (1.5)$$

where measurement error on π_t , $\zeta_{\pi,t}$, is uncorrelated with τ_t and ε_t and the innovations, η_t and ν_t , these innovations are afflicted by SV, which evolve as RWs in $\ln \zeta_{\eta,t+1}^2$ and $\ln \zeta_{\nu,t+1}^2$, drifting persistence in ε_{t+1} is tied to θ_{t+1} , restricting the RW of $\theta_{t+1} \in (-1, 1)$ ensures stationarity of ε_{t+1} at each date $t+1$, and innovations to the linear state variables, η_t and ν_t , and innovations to nonlinear state variables, $\xi_{\eta,t+1}$, $\xi_{\nu,t+1}$, and ϕ_{t+1} , are uncorrelated.

A special case of the SW-UC-SV-TVP-AR(1) model gives a result about forecasting traced to Muth (1960). Shut down SV, $\sigma_{\eta} = \zeta_{\eta,t}$ and $\sigma_{\nu} = \zeta_{\nu,t}$, and eliminate gap inflation persistence, $\theta_t = 0$, for all dates t . The result is a fixed coefficient SW-UC model with an IMA(1, 1) reduced form, $(1 - \mathbf{L})\pi_t = (1 - \varpi\mathbf{L})\nu_t$, where the MA1 coefficient $\varpi \in (-1, 1)$, \mathbf{L} is the lag operator, $\pi_{t-1} = \mathbf{L}\pi_t$, and the one-step ahead forecast error $\nu_t = \eta_t + \varepsilon_t + \tau_t - \tau_{t-1|t-1}$.⁵ The IMA(1, 1) implies a RE inflation updating equation, $\mathbf{E}\{\pi_{t+1} | \pi^t, \sigma_{\eta}, \sigma_{\nu}\} = (1 - \varpi)\pi_t + \varpi\mathbf{E}\{\pi_t | \pi^{t-1}, \sigma_{\eta}, \sigma_{\nu}\}$, where π^t is the date t history of inflation, π_t, \dots, π_1 .

⁵Stock and Watson (2007), Grassi and Proietti (2010), and Shephard (2013) tie ϖ to the autocovariance functions (ACFs) of the IMA(1, 1) and fixed coefficient SW-UC model. At lags zero and one, the ACFs set $(1 + \varpi^2)\sigma_{\nu}^2 = \sigma_{\eta}^2 + 2\sigma_{\varepsilon}^2$ and $-\varpi\sigma_{\nu}^2 = -\sigma_{\varepsilon}^2$. Substitute for σ_{ν}^2 to find $\varpi^2 - (2 + \sigma_{\eta}^2/\sigma_{\varepsilon}^2)\varpi + 1 = 0$. The solution is $\varpi = [1 + 0.5\sigma_{\eta}^2/\sigma_{\varepsilon}^2] - \frac{\sigma_{\eta}}{\sigma_{\varepsilon}}\sqrt{1 + 0.25\sigma_{\eta}^2/\sigma_{\varepsilon}^2}$, given $\varpi \in (-1, 1)$ and $\sigma_{\eta}, \sigma_{\varepsilon} > 0$.

Stock and Watson (2007), Grassi and Prioretto (2010) and Shephard (2013) note the SW-UC-SV model replaces ϖ with the time-varying local weight ϖ_t in the reduced form IMA(1, 1). The result is an exponentially weighted moving average (EWMA) updating recursion or smoother

$$\mathbf{E} \left\{ \boldsymbol{\pi}_{t+1} \mid \boldsymbol{\pi}^t, \zeta_{\eta,t}, \zeta_{\nu,t} \right\} = \sum_{j=0}^{\infty} \mu_{\varpi,t-j} \left(\prod_{\ell=0}^j \varpi_{t-\ell} \right) \boldsymbol{\pi}_{t-j}, \quad (2)$$

in which the discount ϖ_t adjusts to changes in the latest data, where $\mu_{\varpi,t} = (1 - \varpi_t)/\varpi_t$.

2.2 The SI-Prediction Mechanism

This section begins by reproducing the SPF observation equation, the nonlinear SI-law of motion, and the random walk law of motion of λ_t . These elements form the system of equations

$$\boldsymbol{\pi}_{t,t+h}^{SPF} = F_t \boldsymbol{\pi}_{t+h} + \sigma_{\zeta,h} \boldsymbol{\zeta}_{h,t}, \quad \boldsymbol{\zeta}_{h,t} \sim \mathcal{N}(0, 1), \quad (3.1)$$

$$F_t \boldsymbol{\pi}_{t+h} = \lambda_t F_{t-1} \boldsymbol{\pi}_{t+h} + (1 - \lambda_t) \mathbf{E}_t \boldsymbol{\pi}_{t+h}, \quad h = 1, \dots, \mathcal{H}, \quad (3.2)$$

$$\lambda_{t+1} = \lambda_t + \sigma_{\kappa} \kappa_{t+1}, \quad \kappa_{t+1} \sim \mathcal{N}(0, 1), \quad (3.3)$$

where $\mathbf{E}_t \boldsymbol{\pi}_{t+h}$ is conditional on the average SPF participant's statistical model of inflation and $\lambda_t \in (0, 1)$ for all dates t . Equations (3.1)–(3.3) define the SI-prediction mechanism through which shocks to λ_t and movements in other state variables generate fluctuations in $\boldsymbol{\pi}_{t,t+h}^{SPF}$.

The SI-law of motion (3.2) implies an EWMA smoother. Iterate (3.2) backwards, substitute the result into (3.2), and repeat the process many times to produce the SI-EWMA smoother

$$F_t \boldsymbol{\pi}_{t+h} = \sum_{j=0}^{\infty} \mu_{\lambda,t-j} \left(\prod_{\ell=0}^j \lambda_{t-\ell} \right) \mathbf{E}_{t-j} \boldsymbol{\pi}_{t+h}, \quad (4)$$

where the discount rate is the SI-TVP, λ_t , and $\mu_{\lambda,t} = (1 - \lambda_t)/\lambda_t$. The SI-EWMA smoother (4) nests the RE forecast, $\lim_{\lambda_t \rightarrow 0} F_t \boldsymbol{\pi}_{t+h} = \mathbf{E}_t \boldsymbol{\pi}_{t+h}$, and the pure SI update, $\lim_{\lambda_t \rightarrow 1} F_t \boldsymbol{\pi}_{t+h} = \sum_{j=1}^{\infty} \mu_{\lambda,t-j} \left(\prod_{\ell=1}^j \lambda_{t-\ell} \right) \mathbf{E}_{t-j} \boldsymbol{\pi}_{t+h}$. The former limit shuts down SI as λ_t falls to zero because the discount on $\mathbf{E}_{t-j} \boldsymbol{\pi}_{t+h}$ increases with j . In this case, SI inflation forecast updates rely only on $\mathbf{E}_t \boldsymbol{\pi}_{t+h}$ period by period. At the other extreme, less weight is placed on $\mathbf{E}_t \boldsymbol{\pi}_{t+h}$ and more on $\mathbf{E}_{t-j} \boldsymbol{\pi}_{t+h}$, $j > 1$, as λ_t rises to one. Thus, $F_{t-1} \boldsymbol{\pi}_{t+h}$ summarizes the SI inflation forecast.

Between these polar cases, shocks to λ_t alter the discount applied to the history of $\mathbf{E}_t \pi_{t+h}$ in the SI-EWMA smoother (4). This information aids in identifying movements in $\pi_{t,t+h}^{SPF}$ with respect to innovations in λ_t . The EWMA smoother (2) shows a similar relationship exists between $\mathbf{E}_t \pi_{t+h}$, π_t , and the time-varying discount generated by $\zeta_{\eta,t}$, $\zeta_{\nu,t}$, and θ_t . This gives us several sources of information to identify movements in π_t and $\pi_{t,t+h}^{SPF}$ within the joint DGP of the SI-prediction mechanism and the SW-UC-SV-TVP-AR(1) model.⁶

2.3 The State Space Model of the Joint DGP

Drift in inflation gap persistence complicates building a SSM for the joint DGP of the SI-prediction mechanism and SW-UC-SV-TVP-AR(1) model. The SSM rests on the RE and SI term structures of inflation forecasts for which the latent factors are the RE state variables $\mathcal{X}_t = [\tau_t \ \varepsilon_t]'$ and SI analogues $F_t \mathcal{X}_t = [F_t \tau_t \ F_t \varepsilon_t]'$. The problem is the law of iterated expectation (LIE) cannot be employed to create predictions of \mathcal{X}_{t+h} or $F_t \mathcal{X}_{t+h}$ because forecasts of θ_t are needed. Instead, we construct RE and SI term structures of inflation forecasts in the presence of drifting gap inflation persistence by invoking the anticipated utility model (AUM).

The RE term structure of inflation forecasts is based on the observation and state equations of the SW-UC-SV-TVP-AR(1) model. The observation equation (1.1) of the SW-UC-SV-TVP-AR(1) model links π_t to τ_t , ε_t , and $\zeta_{\pi,t}$, which can be rewritten as

$$\pi_t = \delta_{\mathcal{X}} \mathcal{X}_t + \sigma_{\zeta,\pi} \zeta_{\pi,t}, \quad (5.1)$$

where $\delta_{\mathcal{X}} = [1 \ 1]$. The state equations of the SW-UC-SV-TVP-AR(1) model are created by placing the random walk (1.3) below the TVP-AR(1) (1.3)

$$\mathcal{X}_{t+1} = \Theta_{t+1} \mathcal{X}_t + \mathbf{Y}_{t+1} \mathcal{W}_t, \quad (5.2)$$

where $\Theta_{t+1} = \begin{bmatrix} 1 & 0 \\ 0 & \theta_{t+1} \end{bmatrix}$, $\mathbf{Y}_{t+1} = \begin{bmatrix} \zeta_{\eta,t+1} & 0 \\ 0 & \zeta_{\nu,t+1} \end{bmatrix}$, and $\mathcal{W}_t = \begin{bmatrix} \eta_t \\ \nu_t \end{bmatrix}$. The transition dy-

⁶Krane (2011), Coibion and Gorodnichenko (2012), and Jain (2013) use forecast revisions to identify the responses of professional forecasters to disparate shocks, which is an alternative to our approach.

namics of the state equations (5.2) are nonlinear in Θ_{t+1} . These nonlinearities rule out applying the LIE to construct the RE term structure of inflation forecasts.

We appeal to two aspects of the AUM to solve the problem. The AUM resurrects the LIE by (i) assuming agents are ignorant of the true DGP and (ii) treating the TVPs of the joint DGP as fixed (locally) at each date t .⁷ These assumptions are instructions to hold TVPs at their date t values within RE and SI forecasts that condition on date t information.⁸

The RE term structure of inflation forecasts is easy to construct under the AUM. First, generate h -step ahead RE forecasts of \mathcal{X}_t by iterating the state equations (5.2) forward, apply the expectations operator, and invoke the AUM to find $\mathbf{E}_t \Theta_{t+h} \mathcal{X}_{t+h} = \Theta_{t|t}^h \mathcal{X}_t$. Next, push the observation equation (5.1) h -steps ahead and take expectations to obtain the RE term structure of inflation forecasts under the SW-UC-SV-TVP-AR(1) model and AUM

$$\mathbf{E}_t \pi_{t+h} = \delta_x \Theta_{t|t}^h \mathcal{X}_t. \quad (6)$$

The AUM restricts the impact of drifting inflation gap persistence on these RE forecasts by conditioning θ_t on date t information.

Next, we show the SI term structure of inflation forecasts is built on the SI-EWMA formula (4), RE term structure of inflation forecasts (6), and a conjecture about the law of motion of the SI vector $F_{t+1} \mathcal{X}_{t+1}$. The SI-EWMA formula (4) depends on the RE inflation forecasts $\mathbf{E}_{t-j} \pi_{t+h}$. Since these RE forecasts are $\mathbf{E}_{t-j} \pi_{t+h} = \delta_x \Theta_{t|t}^{h+j} \mathcal{X}_{t-j}$ under the AUM, other RE forecasts are needed to replace $\mathbf{E}_{t-j} \pi_{t+h}$ in the SI-EWMA smoother (4). Our solution assumes the average member of the SPF fixes drift in inflation gap persistence at its current value when iterating SI recursions backwards. Under this assumption, $\mathbf{E}_{t-j} \pi_{t+h} = \delta_x \Theta_{t|t}^{h+j} \mathcal{X}_{t-j}$ in the SI-EWMA smoother (4). The result is $F_t \pi_{t+h} = \delta_x \Theta_{t|t}^h \sum_{j=0}^{\infty} \mu_{\lambda,t-j} \left(\prod_{\ell=0}^j \lambda_{t-\ell} \right) \Theta_{t|t}^j \mathcal{X}_{t-j}$. Next, we conjecture the law of motion of the SI state vector is $F_t \mathcal{X}_{t+h} = (1 - \lambda_t) \mathbf{E}_t \mathcal{X}_{t+h} + \lambda_t F_{t-1} \mathcal{X}_{t+h}$. An implication is the SI-EWMA smoother $F_t \mathcal{X}_{t+h} = \sum_{j=0}^{\infty} \mu_{\lambda,t-j} \left(\prod_{\ell=0}^j \lambda_{t-\ell} \right) \mathbf{E}_{t-j} \mathcal{X}_{t+h}$. Condition on the date t drift in inflation

⁷For example, Cogley and Sbordone (2008) employ the AUM model to study the dynamics of trend and gap inflation within a TVP-new Keynesian Phillips curve.

⁸The AUM assumptions result in decision making that is consistent with Bayesian forecasting, according to Cogley and Sargent (2008). They also note Kreps (1998) argues agents engaging in AUM-like behavior are acting rationally when seeing through to the true model is costly.

gap persistence on date t information, $\theta_{t|t}$, to find $F_t \mathcal{X}_{t+h} = \sum_{j=0}^{\infty} \mu_{\lambda,t-j} \left(\prod_{\ell=0}^j \lambda_{t-\ell} \right) \Theta_{t|t}^{h+j} \mathcal{X}_{t-j}$. When $h = 0$, we have $F_t \mathcal{X}_t = \sum_{j=0}^{\infty} \mu_{\lambda,t-j} \left(\prod_{\ell=0}^j \lambda_{t-\ell} \right) \Theta_{t|t}^j \mathcal{X}_{t-j}$. By combining the SI-EWMA smoothers of $F_t \pi_{t+h}$ and $F_t \mathcal{X}_t$, the SI term structure of inflation forecasts is the result

$$F_t \pi_{t+h} = \delta_x \Theta_{t|t}^h F_t \mathcal{X}_t. \quad (7)$$

The online appendix has details about the SI term structure of inflation forecasts (7).

The online appendix also develops state equations for $F_{t+1} \mathcal{X}_{t+1}$. Remember the SI-EWMA smoother of $F_t \mathcal{X}_t$ is $\sum_{j=0}^{\infty} \mu_{\lambda,t-j} \left(\prod_{\ell=0}^j \lambda_{t-\ell} \right) \Theta_{t|t}^j \mathcal{X}_{t-j}$, which by induction gives a law of motion, $F_t \mathcal{X}_t = (1 - \lambda_t) \mathcal{X}_t + \lambda_t \Theta_{t|t} F_{t-1} \mathcal{X}_{t-1}$. We create state equations for $F_{t+1} \mathcal{X}_{t+1}$ by pushing this law of motion forward a period and substituting for \mathcal{X}_{t+1} with the state equations (5.2) of the SW-UC-SV-TVP-AR(1) model. Stack the latter equations on top of the former to obtain the state equations of the SSM of the joint DGP

$$S_{t+1} = \mathcal{A}_{t+1} S_t + \mathcal{B}_{t+1} \mathcal{W}_t, \quad (8.1)$$

where $S_t = \begin{bmatrix} \mathcal{X}'_t \\ F_t \mathcal{X}'_t \end{bmatrix}$, $\mathcal{A}_{t+1} = \begin{bmatrix} \Theta_{t+1} & \mathbf{0}_{2 \times 2} \\ (1 - \lambda_{t+1}) \Theta_{t+1} & \lambda_{t+1} \Theta_{t+1} \end{bmatrix}$, $\mathcal{B}_{t+1} = \begin{bmatrix} \mathbf{Y}_{t+1} \\ (1 - \lambda_{t+1}) \mathbf{Y}_{t+1} \end{bmatrix}$, and the conditioning time subscript on Θ_{t+1} is dropped. The state equations (8.1) show shocks to λ_t alter the transition and impulse dynamics only of $F_t \tau_t$ and $F_t \varepsilon_t$. Changes in θ_t shift the transition dynamics of all elements of S_t while its impulse dynamics react to $\zeta_{\eta,t}$, and $\zeta_{\nu,t}$.

We complete the SSM by constructing its observation equations. First, replace $F_t \pi_{t+h}$ in the SPF measurement equation (3.1) with the SI term structure of inflation forecasts (7) for $h = 1, \dots, \mathcal{H}$. Place the results below the observation equation (5.1) of the SW-UC-SV-TVP-AR(1) model to form the SSM's observation equations

$$\mathcal{Y}_t = \mathcal{C}_t S_t + \mathcal{D} \mathcal{U}_t, \quad (8.2)$$

$$\text{where } \mathcal{Y}_t = \begin{bmatrix} \pi_t \\ \pi_{1,t}^{SPF} \\ \vdots \\ \pi_{\mathcal{H},t}^{SPF} \end{bmatrix}, \quad \mathcal{C}_t = \begin{bmatrix} \delta_x & \mathbf{0}_{1 \times 2} \\ \mathbf{0}_{1 \times 2} & \delta_x \Theta_t \\ \vdots & \vdots \\ \mathbf{0}_{1 \times 2} & \delta_x \Theta_t^{\mathcal{H}} \end{bmatrix}, \quad \mathcal{D} = \begin{bmatrix} \sigma_{\zeta,\pi} & 0 & \dots & 0 \\ 0 & \sigma_{\zeta,1} & \dots & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & \dots & \sigma_{\zeta,\mathcal{H}} \end{bmatrix}, \quad \mathcal{U}_t =$$

$[\zeta_{\pi,t} \ \zeta_{1,t} \ \dots \ \zeta_{\mathcal{I},t}]'$, and $\mathbf{\Omega}_u = \mathcal{D}\mathcal{D}'$. The SSM integrates $F_t\pi_{t+h}$ out of the observation equations (8.2) using the SI term structure of inflation forecasts (7). As a result, $F_t\varepsilon_t$ produces mean reversion in $\pi_{t,t+h}^{SPF}$ while permanent movements are tied directly to $F_t\tau_t$ and Θ_t and indirectly to $\zeta_{\eta,t}$, $\zeta_{v,t}$, λ_t , and Θ_t . The direct response of $\pi_{t,t+h}^{SPF}$ to Θ_t is produced by the observation equations (8.2). Drift in Θ_t also alters transition dynamics in the state equations (8.1), which generates movements in $F_t\tau_t$ and $F_t\varepsilon_t$, and hence, $\pi_{t,t+h}^{SPF}$.

3 Econometric Methods

We combine a RB-APF algorithm adapted from Lopes and Tsay (2011) with the PLE of Storvik (2002) to estimate the SSM (8.1) and (8.2); also see Carvalho, Johannes, Lopes, and Polson (2010), Creal (2012), and Herbst and Schorfheide (2016). The RB-APF and PLE produce filtered estimates of τ_t , ε_t , $F_t\tau_t$, $F_t\varepsilon_t$, $\zeta_{\eta,t}$, $\zeta_{v,t}$, λ_t , and θ_t . Lindsten, Bunch, Särkkä, Schön, and Godsill (2016) give instructions for a PS that generates smoothed estimates of these state variables.

3.1 Rao-Blackwellization of a Nonlinear State Space Model

Lopes and Tsay (2011, p. 173) and Creal (2012, section 2.5.7) outline APF algorithms that rely on the Rao-Blackwellization procedure of Chen and Liu (2000). The first step in Rao-Blackwellizing the SSM (8.1) and (8.2) gathers the nonlinear state variables in $\mathcal{V}_t = [\ln \zeta_{\eta,t}^2 \ \ln \zeta_{v,t}^2 \ \theta_t \ \lambda_t]'$. We generate updates of the nonlinear states by simulating the multivariate RW process

$$\mathcal{V}_{t+1} = \mathcal{V}_t + \mathbf{\Omega}_{\varepsilon}^{0.5} \varepsilon_{t+1}, \quad (9)$$

where $\mathcal{D}_{\varepsilon} = [\sigma_{\eta}^2 \ \sigma_v^2 \ \sigma_{\phi}^2 \ \sigma_{\kappa}^2]$ is the vector of non-zero elements of the diagonal covariance matrix $\mathbf{\Omega}_{\varepsilon}$ and $\varepsilon_{t+1} = [\xi_{\eta,t+1} \ \xi_{v,t+1} \ \phi_{t+1} \ \kappa_{t+1}]'$.⁹ The RB-APF uses the KF to create an analytic distribution of S_t using the SSM (8.1) and (8.2), given simulated values of \mathcal{V}_t . Analytic integration endows the RB-APF estimator of the linear state variables with greater numerical efficiency.

⁹The innovations vector $\varepsilon_{t+1} \sim \mathcal{N}(\mathbf{0}_{4 \times 1}, \mathbf{I}_4)$ conditions on $\theta_{t+1} \in (-1, 1)$ and $\lambda_{t+1} \in (0, 1)$.

3.2 Priors and Initial Conditions

We posit priors for the static volatility parameters and initial conditions to generate synthetic samples of linear and nonlinear states using the SSM (8.1) and (8.2) and multivariate RW (9). The static scale volatility parameters are collected in $\Psi = [\sigma_\eta^2 \ \sigma_\nu^2 \ \sigma_\phi^2 \ \sigma_\kappa^2 \ \sigma_{\zeta,\pi}^2 \ \sigma_{\zeta,1}^2 \ \dots \ \sigma_{\zeta,5}^2]'$. Priors on Ψ are grounded in restrictions of the joint DGP of the SI-prediction mechanism and SW-UC-SV-TVP-AR(1) model while remaining consistent with the PLE of Storvik (2002). The PLE requires priors for Ψ to have analytic posterior distributions. The posterior distributions serve as transition equations to update or “learn” about the joint distribution of \mathcal{S}_t and Ψ .

Table 1 lists our priors for the static volatility parameters found in Ψ . We endow these parameters with inverse gamma (\mathcal{IG}) priors. Columns labeled α_ℓ and β_ℓ denote the scale and shape parameters of the \mathcal{IG} priors of the elements of Ψ , the mean is $0.5\beta_\ell/(0.5\alpha_\ell - 1)$, and the two right most columns display the associated 2.5 and 97.5 percent quantiles, where $\ell = \eta, \nu, \phi, \kappa, \zeta_\pi, \zeta_h$, and $h = 1, \dots, 5$.

Table 1. Inverse Gamma Priors on the Static Coefficients

$$\Psi = \left[\sigma_\eta^2 \ \sigma_\nu^2 \ \sigma_\phi^2 \ \sigma_\kappa^2 \ \sigma_{\zeta,\pi}^2 \ \sigma_{\zeta,1}^2 \ \dots \ \sigma_{\zeta,5}^2 \right]'$$

		Quantiles			
Scale Volatility on Innovation to	α_ℓ	β_ℓ	Mean	2.5%	97.5%
Trend Inflation SV, $\ln \zeta_{\eta,t+1}$:	σ_η^2	3.0	0.04	0.04	[0.004, 0.186]
Gap Inflation SV, $\ln \zeta_{\nu,t+1}$:	σ_ν^2	3.0	0.04	0.04	[0.004, 0.186]
TVP-AR1 Coefficient, θ_{t+1} :	σ_ϕ^2	3.0	0.01	0.01	[0.001, 0.046]
SI Coefficient, λ_{t+1} :	σ_κ^2	3.0	0.01	0.01	[0.001, 0.046]
Measurement Error on π_t :	$\sigma_{\zeta,\pi}^2$	20.0	2.88	0.16	[0.084, 0.300]
Measurement Error on $\pi_{t,t+h}^{SPF}$:	$\sigma_{\zeta,h}^2$	20.0	2.88	0.16	[0.084, 0.300]

Priors on the static volatility coefficients are $\sigma_\ell^2 \sim \mathcal{IG}\left(\frac{\alpha_\ell}{2}, \frac{\beta_\ell}{2}\right)$, where α_ℓ and β_ℓ are scale and shape parameters, $\ell = \eta, \nu, \phi, \kappa, \zeta_\pi, \zeta_h$, and $h = 1, \dots, 5$

Two features are worth discussing about our priors on the scale volatility coefficients of \mathcal{D}_ε . First, we give σ_η^2 and σ_v^2 prior means equal to 0.04. These prior means are larger than the prior mean of 0.01 placed on σ_ϕ^2 and σ_κ^2 . Second, our priors on σ_η^2 , σ_v^2 , σ_ϕ^2 , and σ_κ^2 deliver 2.5 and 97.5 percent quantiles that exhibit greater variation in innovations to $\ln \zeta_{\eta,t+1}^2$ and $\ln \zeta_{v,t+1}^2$ compared with variation in innovations to θ_{t+1} and λ_{t+1} . Nonetheless, the 2.5 and 97.5 percent quantiles of $\sigma_{\zeta,\pi}^2$, $\sigma_{\zeta,1}^2, \dots, \sigma_{\zeta,5}^2$ reveal our belief that volatility in the measurement errors of π_t and $\pi_{t,t+h}^{SPF}$ dominate shock volatility in the joint DGP of the SI-prediction mechanism and SW-UC-SV-TVP-AR(1) model.

Priors on initial conditions of the linear state variables appear in the left two columns of table 2. We draw τ_0 and $F_0\tau_0$ from normal priors with a mean of two percent, which is about the mean of GNP deflator inflation on a 1958Q1 to 1967Q4 training sample. A variance of 100^2 indicates a flat prior over a wide range of values for τ_0 and $F_0\tau_0$. The joint prior of ε_0 and $F_0\varepsilon_0$ is drawn from $\mathcal{N}(\mathbf{0}_{2 \times 1}, \mathbf{\Sigma}_0)$, which equates the prior means to zero (*i.e.*, unconditional means). Prior variances are produced by the ergodic bivariate normal distribution of particle draws of $\zeta_{v,0}$, θ_0 , and λ_0 ; see the notes to table 2. We also restrict priors on τ_0 , ε_0 , $F_0\tau_0$, and $F_0\varepsilon_0$ by splitting the training sample variance of the first difference of GNP deflator inflation between trend (one-third) and gap (two-thirds) shocks.

The last two columns of table 2 lists priors on initial conditions of the nonlinear state variables. We endow priors of $\ln \zeta_{v,0}^2$ and $\ln \zeta_{\eta,0}^2$ with normal distributions. Prior means are calibrated to pre-1968 inflation data similar to Stock and Watson (2007). Uncertainty about $\ln \zeta_{v,0}^2$ and $\ln \zeta_{\eta,0}^2$ is reflected in prior variances of ten. Table 2 shows that θ_0 is drawn from a standard normal, subject to truncation at $(-1, 1)$, and another truncated normal bounds $\lambda_0 \in (0, 1)$ with (untruncated) mean of 0.5 and a unit variance. These priors are in essence uninformative about values inside the bounds.

Table 2. Priors on Initial Conditions of the Linear and Nonlinear States

$$\mathcal{S}_0 = \left[\tau_0 \quad \varepsilon_0 \quad F_0 \tau_0 \quad F_0 \varepsilon_0 \right]' \quad \text{and} \quad \mathcal{V}_0 = \left[\ln \zeta_{\eta,0}^2 \quad \ln \zeta_{v,0}^2 \quad \theta_0 \quad \lambda_0 \right]'$$

Initial State	Prior Distribution	Initial State	Prior Distribution
τ_0 :	$\mathcal{N}(2.0, 100.0^2)$	$\ln \zeta_{\eta,0}^2$:	$\ln \mathcal{N}(\ln 0.2 - 5.0, 10.0)$
$F_0 \tau_0$:	$\mathcal{N}(2.0, 100.0^2)$	$\ln \zeta_{v,0}^2$:	$\ln \mathcal{N}(\ln 0.4 - 5.0, 10.0)$
ε_0 :	$\mathcal{N}(0.0, \sigma_{\varepsilon_0}^2)$	θ_0 :	$\mathcal{TN}(0.0, 1.0, -1.0, 1.0)$
$F_0 \varepsilon_0$:	$\mathcal{N}(0.0, \sigma_{F_0 \varepsilon_0}^2)$	λ_0 :	$\mathcal{TN}(0.5, 1.0, 0.0, 1.0)$

The truncated normal distribution is denoted \mathcal{TN} , where the first two entries are the mean and variance of the prior and the last two entries restrict the range of the prior. The priors on ε_0 and $F_0 \varepsilon_0$ are drawn jointly from $\mathcal{N}(\mathbf{0}_{2 \times 1}, \boldsymbol{\Sigma}_0^{(i)})$, where $\sigma_{\varepsilon_0}^2$ and $\sigma_{F_0 \varepsilon_0}^2$ are diagonal elements of

$$\boldsymbol{\Sigma}_0^{(i)} = \sum_{j=0}^{\infty} \begin{bmatrix} \theta_0^{(i)} & 0 \\ (1 - \lambda_0^{(i)}) \theta_0^{(i)} & \lambda_0^{(i)} \theta_0^{(i)} \end{bmatrix}^j \begin{bmatrix} \zeta_{v,0}^{2,(i)} & \lambda_0^{(i)} \zeta_{v,0}^{2,(i)} \\ \lambda_0^{(i)} \zeta_{v,0}^{2,(i)} & \lambda_0^{2,(i)} \zeta_{v,0}^{2,(i)} \end{bmatrix} \begin{bmatrix} \theta_0^{(i)} & (1 - \lambda_0^{(i)}) \theta_0^{(i)} \\ 0 & \lambda_0^{(i)} \theta_0^{(i)} \end{bmatrix}^j,$$

and $\lambda_0^{(i)}$, $\theta_0^{(i)}$, and $\zeta_{v,0}^{2,(i)}$ are the i th particle draws from priors on the associated initial conditions. If $\theta_0^{(i)} = 0$, the formula computing $\boldsymbol{\Sigma}_0^{(i)}$ remains valid.

3.3 The Auxiliary Particle Filter

Section 3.1 applies the RB process to the SSM (8.1) and (8.2). This process increases the numerical efficiency of the estimator of the linear states, \mathcal{S}_t , by shrinking Monte Carlo error. Another method to improve the efficiency of this estimator is the APF of Pitt and Shephard (1999, 2001). In this section, we sketch a RB-APF to estimate the linear and nonlinear states that begins from algorithm 2 of Lopes and Tsay (2011, p.173); also see Creal (2012, section 2.5.7).¹⁰ The online appendix provides a complete exposition of our implementation of the RB-APF.

¹⁰We sketch a RB-APF separate from the PLE for clarity, but recognize the RB-APF is integral to the PLE.

A RB-APF obtains estimates of the likelihood by running the prediction step of the KF on the SSM (8.1) and (8.2) particle by particle. At date t , the KF predictive step produces the log likelihood, $\ell_t^{(i)}$, and particle weights, $\widehat{\omega}_t^{(i)} = \exp\{\ell_t^{(i)}\} / \sum_i^M \exp\{\ell_t^{(i)}\}$, $i = 1, \dots, M$. Stratified resampling of $\{\widehat{\omega}_t^{(i)}\}_{i=1}^M$ yields indexes that are used to regroup $\mathcal{S}_{t-1|t-1}^{(i)}$, its mean square error (MSE), $\boldsymbol{\Sigma}_{t-1|t-1}^{(i)}$, and $\mathcal{V}_t^{(i)}$; see steps 3(a) and 3(b) of section A3.1 of the online appendix and Hol, Schön, and Gustafsson (2006). This step aims to prevent a particle from receiving all the probability mass as M becomes large. The ensemble of weights $\{\widehat{\omega}_t^{(i)}\}_{i=1}^M$ are also resampled generating $\{\widetilde{\omega}_t^{(i)}\}_{i=1}^M$; see step 3(d) of section A3.1 of the online appendix. The resampled particles $\mathcal{S}_{t-1|t-1}^{(i)}$, $\boldsymbol{\Sigma}_{t-1|t-1}^{(i)}$, and $\mathcal{V}_t^{(i)}$ are employed in the entire KF to update $\{\mathcal{S}_{t|t}^{(i)}, \boldsymbol{\Sigma}_{t|t}^{(i)}, \ell_t^{(i)}\}_{i=1}^M$ and produce new weights $\omega_t^{(i)} = \exp\{\ell_t^{(i)}\} / \sum_i^M \exp\{\ell_t^{(i)}\}$, $i = 1, \dots, M$; see step 3(e) of section A3.1 of the online appendix. By simulating the multivariate RW (9), the nonlinear states are updated to $\mathcal{V}_{t+1}^{(i)}$ across the M particles. Estimates of $\mathcal{S}_{t|t}$, $\boldsymbol{\Sigma}_{t|t}$, and $\mathcal{V}_{t|t}$ rely on the weights $\omega_t^{(i)} = \omega_t^{(i)} / \widetilde{\omega}_t^{(i)}$; see step 4 of section A3.1 of the online appendix.

As already noted, a useful product of the RB-APF is the likelihood of the conditionally linear SSM (8.1) and (8.2). Since the M particles have been reweighted at every step using information contained in the likelihood of the KF, the estimate of the date t data density is

$$\mathcal{P}(y_t | y_{1:t-1}; \Psi) \propto \frac{1}{M} \sum_{i=1}^M \exp\{\ell_t^{(i)}\}, \quad t = 1, \dots, T. \quad (10)$$

Sum the data density (10) over the $t = 1, \dots, T$ observations to compute the log likelihood of the SSM

$$\mathcal{L}(\Psi | y_{1:T}) = \sum_{t=1}^T \ln \left(\mathcal{P}(y_t | y_{1:t-1}; \Psi) \right). \quad (11)$$

Section 4 reports estimates of the joint DGP of the SI-prediction mechanism and SW-UC-SV-TVP-AR(1) model. Its estimated log likelihood is compared with the log likelihood of a joint DGP estimated conditional on setting $\theta_t = 0$ or estimating a constant SI parameter, $\lambda_t = \lambda$. Thus, we use log likelihood (11) to evaluate competing joint DGPs, but only after marginalizing Ψ . The next section discusses the PLE used to estimate Ψ .

3.4 The Particle Learning Estimator

We estimate the joint posterior distribution of \mathcal{S}_t , $\mathbf{\Sigma}_t$, \mathcal{V}_t , and Ψ by embedding the RB-APF in the PLE of Storvik (2002), given priors on the joint DGP of the SI-prediction mechanism and SW-UC-SV-TVP-AR(1) model.¹¹ The PLE rests on two insights. First, choosing conjugate priors for Ψ yields an analytic solution of its posterior distributions. The posterior distribution is recovered conditional on the states and sample data. The idea is to draw Ψ from particle streams of a vector of sufficient statistics, $\Gamma_t^{(i)}$ that depend on $\mathcal{V}_t^{(i)}$, given $\mathcal{Y}_{1:t}$. Since the sufficient statistics are grounded in the \mathcal{JG} priors of Ψ , the mapping to the analytic posterior distributions is a system of transition equations that simulate M particles to learn about or update from $\Gamma_{t-1}^{(i)}$ to $\Gamma_t^{(i)}$. The transition equations are appended to the process that draws $\mathcal{V}_t^{(i)}$ to sample $\Psi^{(i)} \sim \mathcal{P}(\Psi | \Gamma_t^{(i)})$, which in essence equates $\mathcal{P}(\Psi | \mathcal{Y}_{1:t}, \mathcal{V}_t^{(i)})$ to $\mathcal{P}(\Psi | \Gamma_t^{(i)})$. We denote the system of transition equations $\Gamma_t^{(i)} = \mathcal{F}(\Gamma_{t-1}^{(i)}, \mathcal{Y}_{1:t}, \mathcal{V}_t^{(i)}, \mathcal{V}_{t-1}^{(i)})$, $i = 1, \dots, M$.

Second, the PLE marginalizes Ψ out of the posterior of the states produced by the RB-APF. The idea is to update $\Gamma_t^{(i)}$ at the same time the RB-APF generates $\mathcal{S}_t^{(i)}$, $\mathbf{\Sigma}_t^{(i)}$, and $\mathcal{V}_t^{(i)}$. Thus, Ψ is estimated by the PLE jointly with $\mathcal{S}_{t|t}$, $\mathbf{\Sigma}_{t|t}$, and $\mathcal{V}_{t|t}$.

As noted, we place \mathcal{JG} priors on Ψ to expedite Storvik's PLE. The priors, which are reviewed in section 3.2 and table 1, are $\sigma_\ell^2 \sim \mathcal{JG}\left(\frac{\alpha_\ell}{2}, \frac{\beta_\ell}{2}\right)$, where ℓ indexes the elements of Ψ . The \mathcal{JG} priors are useful because the associated posterior distributions are solved analytically. For example, the posterior distribution of the static volatility coefficient of the RW of θ_{t+1} is $\sigma_\phi^{2(i)} \sim \mathcal{JG}\left(\frac{\alpha_t}{2}, \frac{\beta_{\phi,t}^{(i)}}{2}\right)$, where $\alpha_t = \alpha_{t-1} + t - 1$ and $\beta_{\phi,t}^{(i)} = \sum_{\ell=1}^t [\theta_\ell^{(i)} - \theta_{\ell-1}^{(i)}]^2$. The process generating $\beta_{\phi,t}^{(i)}$ suggests conditioning the posterior $\sigma_\phi^{2(i)} | \mathcal{V}_t^{(i)}, \mathcal{V}_{t-1}^{(i)} \sim \mathcal{JG}\left(\frac{\alpha_t}{2}, \frac{\beta_{\phi,t}^{(i)}}{2}\right)$, where the shape parameter $\beta_{\phi,t}^{(i)}$ is a sufficient statistic for σ_ϕ^2 .¹² We extend the idea of identifying $\beta_{\phi,t}^{(i)}$ as sufficient statistics to the entire collection of static volatility parameters in Ψ .

¹¹Another method to estimate Ψ is to wrap a Metropolis-Hasting Markov chain Monte Carlo (MCMC) simulator around a PF. Andrieu, Doucet, and Holenstein (2010) prove the distribution of a MCMC simulator is independent of the error created by a particle in a SMC algorithm. Hence, a PF gives an unbiased estimator of the likelihood (11).

¹²The shape parameter is the numerator of the standard deviation of a random variable distributed \mathcal{JG} .

The online appendix gives procedures to simulate and update $\beta_{\eta,t}^{(i)}$, $\beta_{\nu,t}^{(i)}$, $\beta_{\phi,t}^{(i)}$, and $\beta_{\kappa,t}^{(i)}$ in steps 2 and 3.(a) of the RB-APF algorithm. The algorithm samples $\sigma_{\eta,t}^{2(i)}$, $\sigma_{\nu,t}^{2(i)}$, $\sigma_{\phi,t}^{2(i)}$, and $\sigma_{\kappa,t}^{2(i)}$ from particle streams of sufficient statistics. The law of motion of sufficient statistic $\beta_{\ell,t}^{(i)}$ matches the transition equation $\beta_{\ell,t}^{(i)} = \mathcal{F}_{\ell,t}(\beta_{\ell,t-1}^{(i)}, \mathcal{Y}_{1:t}, \mathcal{V}_t^{(i)}, \mathcal{V}_{t-1}^{(i)})$, for $\ell = \eta, \nu, \phi$, and κ .

This leaves us to describe the routines that sample the measurement error scale volatility parameters, $\sigma_{\zeta,\pi}^2$ and $\sigma_{\zeta,h}^2$, $h = 1, \dots, 5$. Since these variances lack laws of motion that can be employed to build transition equations, the relevant shape parameters are updated on information obtained from KF operations of the RB-APF. For example, we sample $\sigma_{\zeta,h}^{2(i)} \mid \mathcal{Y}_{1:t} \sim \mathcal{JG}\left(\frac{\alpha_t}{2}, \frac{\beta_{\zeta,h,t}^{(i)}}{2}\right)$, where updates of $\beta_{\zeta,h,t}^{(i)}$ are calculated using information from step 3.(b) of the RB-APF; see the online appendix. Thus, updates of the shape parameters of the posterior distributions of $\sigma_{\zeta,\pi}^2$ and $\sigma_{\zeta,h}^2$, which are the sufficient statistics $\beta_{\zeta,\pi,t}^{(i)}$ and $\beta_{\zeta,h,t}^{(i)}$, are driven by the KF prediction error of \mathcal{Y}_t weighted by the “gain” of these innovations.

We summarize the PLE and the way it interacts with RB-APF with the following algorithm.

1. Before initializing the RB-APF at date 0, draw $\Psi^{(i)} = \mathcal{P}(\Gamma_0^{(i)})$.
2. Next, carry out steps 1, 2, and 3.(a)–3.(c) of the RB-APF algorithm (that appear in the online appendix) to obtain the KF predictive likelihood $\ell_t^{(i)} \propto \mathcal{P}(\mathcal{Y}_t \mid \mathcal{S}_{t-1|t-1}^{(i)}, \Sigma_{t-1|t-1}^{(i)}, \mathcal{V}_t^{(i)}, \Psi^{(i)})$ and calculate the particle weights, $\widehat{\omega}_t^{(i)}$.
3. Update the particles $\Psi^{(i)}$, $i = 1, \dots, M$, using the system of transition equations $\Gamma_t^{(i)} = \mathcal{F}_{\Gamma_t}(\Gamma_{t-1}^{(i)}, \mathcal{Y}_{1:t}, \mathcal{V}_t^{(i)}, \mathcal{V}_{t-1}^{(i)})$, which guide the evolution of this vector of sufficient statistics.
4. Engage $\{\widehat{\omega}_t^{(i)}\}_{i=1}^M$ to resample $\{\Gamma_t^{(i)}\}_{i=1}^M$ and perform steps 3.(d)–3.(f), 4, and 5 of the RB-APF (that are listed in the online appendix).
5. Resample $\Delta_{\ell,t}^{(i)}$, which are changes to $\beta_{\ell,t}^{(i)}$, $\ell = \eta, \nu, \phi$, and κ , as described in step 3.(d) of the RB-APF discussed in the online appendix, but “innovations” to $\beta_{\zeta,\pi,t-1}^{(i)}$ and $\beta_{\zeta,h,t-1}^{(i)}$, $\Delta_{\zeta,\pi,t}^{(i)}$ and $\Delta_{\zeta,h,t}^{(i)}$, are not resampled.
6. Repeat steps 1 to 5 of the PLE starting at date $t = 1$ and stopping at date T .

7. Full sample estimates of the static volatility parameters are computed according to $\hat{\Psi} = \sum_{i=1}^M \omega_T^{(i)} \Psi_T^{(i)} = \sum_{i=1}^M \omega_T^{(i)} \mathcal{P}(\Gamma_T^{(i)})$.

By repeating step 5 at dates $t = 1, \dots, T-1$, the PLE produces information about the content of \mathcal{Y}_t for the way the RB-APF “learns” about Ψ .

3.5 A Rao-Blackwellized Particle Smoother

Lindsten, Bunch, Särkkä, Schön, and Godsill (2016) develop an algorithm to compute smoothed estimates of \mathcal{S}_t and \mathcal{V}_t , given $\mathcal{Y}_{1:T}$ and $\hat{\Psi}$. The algorithm is a forward filter-backward smoother (FFBS) for SSMs amenable to Rao-Blackwellization. The forward filter is the RB-APF described in section 3.3 and online appendix. The FFBS applies Rao-Blackwellization methods moving from date T to date 1 to generate smoothed estimates of \mathcal{V}_t conditional on forward filtered particles. Forward filtering operations are conducted using the SSM (8.1) and (8.2) to produce smoothed estimates of \mathcal{S}_t , given smoothed estimates of \mathcal{V}_t .¹³ Lindsten, Bunch, Särkkä, Schön, and Godsill (LBSSG) refer to the entire process as a forward-backward-forward smoother.

The RB-PS operates only on the nonlinear states of the joint DGP of the SI-prediction mechanism and SW-UC-SV-TVP-AR(1) model. The problem is, when moving backwards from date t to date $t-1$, smoothing \mathcal{V}_t can cause its Markov structure to be lost. A reason is that marginalizing the linear states produces a likelihood that depends on $\mathcal{V}_{1:t}$ rather than \mathcal{V}_t .

LBSSG solve this sampling problem by decomposing the target density $\mathcal{P}(\mathcal{V}_{1:T} | \mathcal{Y}_{1:T}; \hat{\Psi})$ into $\mathcal{P}(\mathcal{V}_{1:t} | \mathcal{V}_{1:t+1}, \mathcal{Y}_{1:T}; \hat{\Psi}) \mathcal{P}(\mathcal{V}_{t+1:T} | \mathcal{Y}_{1:T}; \hat{\Psi})$. Drawing from $\mathcal{P}(\mathcal{V}_{t+1:T} | \mathcal{Y}_{1:T}; \hat{\Psi})$ yields an incomplete path of the approximate smoothed nonlinear states from date $t+1$ to date T , which is denoted $\tilde{\mathcal{V}}_{t+1:T}$. Since these draws are initialized at date T by sampling from the date T filtered nonlinear states, $\{\mathcal{V}_T^{(i)}\}_{(i)=1}^M$, backward extension to $\tilde{\mathcal{V}}_{T-1:T}$ is drawn probabilistically from the cloud $\{\mathcal{V}_{1:T-1}^{(i)}\}_{(i)=1}^M$. The Rao-Blackwellized particle smoother is repeated for $t = T-2, \dots, 1$.

¹³Alternative PS are found in Lopes and Tsay (2011) and Carvalho, Johannes, Lopes, and Polson (2010). These approaches to smoothing, which build on the PS of Godsill, Doucet, and West (2004), are applicable to APFs, but not to the RB-APF we employ.

The aforesaid factorization of $\mathcal{P}(\mathcal{V}_{1:T} | \mathcal{Y}_{1:T}; \hat{\Psi})$ is also useful because there is information in $\mathcal{P}(\mathcal{V}_{1:t} | \mathcal{V}_{1:t+1}, \mathcal{Y}_{1:T}; \hat{\Psi})$ about the probabilities (*i.e.*, normalized weights) needed to draw smoothed nonlinear states. Gaining access to this information is difficult because the conditional density of $\mathcal{V}_{1:t}$ is not easy to evaluate.¹⁴ LBSSG's propose simulation methods to perform the backward filtering implicit in $\mathcal{P}(\mathcal{V}_{1:t} | \mathcal{V}_{1:t+1}, \mathcal{Y}_{1:T}; \hat{\Psi})$. This density can be decomposed into

$$\mathcal{P}(\mathcal{V}_{1:t} | \mathcal{V}_{1:t+1}, \mathcal{Y}_{1:T}; \hat{\Psi}) \propto \mathcal{P}(\mathcal{Y}_{t+1:T}, \mathcal{V}_{t+1:T} | \mathcal{V}_{1:t}, \mathcal{Y}_{1:t}; \hat{\Psi}) \mathcal{P}(\mathcal{V}_{1:t} | \mathcal{Y}_{1:t}; \hat{\Psi}),$$

where the object of interest is the predictive density $\mathcal{P}(\mathcal{Y}_{t+1:T}, \mathcal{V}_{t+1:T} | \mathcal{V}_{1:t}, \mathcal{Y}_{1:t}; \hat{\Psi})$. LBSSG show this density equals $\int \mathcal{P}(\mathcal{Y}_{t+1:T}, \mathcal{V}_{t+1:T} | \mathcal{S}_t, \mathcal{V}_t; \hat{\Psi}) \mathcal{P}(\mathcal{S}_t | \mathcal{Y}_{1:t}, \mathcal{V}_{1:t}; \hat{\Psi}) d\mathcal{S}_t$. Hence, run the KF forward to obtain estimates of \mathcal{S}_t and $\boldsymbol{\Sigma}_t$ by drawing from $\mathcal{P}(\mathcal{S}_t | \mathcal{Y}_{1:t}, \mathcal{V}_{1:t}; \hat{\Psi})$. The mean and MSE of \mathcal{S}_t are employed in simulations to generate sufficient statistics that approximate the density of the SSM (8.1) and (8.2), which when normalized are the probabilities of drawing a path of $\tilde{\mathcal{V}}_{1:T}$. The upshot is, although \mathcal{S}_t does not enter $\mathcal{P}(\mathcal{V}_{1:t} | \mathcal{V}_{1:t+1}, \mathcal{Y}_{1:T}; \hat{\Psi})$, the conditionally linear states are relevant for estimating the probability of sampling $\tilde{\mathcal{V}}_{1:T}$. In a final step that is conditional on the path of $\tilde{\mathcal{V}}_{1:t}$, the linear states are smoothed by iterating the KF forward.

Our implementation of LBSSG's RB-PS is described in the next algorithm.

1. Retrieve a stored ensemble of M particles, $\left\{ \left\{ \mathcal{S}_{t|t}^{(i)}, \boldsymbol{\Sigma}_{t|t}^{(i)}, \mathcal{V}_t^{(i)}, \omega_t^{(i)} \right\}_{t=1}^T \right\}_{i=1}^M$, created by running the RB-APF on the SSM (8.1) and (8.2), given the PLE of $\Psi, \hat{\Psi}$.
2. Initialize the PS at date T
 - (a) by drawing $\tilde{\mathcal{V}}_T^{(i)}$ for each i from the filtered particle draws $\left\{ \mathcal{V}_T^{(i)} \right\}_{i=1}^M$ that have been resampled using the weights $\omega_T^{(i)}$, and
 - (b) compute $\tilde{\mathbf{O}}_T^{(i)} = (\mathbf{c}_T^{(i)})' \hat{\boldsymbol{\Omega}}_u^{-1} \mathbf{c}_T^{(i)}$, $\tilde{\boldsymbol{\Sigma}}_T^{(i)} = (\mathbf{c}_T^{(i)})' \hat{\boldsymbol{\Omega}}_u^{-1} \mathcal{Y}_T$, and $\tilde{\mathbf{U}}_T^{(i)} = \mathbf{I} + (\mathbf{B}_T^{(i)})' \tilde{\mathbf{O}}_T^{(i)} \mathbf{B}_T^{(i)}$ across M resampled particles, which are used to build sufficient statistics of \mathcal{S}_T .

¹⁴The KF creates an exact predictive density (up to a normalizing constant). However, computing the density involves iterating the filter forward from dates $t = 1, \dots, T-1$ to date T across M particle streams. These calculations are computationally costly, which motivate LBSSG to approximate the predictive density with simulated sufficient statistics.

3. For each particle $i = 1, \dots, M$, iterate from date $t = T-1$ to date $t = 1$ to calculate the unnormalized weights, $\mathcal{W}_t^{(i)} = \omega_t^{(i)} \delta_t^{(i)} \left| \mathbf{U}_t^{(i)} \right|^{-0.5} \exp \left\{ \frac{1}{2} \mathfrak{g}_t^{(i)} \right\}$, which generate smoothed normalized weights, $\tilde{\omega}_{t|T}^{(i)} = \mathcal{W}_t^{(i)} / \sum_{m=1}^M \mathcal{W}_t^{(m)}$, where

(a) $\Pr(j = i) = \tilde{\omega}_{t|T}^{(i)}$ counts the number of instances $\mu > \sum_{j=1}^m \tilde{\omega}_{t|T}^{(j)}$, $\mu \sim \mathcal{U}(0, 1)$, and $m = 1, \dots, M$,

(b) $\delta_t^{(i)} \sim \mathcal{P}(\tilde{\mathcal{V}}_{t+1} | \mathcal{V}_t^{(i)}; \hat{\Psi})$, which is implemented by $\delta_t^{(i)} = \exp \{ \mathfrak{v}_{\eta,t}^{(i)} + \mathfrak{v}_{\nu,t}^{(i)} \} \times \mathfrak{v}_{\theta,t}^{(i)} \times \mathfrak{v}_{\lambda,t}^{(i)}$,

$$\mathfrak{v}_{\ell,t}^{(i)} = -\frac{1}{2} \left[\frac{\ln \tilde{\xi}_{\ell,t+1}^{2(i)} - \ln \xi_{\ell,t}^{2(i)}}{\hat{\sigma}_\ell} \right]^2, \quad \ell = \eta, \nu,$$

$$\mathfrak{v}_{\theta,t}^{(i)} = \frac{\mathcal{N}(\tilde{\theta}_{t+1}, \theta_t^{(i)}, \hat{\sigma}_\phi)}{\mathcal{N}_{\Delta,\theta,t}}, \quad \mathfrak{v}_{\lambda,t}^{(i)} = \frac{\mathcal{N}(\tilde{\lambda}_{t+1}, \lambda_t^{(i)}, \hat{\sigma}_\phi)}{\mathcal{N}_{\Delta,\lambda,t}},$$

$$\mathcal{N}_{\Delta,\theta,t} = \Phi\left(\frac{1 - \theta_t^{(i)}}{\hat{\sigma}_\phi}\right) - \Phi\left(\frac{-1 - \theta_t^{(i)}}{\hat{\sigma}_\phi}\right), \quad \mathcal{N}_{\Delta,\lambda,t} = \Phi\left(\frac{1 - \lambda_t^{(i)}}{\hat{\sigma}_\kappa}\right) - \Phi\left(\frac{-\lambda_t^{(i)}}{\hat{\sigma}_\kappa}\right),$$

$\Phi(\cdot)$ is the CDF of the normal distribution,

(c) $\mathbf{U}_t^{(i)} = \mathbf{I} + (\boldsymbol{\Sigma}_t^{(i)})' \tilde{\mathbf{O}}_t^{(i)} \boldsymbol{\Sigma}_t^{(i)}$, which depends on the backwards transition equations¹⁵

$$\begin{aligned} \bar{\mathbf{O}}_t^{(i)} &= (\mathcal{A}_{t+1}^{(i)})' \left[\mathbf{I} - \tilde{\mathbf{O}}_{t+1}^{(i)} \boldsymbol{\Sigma}_{t+1}^{(i)} (\mathbf{U}_{t+1}^{(i)})^{-1} (\boldsymbol{\Sigma}_{t+1}^{(i)})' \right] \tilde{\mathbf{O}}_{t+1}^{(i)} \mathcal{A}_{t+1}^{(i)}, \\ \tilde{\mathbf{O}}_t^{(i)} &= \bar{\mathbf{O}}_t^{(i)} + (\mathbf{e}_t^{(i)})' \hat{\boldsymbol{\Omega}}_u^{-1} \mathbf{e}_t^{(i)}, \end{aligned}$$

and

$$(d) \mathfrak{g}_t^{(i)} = (\tilde{\mathcal{S}}_t^{(i)})' \left(\tilde{\mathbf{O}}_t^{(i)} \right)^{\frac{1}{2}} \left(\tilde{\mathbf{O}}_t^{(i)} \right)^{\frac{1}{2}'} \tilde{\mathcal{S}}_t^{(i)} - 2 (\tilde{\mathcal{S}}_t^{(i)})' \mathfrak{S}_t^{(i)} - \left[\boldsymbol{\Sigma}_t^{(i)} \left(\tilde{\mathcal{S}}_t^{(i)} - \tilde{\mathbf{O}}_t^{(i)} \mathfrak{S}_t^{(i)} \right) \left(\tilde{\mathbf{O}}_t^{(i)} \right)^{-1} \right]^2,$$

where the backwards laws of motion consist of $\tilde{\mathcal{S}}_t^{(i)} = \bar{\mathcal{S}}_t^{(i)} + (\mathbf{e}_t^{(i)})' \hat{\boldsymbol{\Omega}}_u^{-1} \mathbf{y}_t$ and

$$\bar{\mathcal{S}}_t^{(i)} = (\mathcal{A}_{t+1}^{(i)})' \left[\mathbf{I} - \tilde{\mathbf{O}}_{t+1}^{(i)} \boldsymbol{\Sigma}_{t+1}^{(i)} (\mathbf{U}_{t+1}^{(i)})^{-1} (\boldsymbol{\Sigma}_{t+1}^{(i)})' \right] \tilde{\mathcal{S}}_{t+1}^{(i)}.$$

4. Given the draw j in step 3.(a), add \mathcal{V}_t^j to $\tilde{\mathcal{V}}_{t+1:T}$ to produce the approximate (partially) smoothed trajectory $\tilde{\mathcal{V}}_{t:T} = \left\{ \mathcal{V}_t^j, \tilde{\mathcal{V}}_{t+1:T} \right\}$.

¹⁵LBSSG propose a square root KF to ensure numerical stability of the backward filtering operations.

5. Subsequent to iterating backwards from date T to date 1 for each of the M particle paths, calculate the distribution of $\tilde{\mathcal{S}}_t$ by running the KF on the SSM (8.1) and (8.2) to produce

$$\mathcal{P}(\mathcal{S}_{1:T} | \mathcal{Y}_{1:T}; \hat{\Psi}) = \int_{\tilde{\mathcal{V}}_{1:T}} \mathcal{P}(\mathcal{S}_{1:T} | \tilde{\mathcal{V}}_{1:T}^{(i)}, \mathcal{Y}_{1:T}; \hat{\Psi}) dF(\tilde{\mathcal{V}}_{1:T}) \approx \sum_{i=1}^M \mathcal{P}(\mathcal{S}_{1:T} | \tilde{\mathcal{V}}_{1:T}^{(i)}, \mathcal{Y}_{1:T}; \hat{\Psi}), \quad (12)$$

which is conditional on the backward filtered smoothed approximate paths of $\tilde{\mathcal{V}}_{1:T}$.

The density $\mathcal{P}(\mathcal{S}_{1:T} | \tilde{\mathcal{V}}_{1:T}^{(i)}, \mathcal{Y}_{1:T}; \hat{\Psi})$ is the multivariate normal distribution that results from running the Kalman smoother (KS) on the SSM (8.1) and (8.2), where the approximation on the right side of (12) conditions on the M trajectories of $\tilde{\mathcal{V}}_{1:T}^{(i)}$. We use the disturbance smoothing algorithm of Durbin and Koopman (2002) to draw $\tilde{\mathcal{S}}_{1:T}$ from the distribution created by the KS, given a $\tilde{\mathcal{V}}_t^{(i)}$ generated by the PS.

4 The Data and Estimates

We present estimates of the joint DGP of the SI-prediction mechanism and SW-UC-SV-TVP-AR(1) model in this section. These estimates are compared with ones gleaned from joint DGPs that lack inflation gap persistence, $\theta_t = 0$ or drift in SI updating $\lambda_t = \lambda$.¹⁶ The goal is to evaluate the impact of inflation gap persistence or SI on the dynamics of π_t and $\pi_{t,t+h}^{SPF}$, $h = 1, \dots, 5$. The joint DGPs are estimated using a RB-APF, PLE, and PS that engage $M = 100,000$ particles. These estimates are used to study (i) comovement of τ_t and $F_t \tau_t$ with π_t , and $\pi_{t,t+h}^{SPF}$, (ii) fluctuations in ε_t and $F_t \varepsilon_t$, (iii) the history of $\zeta_{\eta,t}$ and $\zeta_{\nu,t}$ since the start of the sample, (iv) movements in θ_t and λ_t over the business cycle, and (v) the contributions of \mathcal{Y}_t , π_t , and $\pi_{t,t+h}^{SPF}$ to variation in τ_t and $F_t \tau_t$.

4.1 The Data

Our estimates rest on a sample of real-time realized inflation, π_t , and h -step ahead average SPF inflation prediction, $\pi_{t,t+h}^{SPF}$. We obtain the data from the Real-Time Data Set for Macroe-

¹⁶When $\theta_t = 0$ ($\lambda_t = \lambda$), σ_ϕ^2 (σ_κ^2) is deleted from Ψ . Fixing the frequency of SI updating also adds λ to Ψ with the prior $\lambda \sim \text{Beta}(1, 1)$.

conomists (RTDSM), which is compiled by the Federal Reserve Bank (FRB) of Philadelphia.¹⁷ The data consist of observations from 1968Q4 through 2017Q2 for real-time realized inflation and average SPF inflation predictions.

Realized inflation is the RTDSM's quarterly real-time vintages of the GNP and GDP deflator.¹⁸ 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 $t-1$. We employ these vintages to compute the quarterly difference in the log levels of real-time observations on the GNP or GDP deflator, P_t . The quarterly price level data are transformed into inflation measured at an annualized rate using $\pi_t = 400[\ln P_t - \ln P_{t-1}]$.


Average SPF inflation predictions include a nowcast of the GNP or GDP deflator's level and forecasts of these price levels 1-, 2-, 3-, and 4-quarters ahead. These surveys are collected at quarter t without full knowledge of π_t . We comply with this timing protocol by assuming the average nowcast, 1-quarter, \dots , and 4-quarter ahead predictions, which are denoted $\pi_{t,t+1}^{SPF}$, $\pi_{t,t+2}^{SPF}$, \dots , and $\pi_{t,t+5}^{SPF}$, are conditional on data available at the end of quarter $t-1$. These inflation predictions are the annualized log differences of the average SPF prediction of the deflator's level and one lag of the real-time realized price level supplied by the RTDSM.

Figure 1 plots π_t and four different average SPF inflation predictions. Plots of π_t and the average SPF inflation nowcast, $\pi_{t,t+1}^{SPF}$, appear in figure 1(a). Realized inflation is also found in figure 1(b), but the 1-quarter ahead average SPF inflation prediction, $\pi_{t,t+2}^{SPF}$, replaces $\pi_{t,t+1}^{SPF}$. Figure 1(c) displays π_t and the 3-quarter ahead average SPF inflation prediction, $\pi_{t,t+4}^{SPF}$, and figure 1(d) has π_t and the 4-quarter ahead average SPF inflation prediction, $\pi_{t,t+5}^{SPF}$. The panels depict π_t with a dot-dash (red) line and average SPF inflation predictions with a solid (blue) line. Vertical gray shaded bars denote NBER recession dates.

¹⁷The data are available at <http://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/>.

¹⁸The SPF measured the output price level with the implicit GNP deflator before 1992Q1. From 1992Q1 to 1996Q4, the implicit GDP deflator played this role. It was replaced by the chain weighted GDP deflator from 1997Q1 to the end of the sample.

FIGURE 1: REALIZED INFLATION AND SPF INFLATION PREDICTIONS, 1968Q4 TO 2017Q2



Note: The four plots contain vertical gray bands that denote NBER dated recessions.

The plots reveal several features of π_t and the average SPF inflation predictions. First, average SPF inflation predictions exhibit less variation than π_t throughout the sample. Next, as h increases, average SPF inflation predictions become smoother and are centered on π_t . All this suggests the average SPF surveys provide useful forecasts of inflation, which is a point made by Ang, Bekaert, and Wei (2007), Faust and Wright (2013), Mertens (2016), and Nason and Smith (2016a), among others.

Differences between the average SPF nowcast and 4-quarter ahead prediction contain information to identify τ_t , ε_t , $F_t\tau_t$, and $F_t\varepsilon_t$. For example, the average SPF inflation nowcast peaks close to 10 percent during the 1973–75 recession and around the double dip recessions of the early 1980s as Figure 1(a) shows. The former peak in inflation falls moving from $\pi_{t,t+2}^{SPF}$ to $\pi_{t,t+5}^{SPF}$ in figures 1(b), 1(c), and 1(d). At a 4-quarter ahead horizon, the average SPF inflation prediction rises steadily from about three percent in the early 1970s to a peak greater than eight percent around the 1980 recession.¹⁹ Our estimates rely on this information, which is a function of the SPF inflation prediction horizon, to identify persistence, stickiness, and volatility in RE and SI trend and gap inflation.

4.2 Posterior Estimates of Ψ and Fit of the Joint DGPs

Table 3 lists full sample estimates of Ψ , $\hat{\Psi}$, for three joint DGPs. The DGPs combine the SI-prediction mechanism and SW-UC-SV-TVP-AR(1) model, SI-prediction mechanism and a SW-UC model in which no persistence, $\theta_t = 0$, only SV drives gap inflation, and a fixed parameter, $\lambda_t = \lambda$, SI-prediction mechanism and the SW-UC-SV-TVP-AR(1) model.

The restrictions on inflation gap persistence and the frequency of SI inflation updating affect $\hat{\Psi}$ in several ways. First, innovations to the RW of trend inflation SV are more volatile than innovations to the RW of gap inflation SV in the DGPs with drifting gap persistence because $\hat{\sigma}_\eta^2 > \hat{\sigma}_v^2$. However, $\hat{\sigma}_\eta^2$ is larger while $\hat{\sigma}_v^2$ is smaller in the DGP that estimates θ_t and λ_t . In

¹⁹Figure 1(d) shows $\pi_{t,t+5}^{SPF}$ is missing observations in 1969, 1970, and 1974. The KF is modified to accommodate the missing observations.

contrast, $\hat{\sigma}_\eta^2$ and $\hat{\sigma}_v^2$ are about equal in the DGP with $\theta_t = 0$ and close to the calibrated values Stock and Watson (2007) and Creal (2012) use to estimate the state of the SW-UC-SV model. Next, there is little variation in estimates of the scale volatility on innovations to the RWs of θ_t and λ_t , $\hat{\sigma}_\phi^2$ and $\hat{\sigma}_\kappa^2$, across the DGPs in which these parameters appear. The DGPs with drifting gap persistence produce estimates of the scale volatility on the measurement errors of SPF inflation predictions, $\hat{\sigma}_{\zeta,h}^2$, $h = 1, \dots, 5$, that are quantitatively similar. The converse is true for estimates of the scale volatility on the measurement errors of π_t , $\hat{\sigma}_{\zeta,\pi}^2$, because it is nearly twice as large in the DGP that estimates θ_t and λ_t compared with the other two DGPs.

Estimates of log marginal data densities (MDDs) appear at the bottom of table 3 for the three joint DGPs. Equation (11) is used to calculate $\mathfrak{L}(\Psi | y_{1:T})$, which is the log MDD for a joint DGP tied to Ψ . Standard errors of the log MDDs are beneath estimates of $\mathfrak{L}(\Psi | y_{1:T})$.²⁰ The estimates of $\mathfrak{L}(\Psi | y_{1:T})$, indicate the data have, at a minimum, a very strong preference for the joint DGP of the SI prediction mechanism and SW-UC-SV-TVP-AR(1) model. Hence, the rest of the paper reports evidence this joint DGP has for the stickiness, persistence, and volatility of τ_t , $F_{t|t}\tau_t$, ε_t , and $F_{t|t}\varepsilon_t$.

Figure 2 plots the PLE paths of $\hat{\sigma}_\eta^2$, $\hat{\sigma}_v^2$, $\hat{\sigma}_\phi^2$, and $\hat{\sigma}_\kappa^2$ consistent with the joint DGP favored by the data. The scale volatility parameters are plotted with solid (navy blue) lines and 68 and 90 percent uncertainty bands appear as dark and light shading in figures 2(a)–2(d). These figures show $\hat{\sigma}_\eta^2$ more than doubles, $\hat{\sigma}_v^2$ falls by about a third, σ_ϕ^2 rises by about a quarter, and σ_κ^2 changes little from the start to the end of sample. The PLE path of $\hat{\sigma}_\eta^2$ drifts up for much of the sample as seen in figure 2(a). However, the PLE paths of these parameters are smooth from the 2001 recession to the end of the sample. Also, the 68 percent uncertainty bands are tight for the most part in figure 2, but the 90 percent uncertainty bands are wider and on occasion display substantial variation.

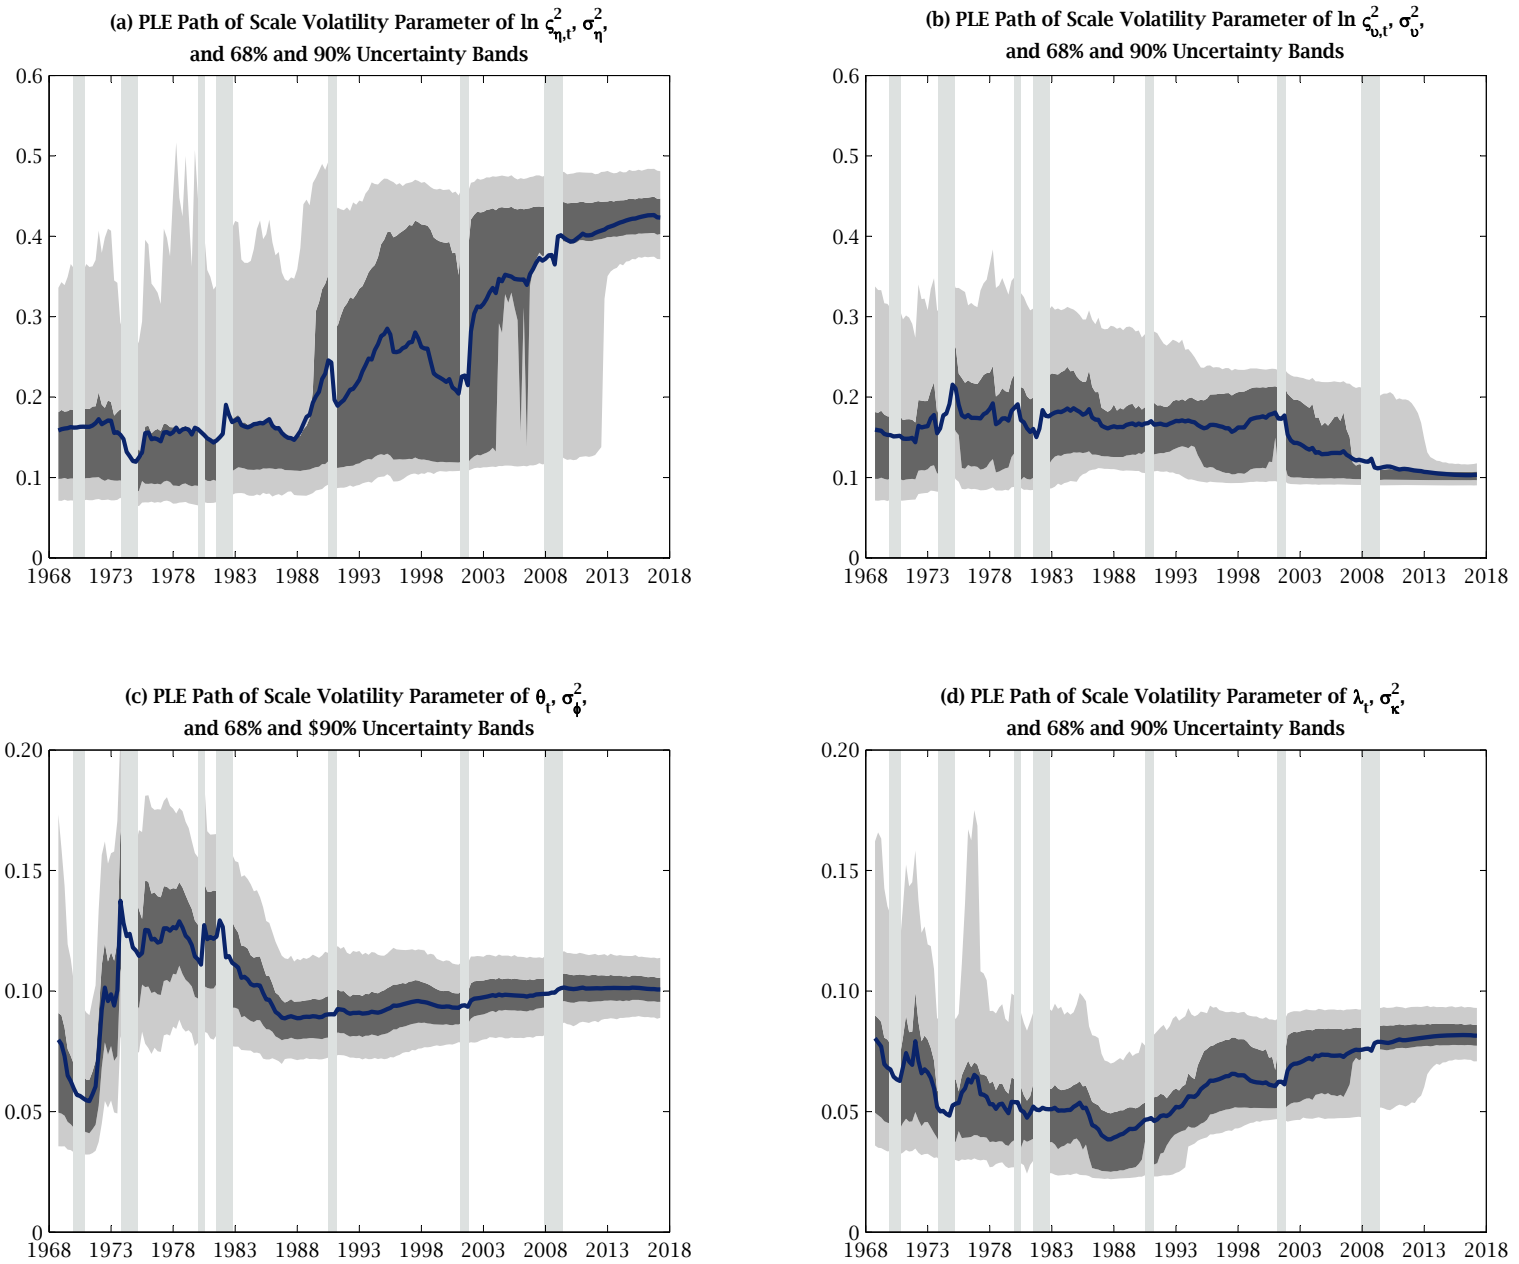
²⁰The standard errors are standard deviations of estimates of the log MDDs obtained from rerunning the PF using different random seeds across the three DGPs. Hence, the approximation error of the PF is measured by the standard errors of $\mathfrak{L}(\Psi | y_{1:T})$ and not the sampling uncertainty of a joint DGP.

Table 3: Posterior Estimates of the Joint DGPs of the SI Prediction Mechanism and SW-UC-SV Models

Parameter	TVP-SI: λ_t TVP-AR(1): θ_t	TVP-SI: λ_t Gap-SV: $\theta_t = 0$	Fixed SI: $\lambda_t = \lambda$ TVP-AR(1): θ_t
σ_η^2	0.423 [0.372, 0.481]	0.194 [0.167, 0.226]	0.336 [0.294, 0.385]
σ_v^2	0.103 [0.090, 0.118]	0.193 [0.163, 0.302]	0.191 [0.168, 0.218]
σ_ϕ^2	0.101 [0.089, 0.114]	-	0.107 [0.095, 0.121]
σ_κ^2	0.081 [0.071, 0.093]	0.084 [0.038, 0.099]	-
$\sigma_{\zeta,\pi}^2$	0.213 [0.171, 0.263]	0.115 [0.085, 0.148]	0.115 [0.090, 0.146]
$\sigma_{\zeta,1}^2$	0.148 [0.126, 0.173]	0.314 [0.266, 0.371]	0.148 [0.126, 0.175]
$\sigma_{\zeta,2}^2$	0.070 [0.059, 0.082]	0.093 [0.079, 0.110]	0.068 [0.057, 0.080]
$\sigma_{\zeta,3}^2$	0.052 [0.044, 0.061]	0.053 [0.044, 0.062]	0.052 [0.044, 0.061]
$\sigma_{\zeta,4}^2$	0.046 [0.039, 0.055]	0.068 [0.058, 0.080]	0.047 [0.040, 0.55]
$\sigma_{\zeta,5}^2$	0.048 [0.040, 0.056]	0.098 [0.083, 0.116]	0.048 [0.041, 0.056]
$\mathcal{L}(\Psi y_{1:T})$	-473.132 (7.068)	-669.150 (5.823)	-483.996 (6.691)

The table presents posterior means of the elements of Ψ , which are calculated using the full sample at date $T = 2017Q2$. The values in brackets below the posterior means are 5 and 95 percent quantiles. The model in which the SI parameter is fixed yields the posterior mean $\hat{\lambda} = 0.304$ with 5 and 95 percent quantiles of 0.250 and 0.360 conditional on the data and priors. The log MDDs are computed using the formula for $\mathcal{L}(\Psi | y_{1:T})$ described by equation (11) in section 3.3. Volatility over the log MDDs are measured by standard errors that appear in parentheses. The estimates of the static scale volatility parameters and log marginal data densities are created using $M = 100,000$ particles.

FIGURE 2: ESTIMATES OF STATIC VOLATILITY PARAMETERS, 1968Q4 TO 2017Q2



Note: The dark (light) gray areas surrounding estimates of the static scale volatility parameters, σ_{η}^2 , σ_v^2 , σ_{ϕ}^2 , and σ_{κ}^2 cover 68 (90) percent uncertainty bands. The four plots contain vertical gray bands that denote NBER dated recessions.

4.3 Trend and Gap Inflation

Figure 3 contains π_t , the average SPF inflation nowcast and 4-quarter ahead inflation prediction, $\pi_{t,t+1}^{SPF}$ and $\pi_{t,t+5}^{SPF}$, filtered RE trend inflation, $\tau_{t|t}$, filtered SI trend inflation, $F_{t|t}\tau_t$, filtered RE gap inflation, $\varepsilon_{t|t}$, and filtered SI gap inflation, $F_{t|t}\varepsilon_t$, on the 1968Q4 to 2017Q2 sample. Plots of $\pi_{t,t+1}^{SPF}$, $F_{t|t}\tau_t$, and its 68 percent uncertainty bands are in figure 3(a). Figure 3(b) is similar, but replaces $\pi_{t,t+1}^{SPF}$ with $\pi_{t,t+5}^{SPF}$. In these figures, solid (blue) lines are average SPF inflation predictions and $F_{t|t}\tau_t$ is the dotted (black) lines. Figure 3(c) displays $\tau_{t|t}$ with a dash (green) line, $F_{t|t}\tau_t$ with a dotted (black) line, and π_t with a dot-dash (red) line. Estimates of RE and SI gap inflation appear in figure 3(d) as a dashed (green) line, $\varepsilon_{t|t}$, and dotted (black) line, $F_{t|t}\varepsilon_t$.

Estimates of SI trend inflation are informed by the 1973–75 recession, inflation surge of the late 1970s and early 1980s, and Volcker disinflation.²¹ In 1974Q4, figure 3(a) displays a spike in $\pi_{t,t+1}^{SPF}$ of nearly 10 percent, but $F_{t|t}\tau_t$ is only 3.8 percent. At the same time, $\pi_{t,t+5}^{SPF}$ is 6.1 percent. The peaks in $\pi_{t,t+5}^{SPF}$ and $F_{t|t}\tau_t$, which occur a year and a half later, are close to 6.5 percent. The next peaks in $\pi_{t,t+1}^{SPF}$ and $\pi_{t,t+5}^{SPF}$ are 9.5 in 1979Q4 and 8.3 percent in 1980Q1. However, only in 1981Q2 does $F_{t|t}\tau_t$ peak at 7.5 percent. After 1983, $\pi_{t,t+1}^{SPF}$, $\pi_{t,t+5}^{SPF}$, and $F_{t|t}\tau_t$ fall steadily before leveling off in the late 1990s as figures 3(a) and 3(b) show. However, $F_{t|t}\tau_t$ often deviates from $\pi_{t,t+1}^{SPF}$ between 1983 and 2000. As a result, $\pi_{t,t+1}^{SPF}$ often is outside the 68 percent uncertainty bands of $F_{t|t}\tau_t$ during this period while $\pi_{t,t+5}^{SPF}$ falls within the 68 percent uncertainty bands of $F_{t|t}\tau_t$ after the Volcker disinflation in figure 3(b).

Figure 3(c) has several interesting features. First, π_t is volatile compared with $\tau_{t|t}$ and $F_{t|t}\tau_t$. Another striking aspect of figure 3(c) is $\tau_{t|t}$ and $F_{t|t}\tau_t$ are nearly identical for much of the sample. This is not true for π_t and $F_{t|t}\tau_t$ (or $\tau_{t|T}$) from 1968Q4 to 2000. For example, $\tau_{t|t}$ and $F_{t|t}\tau_t$ are less than a third of π_t during the first oil price shock. However, $F_{t|t}\tau_t$ explains much of the increases in π_t and $\pi_{t,t+1}^{SPF}$ by the late 1970s and early 1980s. Hence, $\tau_{t|t}$ and $F_{t|t}\tau_t$ respond slowly to the first oil price shock, but the inflation shock of the late 1970s and early 1980s produces quicker responses in $\tau_{t|t}$ and $F_{t|t}\tau_t$. Subsequently, π_t is often less than $\tau_{t|t}$

²¹Meltzer (2014, p. 1209) argues that 1986 marks the end of the Volcker disinflation.

and $F_{t|t}\tau_t$ from 1983 to 2000. Beginning in 2003, $\tau_{t|t}$ and $F_{t|t}\tau_t$ are often centered on π_t .

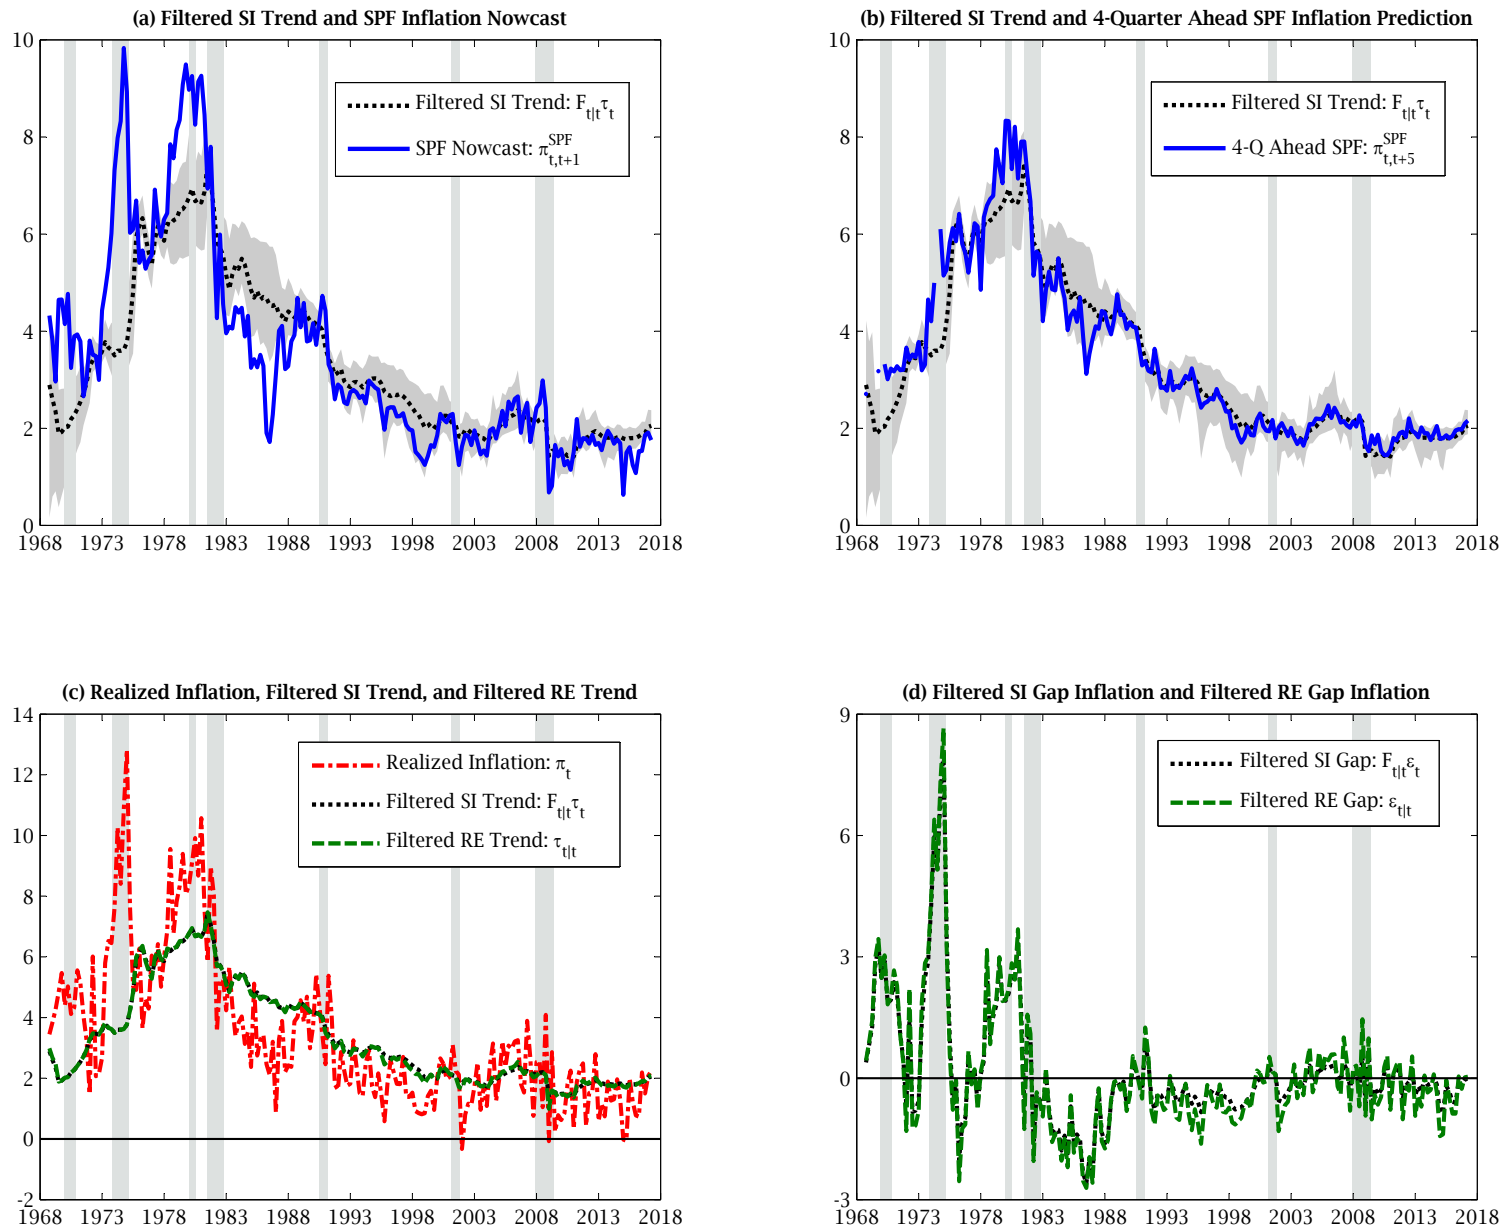
The estimates RE and SI trend inflation are a counterpoint to studies in which gap inflation dominates movements in inflation; see Cogley and Sbordone (2008) among others. One reason is $\tau_{t|t}$ and $F_{t|t}\tau_t$ condition on $\pi_{t,t+h}^{SPF}$, $h = 1, \dots, 5$. This differs from studies that rely on univariate SW-UC-SV models; for example see Grassi and Prioretto (2010), Creal (2012), and Shephard (2013).

We plot $\varepsilon_{t|t}$ and $F_{t|t}\varepsilon_{t|t}$ in figure 3(d). These plots show $\varepsilon_{t|t}$ and $F_{t|t}\varepsilon_t$ are nearly identical for the 1968Q4-2017Q2 sample. These estimates of gap inflation rise from less than one percent in 1968Q4 to about 3.5 percent in 1970. Thereafter, $\varepsilon_{t|t}$ and $F_{t|t}\varepsilon_t$ turn negative before the 1973-75 recession, which coincides with the largest spikes in $\varepsilon_{t|t}$ and $F_{t|t}\varepsilon_t$ of nearly nine percent. These spikes are followed by $\varepsilon_{t|t}$ and $F_{t|t}\varepsilon_t$ falling to about -2.5 percent by 1976. From the late 1970s to 1981, $\varepsilon_{t|t}$ and $F_{t|t}\varepsilon_t$ range from about zero to 3.7 percent.

There are two more aspects of figure 3(d) worth discussing. First, $\varepsilon_{t|t}$ and $F_{t|t}\varepsilon_t$ are less volatile subsequent to the Volcker disinflation compared with the 1970s. After 1983, (the absolute values of) $\varepsilon_{t|t}$ and $F_{t|t}\varepsilon_t$ are never larger than three percent. Second, $\varepsilon_{t|t}$ and $F_{t|t}\varepsilon_t$ are often negative from 1983 to 2000, which leads the average SPF participant to expect an increase in future growth in realized inflation. Nelson (2008) explains this prediction is an implication of the Beveridge and Nelson (1981) decomposition, which is built into the SW-UC-SV-TVP-AR(1) model of the joint DGP. Hence, the average SPF participant believes the Volcker disinflation produced only a transitory drop in realized inflation.

Movements in $F_{t|t}\varepsilon_t$ have parallels in monetary policy. Remember the average SPF participant expects mean reversion in π_t during the 1973-75 recession. However, in the late 1970s the average SPF participant believe unit root dynamics dominates π_t . An explanation for this shift in the average SPF participant's beliefs about the inflation regime is discussed by Meltzer (2014, pp. 1006-1007). He notes that in the 1970s U.S. monetary policy makers did not distinguish permanent from transitory shocks. As a result, their responses to the first oil price shock contributed to unanchored inflation expectations by the late 1970s.

FIGURE 3: REALIZED INFLATION, SPF INFLATION PREDICTIONS, AND ESTIMATES OF TREND AND GAP INFLATION, 1968Q4 TO 2017Q2



Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation, $F_{t|t}\tau_t$. The vertical gray bands denote NBER dated recessions in the four charts.

The Volcker disinflation is another example. After 1983, π_t and $F_{t|t}\tau_t$ began to fall, but the drop in π_t is steeper as figure 3(c) shows. These plots are consistent with mostly negative realizations for $F_{t|t}\varepsilon_t$ from 1983 to 2000 as in figure 3(d). As discussed previously, we assign these movements in $F_{t|t}\tau_t$ and $F_{t|t}\varepsilon_t$ to the average SPF participant expecting a temporary fall in π_t during and after the Volcker disinflation. The assessment agrees with Goodfriend and King (2005) and Meltzer (2014, p. 1131). They argue households, firms, and investors expected only a transitory drop in inflation after 1983.

4.4 Trend and Gap Inflation Volatilities


Estimates of filtered and smoothed trend and gap inflation SVs appear in figure 4. Figures 4(a) and 4(c) contain dotted lines, which are $\varsigma_{\eta,t|t}$ (purple) and $\varsigma_{\nu,t|t}$ (teal). Dot-dashed (purple and teal) lines are $\varsigma_{\eta,t|T}$ and $\varsigma_{\nu,t|T}$ in figures 4(b) and 4(d). These figures also include 90 percent uncertainty bands, which are thinner solid (black) lines.

Figure 4 makes several points about $\varsigma_{\eta,t|t}$, $\varsigma_{\nu,t|t}$, $\varsigma_{\eta,t|T}$, and $\varsigma_{\nu,t|T}$. Figures 4(a) shows the largest peaks in $\varsigma_{\eta,t|t}$ occur in 1977, 1983, and 2009 while $\varsigma_{\nu,t|t}$ is dominated by a spike in 1975 in figure 4(c). Figures 4(b) and 4(d) display peaks in $\varsigma_{\eta,t|T}$ and $\varsigma_{\nu,t|T}$ during the 1981–82 recession and in 1975, respectively. Hence, these plots are more evidence shocks to gap inflation dominate movements in π_t and $\pi_{t,t+h}^{SPF}$ during the 1973–75 recession, but in the inflation surge of the late 1970s and early 1980s permanent shocks are more important.

Another revealing feature of figures 4(a) and 4(c) is the behavior of SV around NBER dated recessions. The filtered SVs, $\varsigma_{\eta,t|t}$ and $\varsigma_{\nu,t|t}$, often rise during or after a NBER recessions as depicted in figures 4(a) and 4(c). There are peaks $\varsigma_{\eta,t|T}$ ($\varsigma_{\nu,t|T}$) during the 1990–91 and 2007–09 (1981–82, 1990–91, 2001, and 2007–09) recessions.

Figure 4(b) and 4(d) are also informative about the long run behavior of $\varsigma_{\eta,t|T}$ and $\varsigma_{\nu,t|T}$. These SVs display steady declines for extended periods during the sample. The descent starts in 1983 for $\varsigma_{\eta,t|T}$ while this process starts in 1975 for $\varsigma_{\nu,t|T}$.

FIGURE 4: ESTIMATES OF THE STOCHASTIC VOLATILITY OF TREND AND GAP INFLATION, 1968Q4 TO 2017Q2



Note: The solid thin (black) lines around estimates of filtered and smoothed trend and gap inflation SV are lower and upper bounds on 90 percent uncertainty bands. The four plots contain vertical gray bands that denote NBER dated recessions.

Finally, our estimates show $\zeta_{\eta,t|T}$ is smaller than $\zeta_{\nu,t|T}$ for the entire sample. These estimates differ from Grassi and Prioretto (2010), Stock and Watson (2010), Creal (2012), and Shephard (2013). These authors report trend SV dominates inflation gap SV from the 1970s well into the late 1990s. However, Creal and Shephard find that gap inflation SV is greater than trend SV after 2000.

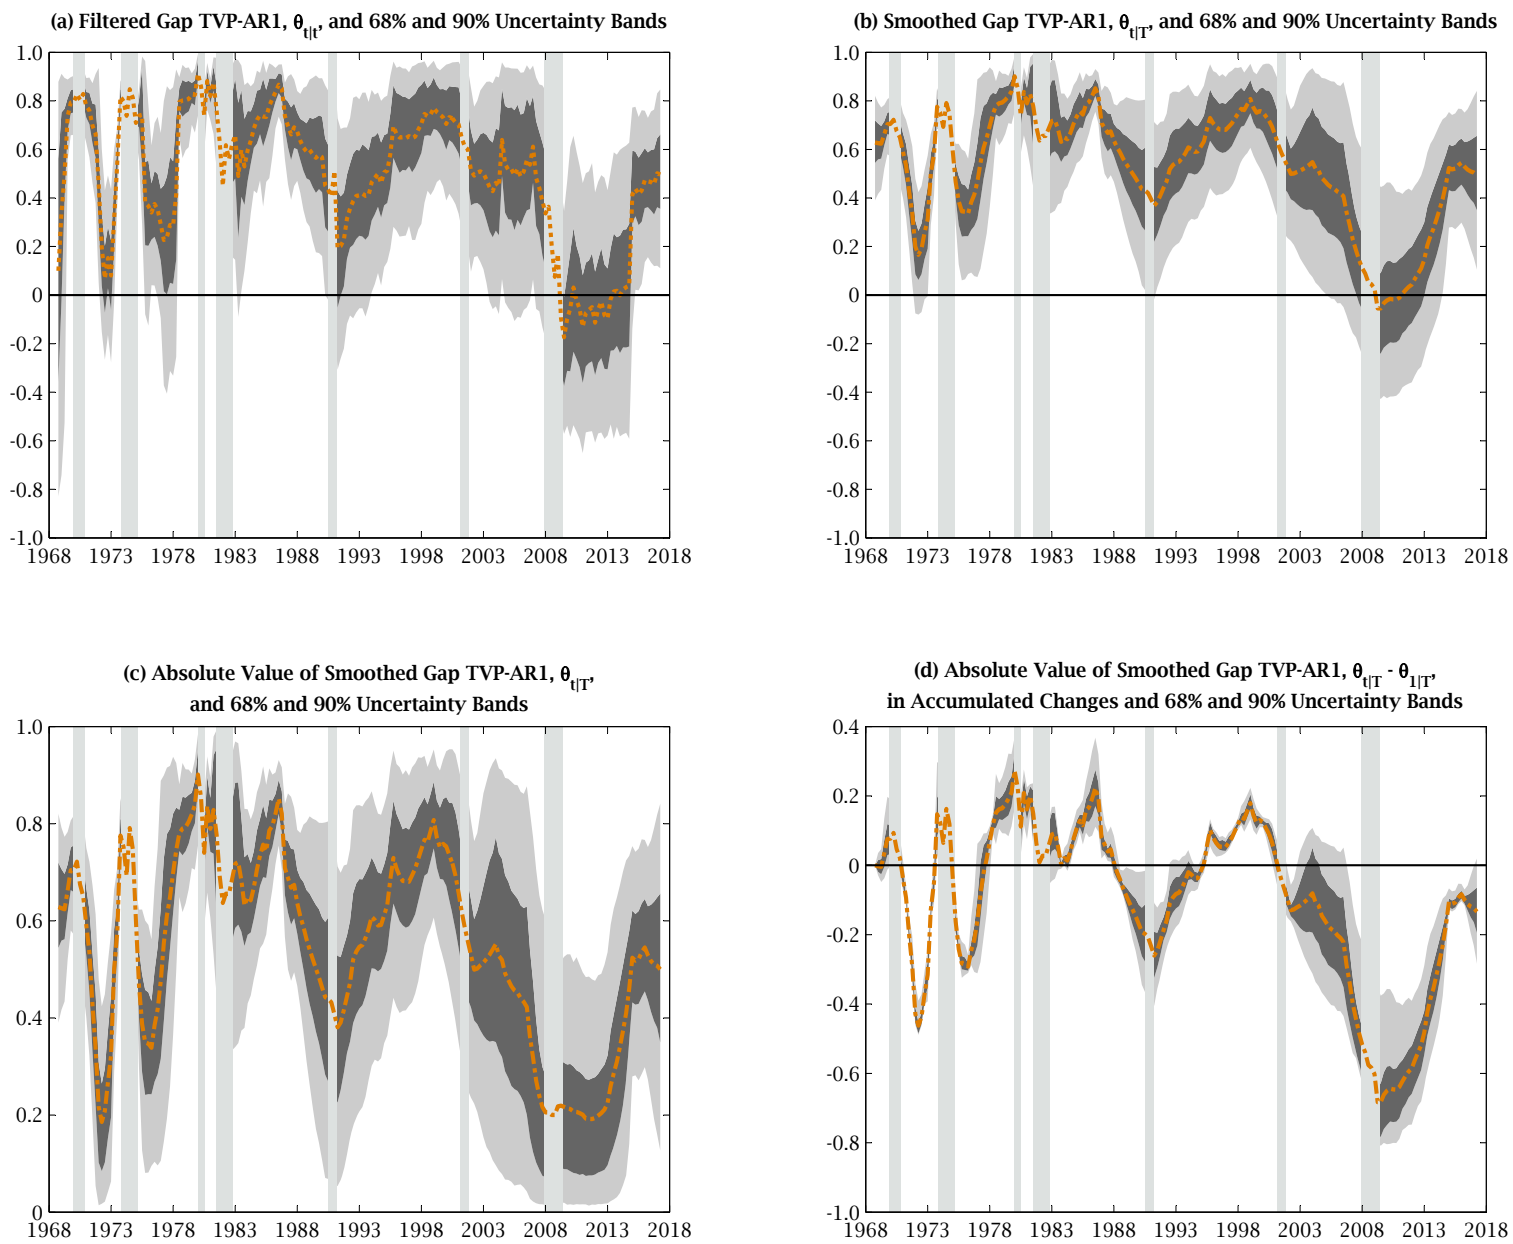
4.5 Drifting Inflation Gap Persistence

Figures 5(a) and 5(b) display filtered and smoothed estimates of drifting inflation gap persistence, $\theta_{t|t}$ and $\theta_{t|T}$. Dotted and dot-dash (orange) lines denote $\theta_{t|t}$ and $\theta_{t|T}$. Surrounding $\theta_{t|t}$ and $\theta_{t|T}$ are 68 and 90 percent uncertainty bands in the dark and light gray shaded areas. Figures 5(c) and 5(d) plot the absolute value of smoothed inflation gap persistence, $|\theta_{t|T}|$, and accumulated changes of this absolute value, $|\theta_{t|T}| - |\theta_{1|T}|$. These plots depict $|\theta_{t|T}|$ and $|\theta_{t|T}| - |\theta_{1|T}|$ with dot-dashed (orange) lines, where the dark and light gray shaded areas are 68 and 90 percent uncertainty bands.

There is co-movement between $\theta_{t|t}$ and $\theta_{t|T}$ with NBER dated cycles in figures 5(a) and 5(b). The co-movement is procyclical during the 1969-70, 1973-75, and 1980 recessions. These recessions see peaks in $\theta_{t|t}$ and $\theta_{t|T}$ while there are troughs between these recession. Post-1981, $\theta_{t|t}$ and $\theta_{t|T}$ turn countercyclical. Filtered and smoothed estimates of drifting inflation gap persistence peak between the recessions of 1981-82, 1990-91, 2001, and 2007-09 while these recessions see troughs in $\theta_{t|t}$ and $\theta_{t|T}$.

Uncertainty bands of $\theta_{t|t}$ and $\theta_{t|T}$ also appear in figures 5(a) and 5(b). The 90 percent quantiles of $\theta_{t|T}$ ($\theta_{t|t}$) cover zero in 1971-72, 1990-91, and 2006-14 (1968-69, 1972-73, 1975, 1976-78, 1983, 1990-93, and 2003-14). Hence, we infer there are episodes in which inflation gap persistence is zero. These results are similar to evidence presented by Cogley, Primiceri, and Sargent (2010). They find inflation gap persistence drops after 1983. However, our evidence is tied to procyclical troughs in $\theta_{t|T}$ before 1983 and to the 2007-09 recession and its aftermath, which occurs more than 20 years after the Volcker disinflation.

FIGURE 5: ESTIMATES OF TIME-VARYING INFLATION GAP PERSISTENCE, 1968Q4 TO 2017Q2



Note: The dark (light) gray areas surrounding estimates of the TVP-AR1 of gap inflation cover 68 (90) percent uncertainty bands. The four plots contain vertical gray bands that denote NBER dated recessions.

Another take on the statistical and economic significance of drifting gap inflation persistence appears in figure 5(c). This figure displays the absolute value of $\theta_{t|T}$, $|\theta_{t|T}|$. The plot of $|\theta_{t|T}|$ gives evidence similar to that found in figure 5(b). There is evidence of a shift in business cycle behavior of $|\theta_{t|T}|$ around the Volcker disinflation. Drift in the absolute value of inflation gap persistence also declines steadily from the late 1990s to 2013.

There remains the inference problem that $\theta_{t|t}$, $\theta_{t|T}$, and $|\theta_{t|T}|$ are not necessarily informative about the statistical and economic content of changes in drifting inflation gap persistence during the sample. We address this problem by plotting accumulating changes in $|\theta_{t|T}|$, $|\theta_{t|T}| - |\theta_{1|T}|$, in figure 5(d). Figure 5(d) shows these changes have tighter uncertainty bands compared with the plots in figures 5(a), 5(b), and 5(c). Nonetheless, the path of $|\theta_{t|T}| - |\theta_{1|T}|$ continues to show peaks that coincide with pre-1981 recessions and troughs occurs between these recessions. The opposite is observed post-1981.

Hence, figure 5 gives evidence that dates a switch from procyclical to countercyclical drift in inflation gap persistence to 1981. This break is consistent with an argument made by Meltzer (2014, p.1006 and p.1207). He contends there was a shift in the pattern of U.S. inflation persistence because of changes to the way the Fed operated monetary policy in the 1980s and 1990s compared with the 1970s.

4.6 Time Variation in the Frequency of SI Updating

Figure 6 presents filtered and smoothed estimates of the time variation in the frequency of SI updating, $\lambda_{t|t}$ and $\lambda_{t|T}$. These panels plot $\lambda_{t|t}$ and $\lambda_{t|T}$ as dotted (light green) and dot-dashed (brick) lines. The thin solid (brick) lines denote 90 percent uncertainty bands of $\lambda_{t|T}$ and 90 percent uncertainty bands of $\lambda_{t|t}$ are depicted with light gray areas. Figures 6(b) and 6(d) plot accumulated changes in $\lambda_{t|T}$, $\lambda_{t|T} - \lambda_{1|T}$. In these panels, dark and light gray areas are 68 and 90 percent uncertainty bands of $\lambda_{t|T} - \lambda_{1|T}$. The top row of figure 6 has $\lambda_{t|t}$, $\lambda_{t|T}$, and $\lambda_{t|T} - \lambda_{1|T}$ estimated using the joint DGP of the SI-prediction mechanism and the SW-UC-SV-TVP-AR(1) model. Figures 6(c) and 6(d) report similar estimates, but the SW-UC-SV model lacks

persistence in gap inflation, or $\theta_t = 0$ for all dates t .


Plots of $\lambda_{t|t}$ and $\lambda_{t|T}$ display a decade long swing from more frequent to less frequent updating beginning in the late 1980s in figure 6(a). From the late 1960s to the 1988, the average SPF inflation respondent is estimated to update almost every quarter to changes in $E_t \pi_{t+h}$ because $\lambda_{t|T}$ varies between 0.01 and 0.35. However, there is uncertainty about these estimates because the 90 percent confidence bands of $\lambda_{t|T}$ range from 0.01 to 0.60.

Figures 6(a) also shows $\lambda_{t|t}$ and $\lambda_{t|T}$ reach a plateau from 1994 to 2007 before falling during the 2007–09 recession. From 1995 to 2008, $\lambda_{t|t}$ and $\lambda_{t|T}$ range between 0.50 and 0.70. The recession of 2007–2009 sees $\lambda_{t|T}$ ($\lambda_{t|t}$) dropping to 0.25 (0.35). Subsequently, $\lambda_{t|T}$ ($\lambda_{t|t}$) recovers to 0.47 (0.60) before 2017Q2. The filtered and smoothed estimates of λ_t are also associated with substantial uncertainty. For example, when $\lambda_{t|T}$ plateaus in the late 1990s, the five percent quantile is as low as 0.20 and the 95 percent quantile is as high as 0.95. Furthermore, the 90 percent uncertainty bands of $\lambda_{t|t}$ and $\lambda_{t|T}$ remain wide in figure 6(a) as the sample moves past the recession of 2001, the “considerable” and “extended” period policy regimes of the Greenspan and Bernanke Feds of the early 2000s, the 2007–09 recession, and unconventional policy regimes of the Bernanke and Yellen Feds.

There are useful inferences to draw from $\lambda_{t|t}$ and $\lambda_{t|T}$, even with the uncertainty surrounding these estimates. For example, infrequent SI inflation updating by the average member of the SPF lets the Fed engage in a policy of “opportunistic disinflation” during the 1990s as described by Meyer (1996) and Orphanides and Wilcox (2002). Orphanides and Wilcox argue that in the mid 1990s Fed policy makers advocated to wait for a state of the world in which there is little cost to monetary policy lowering inflationary expectations rather than to take actions during periods when the potential for a costly disinflation are large. However, since the joint DGP of the SI-prediction mechanism and SW-UC-SV-TVP-AR(1) model is the source of $\lambda_{t|t}$ and $\lambda_{t|T}$, we only have estimates of the average SPF respondent’s beliefs about changes in the inflation regime and not evidence about shifts in the monetary policy regime.²²

²²Information about monetary policy interventions is needed to conduct a monetary policy evaluation of this kind as studied, for example, by Leeper and Zha (2003).

FIGURE 6: ESTIMATES OF THE TIME-VARYING SI PARAMETER, 1968Q4 TO 2017Q2



Note: The dark (light) gray areas surrounding estimates of the SI-TVP cover 68 (90) percent uncertainty bands. The left column of charts also displays solid thin (red) lines around smoothed estimates of the SI-TVP, $\lambda_{t|T}$, that are lower and upper bounds on 90 percent uncertainty bands. The four plots contain vertical gray bands that denote NBER dated recessions.

There is greater support for statistically and economically important time variation in the frequency of SI inflation updating in figure 6(b). This figure plots $\lambda_{t|T} - \lambda_{1|T}$ for the joint DGP in which there is drift in inflation gap persistence. In this case, the path of $\lambda_{t|T} - \lambda_{1|T}$ in figure 6(b) is similar to $\lambda_{t|T}$ displayed in figure 6(a) with respect to level and slope. Another interesting feature of figure 6(b) is the uncertainty bands surrounding $\lambda_{t|T} - \lambda_{1|T}$. Figure 6(b) displays 90 percent uncertainty bands of $\lambda_{t|T} - \lambda_{1|T}$ that are narrower for the entire sample compared with the analogous confidence bands of $\lambda_{t|T}$ in figure 6(a). These estimates strengthen the case that changes in the frequency of SI inflation updating by the average member of the SPF are statistical and economic important.

This message is reinforced by figure 6(d). This figure presents estimates of $\lambda_{t|T} - \lambda_{1|T}$ conditional on a joint DGP in which there is no persistence in the inflation gap. Given θ_t is zero, SI inflation updating is less frequent quarter by quarter, as depicted by $\lambda_{t|t}$ and $\lambda_{t|T}$ in figure 6(c) compared with the estimates found in figure 6(a). Although figure 6(c) suggests that there is useful information about the frequency of SI inflation updating conditional on $\theta_t = 0$, the plot of $\lambda_{t|T} - \lambda_{1|T}$ in figure 6(d) indicates otherwise. Figure 6(d) depicts $\lambda_{t|T} - \lambda_{1|T}$ as fluctuating around zero with 90 percent uncertainty bands that often contain zero under the joint DGP in which inflation gap has no persistence.

This section reports estimates of $\lambda_{t|t}$, $\lambda_{t|T}$, and $\lambda_{t|T} - \lambda_{1|T}$ shows that SI inflation updating by the average SPF respondent is statistically and economically significant for the last 48 years. These results agree with Coibion and Gorodnichenko (2015). Nonetheless, our estimates also reveal shifts in SI inflation updating during the sample. From the 1969 to 1988, the frequency of SI inflation updating occurred almost every quarter. The frequency declines to about once every two to three quarter until 2007, followed by a sharp increase during the 2007-09 recession. Afterwards, the frequency drops by 2017Q2. These shifts in estimates of SI inflation updating indicate the average SPF participant's beliefs about the inflation regime changed within a few years of the end of the Volcker disinflation. The average SPF participant's beliefs about the inflation regime also appear to have been altered by the recession of 2007-09.


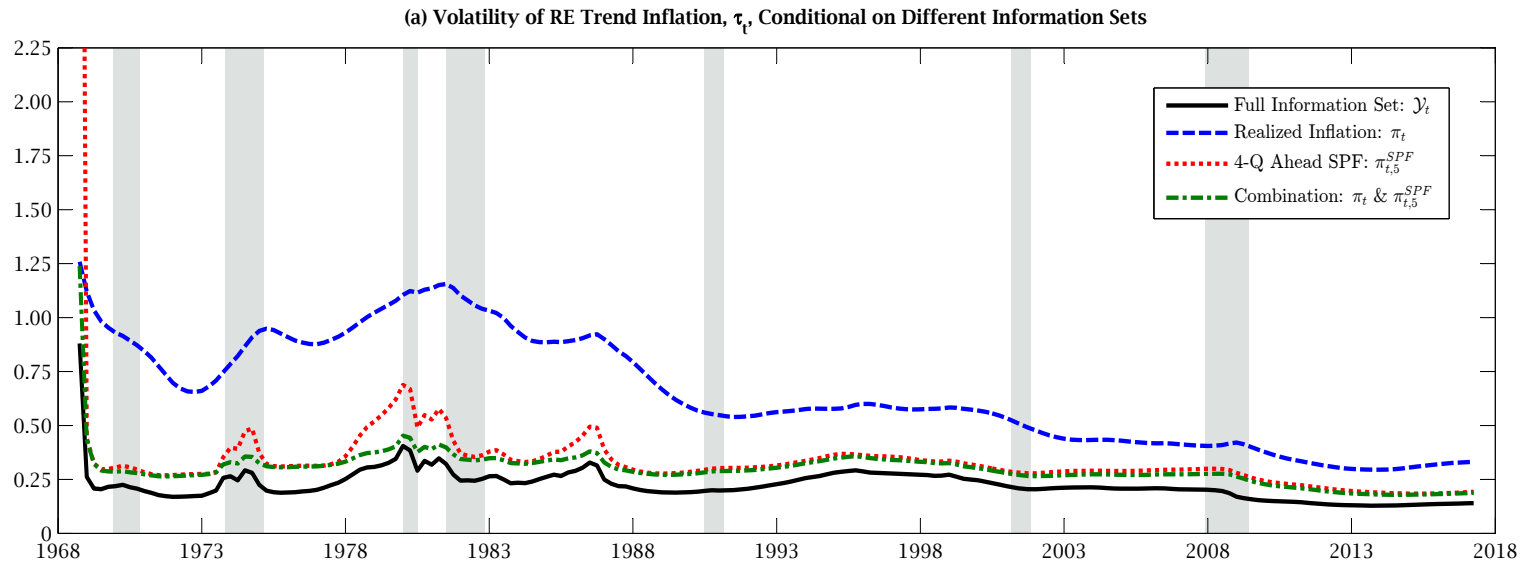
4.7 SPF Inflation Predictions and Trend Inflation Uncertainty

Figure 7 displays conditional volatilities of RE trend inflation, τ_t , and SI trend inflation, $F_t\tau_t$. The plots quantify uncertainty over time in τ_t and $F_t\tau_t$ conditional on the history of \mathcal{Y}_t , or histories of subsets of its elements, smoothed estimates of the nonlinear states, $\tilde{\nu}_{t|T}$, and estimates of the static scale volatility coefficients, $\hat{\Psi}$. The measure of the volatility of τ_t is $\text{Var}(\tau_t | \mathcal{Y}_{1:t}, \tilde{\nu}_{t|T}, \hat{\Psi})$, where the entire information set runs from the first observation to quarter t , the smoothed nonlinear states begin at quarter t and end with quarter T , and estimates of the static scale volatility parameters are full sample. Similar computations are used to produce the conditional volatility of $F_t\tau_t$. Thus, the paths of the nonlinear states and parameter estimates are held fixed across changes in the sample data fed into the KF to produce estimates of the conditional volatilities of τ_t and $F_t\tau_t$.

Figure 7(a) plots the conditional volatilities of τ_t . The conditional volatilities of $F_t\tau_t$ are found in figure 7(b). In these figures, the solid (black) line, dashed (blue) line, dotted (red) line, and dot-dashed (green) line are $\text{Var}(x | \mathcal{Y}_{1:t}, \tilde{\nu}_{t|T}, \hat{\Psi})$, $\text{Var}(x | \pi_{1:t}, \tilde{\nu}_{t|T}, \hat{\Psi})$, $\text{Var}(x | \pi_{1:t}^{SPF}, \tilde{\nu}_{t|T}, \hat{\Psi})$, and $\text{Var}(x | \pi_{1:t}, \pi_{1:t}^{SPF}, \tilde{\nu}_{t|T}, \hat{\Psi})$, respectively, where $x = \tau_t, F_t\tau_t$.

Figures 7(a) and 7(b) reveal π_t and $\pi_{t,t+5}^{SPF}$ jointly contribute the bulk of the information pertinent to estimate τ_t and $F_t\tau_t$. The reason is the dot-dashed (green) lines of figures 7(a) and 7(b) are always near the solid (black) lines. Hence, given only π_t and $\pi_{t,t+5}^{SPF}$, $\text{Var}(\tau_t | \pi_t, \pi_{1:t,t+5}^{SPF}, \tilde{\nu}_{t|T}, \hat{\Psi})$ and $\text{Var}(F_t\tau_t | \pi_t, \pi_{1:t,t+5}^{SPF}, \tilde{\nu}_{t|T}, \hat{\Psi})$ are close to the estimates conditioned on the entire information set, $\text{Var}(\tau_t | \mathcal{Y}_t, \tilde{\nu}_{t|T}, \hat{\Psi})$ and $\text{Var}(F_t\tau_t | \mathcal{Y}_t, \tilde{\nu}_{t|T}, \hat{\Psi})$. In contrast, the dotted (blue) lines are far from the solid (black) and large dot-dashed (green) lines in the first half of the sample. Hence, prior to the Volcker disinflation, there is insufficient information in π_t alone to estimate τ_t and $F_t\tau_t$ without also generating more variation in these estimates compared with estimates conditioning on either \mathcal{Y}_t or π_t and $\pi_{t,t+5}^{SPF}$. However, conditioning only on $\pi_{t,t+5}^{SPF}$ produces substantial variation around $F_t\tau_t$ that is manifested as large differences between plots of $\text{Var}(F_t\tau_t | \pi_{1:t,t+5}^{SPF}, \tilde{\nu}_{t|T}, \hat{\Psi})$ and $\text{Var}(F_t\tau_t | \mathcal{Y}_t, \tilde{\nu}_{t|T}, \hat{\Psi})$ or $\text{Var}(F_t\tau_t | \pi_t, \pi_{1:t,t+5}^{SPF}, \tilde{\nu}_{t|T}, \hat{\Psi})$ in figure 7(b).

FIGURE 7: UNCERTAINTY MEASURE OF TREND INFLATION CONDITIONAL ON DIFFERENT INFORMATION SETS, 1968Q4 TO 2017Q2



Note: The two plots contain vertical gray bands that denote NBER dated recessions.

In summary, figure 7 shows realized inflation and the 4-quarter ahead average SPF inflation prediction contain much of the information useful for reducing uncertainty surrounding τ_t and $F_t \tau_t$ and, hence, efficiently estimating these measures of trend inflation.

5 Conclusions

This paper studies the joint dynamics of realized inflation and inflation predictions of the Survey of Professional Forecasters (SPF). The joint data generating process (DGP) mixes a Stock and Watson (2007) unobserved components (SW-UC) model with the Coibion and Gorodnichenko (2015) version of the Mankiw and Reis (2002) sticky information (SI) model. The SW-UC model with stochastic volatility (SV) in trend and gap inflation is extended to include drift in inflation gap persistence. The SI law of motion is endowed with drift in the SI inflation updating parameter. We estimate the joint DGP on a sample of real-time realized inflation and averages of SPF inflation predictions from 1968Q4 to 2017Q2. The estimator embeds a Rao-Blackwellized auxiliary particle filter into the particle learning estimator of Storvik (2002). Smoothed estimates of the state variables are constructed using an algorithm developed by Lindsten, Bunch, Särkkä, Schön, and Godsill (2016).

There are five key results to draw from our estimates. First, longer horizon average SPF inflation predictions provide useful information for estimating rational expectations (RE) and SI trend inflation and reducing uncertainty around these estimates. Second, RE and SI inflation gaps dominate inflation fluctuations during the first oil price shock. This is reversed during the late 1970s and early 1980s. Third, trend (gap) inflation SV falls steadily after 1983 (1975). We also find that inflation gap persistence is procyclical before 1981 and turns countercyclical afterwards. Fifth, changes in the frequency of SI inflation updating are statistically and economically important. The average SPF participant is updating SI inflation predictions often from the late 1960s through the late 1980s. Subsequently, the frequency of SI inflation updating falls to levels associated with estimates reported by Coibion and Gorodnichenko (2015),

among others, and remains low until the 2007–09 recession.

Our results fit into a literature represented by, among others, Krane (2011), and Nason and Smith (2016a, b). These authors find that the responses of professional forecasters to permanent shocks are greater than for those to transitory shocks when revising their predictions. In the same way that this research inspired us, we hope that this paper stimulates further work on the ways in which professional forecasters and other economic agents process information to form beliefs and predictions about future economic 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.
- Bernanke, B.S. (2007). Inflation Expectations and Inflation Forecasting. Remarks given at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Massachusetts (July 10). Available at <http://www.federalreserve.gov/newsevents/speech/Bernanke20070710a.htm>.
- Beveridge, S., C.R. Nelson. (1981). A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the Business Cycle. *Journal of Monetary Economics* 7, 151–174.
- Carvalho, C.M., M.S. Johannes, H.F. Lopes, and N.G. Polson (2010). Particle learning and smoothing. *Statistical Science* 25, 88–106.
- Chen, R., J.S. Liu (2000). Mixture Kalman filters. *Journal of the Royal Statistical Society, Series B* 62, 493–508.
- 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-level uncertainty and instability in the U.S., 1850–2012. *Review of Economics and Statistics* 97, 827–838.
- 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., 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* 105, 2644–2678.
- 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.
- 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.
- Godsill, S.J., A. Doucet, M. West (2004). Monte Carlo smoothing for nonlinear time series. *Journal of the American Statistical Association* 99, 156–168.
- Goodfriend, M., R.G. King (2005). The incredible Volcker disinflation. *Journal of Monetary Economics* 52, 981–1015.
- 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 (2016). BAYESIAN INFERENCE FOR DSGE MODELS. Princeton, NJ: Princeton University Press.
- Hol, J.D., T.B. Schön, F. Gustafsson (2006). On resampling algorithms for particle filters. In 2006 IEEE NONLINEAR STATISTICAL SIGNAL PROCESSING WORKSHOP, Ng, W. (ed.), 79–82. Red Hook, NY: Curran Associates, Inc.
- 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.
- Lindsten, F., P. Bunch, S. Särkkä, T.B. Schön, S.J. Godsill (2016). Rao-Blackwellized particle smoothers for conditionally linear Gaussian models. *IEEE Journal of Selected Topics in Signal Processing* 10, 353–365.

- Lopes, H.F., R.S. Tsay (2011). Particle filters and Bayesian inference in financial econometrics. *Journal of Forecasting* 30, 168–209.
- 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.
- Meltzer, A.H. (2014). A HISTORY OF THE FEDERAL RESERVE: VOLUME 2, BOOK II, 1970–1986. Chicago, IL: The University of Chicago Press.
- Mertens, E. (2016). Measuring the level and uncertainty of trend inflation. *The Review of Economics and Statistics* 98, 950–967.
- Meyer, L. (1996). Monetary policy objectives and strategy. Remarks given at the National Association of Business Economists 38th Annual Meeting, Boston, MA (September 8). Available at <http://www.federalreserve.gov/boarddocs/speeches/1996/19960908.htm>.
- Muth, J.F. (1960). Optimal properties of exponentially weighted forecasts. *Journal of the American Statistical Association* 55, 299–306.
- Nason, J.M., G.W. Smith (2016a). Sticky professional forecasts and the unobserved components model of US inflation. Manuscript, Department of Economics, Queen's University.
- Nason, J.M., G.W. Smith (2016b). Measuring the slowly evolving trend in US inflation with professional forecasts. Manuscript, Department of Economics, Queen's 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.
- Pitt M.K., N. Shephard, N. (2001). Auxiliary variable based particle filters. In Doucet, A., de Freitas, N., Gordon, N. (eds.) SEQUENTIAL MONTE CARLO METHODS IN PRACTICE, New York: Springer-Verlag.
- Pitt, M. K., N. Shephard (1999). Filtering via simulation: auxiliary particle filters. *Journal of the American Statistical Association* 94, 590–599.
- 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.
- Storvik, G. (2002). Particle filters for state-space models with the presence of unknown static parameters. *IEEE Transactions on Signal Processing* 50, 281–289.