

# Measuring the Slowly Evolving Trend in US Inflation with Professional Forecasts

James M. Nason and Gregor W. Smith<sup>†</sup>

April 2020  
first version September 2012

## Abstract

---

Much research studies US inflation history with a trend-cycle model with unobserved components, where the trend may be viewed as the Fed's evolving inflation target or long-horizon expected inflation. We provide a novel way to measure the slowly evolving trend and the cycle (or inflation gap), by combining inflation predictions from the Survey of Professional Forecasters (SPF) with realized inflation. The SPF forecasts may be treated either as rational expectations (RE) or updating according to a sticky information (SI) law of motion. We estimate RE and SI state space models with stochastic volatility on samples of CPI and GNP/GDP deflator inflation and the associated SPF inflation predictions using a particle Metropolis-Markov chain Monte Carlo sampler. The trend converges to 2% and its volatility declines over time, two tendencies largely complete by the late 1990s.

---

*JEL classification:* E31, E37.

*Keywords:* US inflation, professional forecasts, sticky information, Beveridge-Nelson

<sup>†</sup>Nason: Department of Economics, North Carolina State University;  
jmnason@ncsu.edu. Smith: Department of Economics, Queen's University;  
smithgw@econ.queensu.ca.

Acknowledgements: We thank the Social Sciences and Humanities Research Council of Canada, the Bank of Canada research fellowship programme, and the Jenkins Family Economics Fund of North Carolina State University for support of this research. For helpful comments, we thank Drew Creal, Luca Martino, Elmar Mertens, Umberto Picchini, Michael Pitt, Jonathan Wright (the editor), two anonymous referees, and participants at too many conferences and seminars to mention. The views expressed herein are not necessarily those of the Bank of Canada and are the authors alone. A technical appendix is available at: [https://www.jamesmnason.net/s/NasonSmith\\_appendix\\_26April2020.pdf](https://www.jamesmnason.net/s/NasonSmith_appendix_26April2020.pdf)

## 1. Introduction

For the past thirty years the unobserved-components (UC) model has been an informative lens through which economists have viewed US inflation dynamics. That statistical model decomposes inflation into permanent and transitory components. The permanent component or trend usually (and in this paper) is identified with the Beveridge and Nelson (1981) decomposition, meaning that it is a random walk. This decomposition has been widely adopted in forecasting inflation. For example, Stock and Watson (2007, 2010) use it to isolate changes in the variances of the components and hence in the overall persistence and forecastability of inflation over time. Faust and Wright (2013), in their review of inflation forecasting, list the many studies that feature a slowly evolving trend.

The Beveridge-Nelson (BN) decomposition also sheds light on inflation history. For example, Cogley and Sargent (2015) use the model to identify inflation moderations in a long span of annual data. A key feature of this model is that the trend component serves as a measure of long-horizon inflation expectations, an indicator of the Fed's credibility as well as a constraint on the effect of policy.

We use surveys of professional forecasts to measure the two components of inflation. A simple example illustrates the idea. First, suppose inflation is the sum of a random-walk trend,  $\tau_t$ , and an inflation gap,  $\epsilon_t$ , that is white noise. Thus,  $\tau_t$  also is the expectation of future inflation. Second, suppose professional forecasters report their rational expectations. Thus, their  $h$ -step-ahead forecasts directly provide an estimate of  $\tau_t$ , while  $\epsilon_t$  can be found by subtracting that trend estimate from realized inflation. We then show how to extend this idea to allow for a persistent inflation gap, to incorporate stochastic volatility (SV) into the innovations of both components, to include information from professional forecasts at multiple horizons, and to integrate sticky forecasts into this setting.

As this simple example shows, this approach requires a view on the connection between unobservable,  $h$ -step-ahead, rational-expectations (RE) forecasts of inflation, denoted  $E_t\pi_{t+h}$ , and the mean, reported inflation forecasts of professional forecasters, denoted  $F_t\pi_{t+h}$ . We obtain these inflation forecasts from the Survey of Professional Forecasters (SPF).

We consider two possibilities. One way to extract information from the SPF is to

assume that its mean  $h$ -step ahead inflation forecast coincides with the  $h$ -step ahead prediction made by the UC model under its information set. Second, we assume professional forecasters do not react to new information as prescribed by RE. Instead, new information produces less than a full response by professional forecasters. Thus, professional forecasts are sticky and can be modeled using the sticky information (SI) framework of Mankiw and Reis (2002).

The SI model is motivated by the considerable evidence that panels of professional forecasts are not consistent with full-information RE but rather exhibit bias. How can professional forecasts be useful if they are biased? Precisely because the pattern of forecast errors is systematic, these surveys provide evidence about the underlying expectations of the forecasters. Coibion and Gorodnichenko (2012, 2015) use the SI model to link reported forecasts to the actual conditional expectations of professional forecasters. The SI description of forecasts also is of interest because it has been widely used to close macroeconomic models, for example in studies of the New Keynesian Phillips curve.

We develop procedures to connect mean, reported forecasts to the UC model under either the RE or SI assumption about professional forecasts. In the latter case we estimate the parameter describing stickiness along with those of the UC model. We find that inflation-gap persistence is quite low. Whether inflation is measured by the CPI or by the GDP deflator, the trend converges to near 2% and its volatility declines over time, two tendencies largely complete by the late 1990s. Information stickiness also can be identified and is especially important to describe the 1970s.

Section 2 outlines the unobserved components model and its implications for forecasts under RE and SI respectively. These models of reported forecasts lead to sets of observation equations to be used in extracting the components of inflation. Section 3 briefly describes related research. Section 4 describes the Bayesian sequential Monte Carlo methods we adopt. A reader not interested in the estimator *per se* might skip this section, while one specifically interested in the estimator will find still more detail in the online appendix. Section 5 gives data sources and priors. Section 6 contains the findings.

## 2. The RE- and SI-SPF State Space Models

The state-space models depend on the Stock and Watson (2007, 2010) unobserved components (UC) model to generate realized inflation,  $\pi_t$ . Their model equates  $\pi_t$  to trend inflation,  $\tau_t$ , plus gap inflation,  $\epsilon_t$ . Trend inflation is a random walk while  $\epsilon_t$  is a first-order autoregression with AR1 parameter  $\rho \in (-1, 1)$ . Innovations to  $\tau_t$  and  $\epsilon_t$  are independent, but are subject to stochastic volatility (SV) that in each case evolves as a (log squared) random walk. Thus, the UC model is:

$$\pi_t = \tau_t + \epsilon_t, \tag{1}$$

$$\tau_t = \tau_{t-1} + \xi_{\eta,t-1}\eta_t, \quad \eta_t \sim N(0, 1), \tag{2}$$

$$\epsilon_t = \rho\epsilon_{t-1} + \xi_{v,t-1}v_t, \quad v_t \sim N(0, 1), \tag{3}$$

$$\ln \xi_{\eta,t}^2 = \ln \xi_{\eta,t-1}^2 + \sigma_\eta \phi_{\eta,t}, \quad \phi_{\eta,t} \sim N(0, 1), \tag{4}$$

$$\ln \xi_{v,t}^2 = \ln \xi_{v,t-1}^2 + \sigma_v \phi_{v,t}, \quad \phi_{v,t} \sim N(0, 1). \tag{5}$$

In sum,  $\tau_t$  is the slowly-evolving trend in inflation while  $\epsilon_t$  is the inflation gap, which is stationary, but can be persistent. Their innovations are  $\eta_t$  and  $v_t$ , respectively. These two innovations experience SV depicted with independent log random walks in  $\xi_{\eta,t}^2$  and  $\xi_{v,t}^2$ . We assume the innovations  $\eta_t$ ,  $v_t$ ,  $\phi_{\eta,t}$ , and  $\phi_{v,t}$  form a vector of martingale difference sequences on the filtration  $\mathcal{F}_{t-1}$ , given initial conditions  $\tau_0$ ,  $\epsilon_0$ ,  $\xi_{\eta,0}$ , and  $\xi_{v,0}$ , where  $t = 0, 1, 2, \dots, \infty$ .

This UC model predicts the long-run forecast of  $\pi_t$  is the Beveridge-Nelson (1981) trend. The long-run forecast is  $h$ -step ahead expected inflation, given date  $t$  information, at the infinite horizon,  $\lim_{h \rightarrow \infty} E_t \pi_{t+h} = \tau_t$ . Watson (1986) and Morley, Nelson, and Zivot (2004) show that this result follows from the fact that  $\tau_t$  evolves as a random walk. The prediction is robust to SV in the innovations,  $\eta_t$  and  $v_t$ , to  $\tau_t$  and  $\epsilon_t$  as long as the SVs are ‘predetermined’ or independent of these innovations. We refer to the UC model (1)–(5) as the BNSW model in recognition of these classic contributions.

### 2.1. Rational Expectations

This section describes how to learn about these components using realized inflation,  $\pi_t$ , and the mean  $h$ -quarter ahead inflation predictions from the SPF. Denote these predictions

$\pi_{t,h}^{SPF}$  at horizon  $h$ . Under rational expectations (RE), we assume  $\pi_{t,h}^{SPF}$  equals the  $h$ -quarter ahead RE inflation forecast,  $E_t\pi_{t+h}$ , plus a classical measurement error,  $\sigma_{\psi,h}\psi_{t,h}$ , where  $\psi_{t,h} \sim N(0, 1)$ .

Unbiasedness of professional forecasts constitutes indirect evidence in favor of this assumption. Keane and Runkle (1990) give early evidence of the unbiasedness of price forecasts using disaggregated data from the Livingston Survey. Ang, Bekaert, and Wei (2007) describe an inflation-forecasting tournament in which the median professional forecast is the best predictor of annual inflation. Gil-Alana, Moreno, and Pérez de Gracia (2012) find similarly favorable results for survey-based expectations of quarterly inflation and, specifically, the mean CPI inflation forecasts from the SPF. Croushore (2010) demonstrates the general lack of bias in the SPF forecasts using real-time measures of target variables. Overall, as Faust and Wright (2013, p. 5) note, “Subjective forecasts of inflation seem to outperform model-based forecasts in certain dimensions, often by a wide margin.” Winning tournaments based on mean-squared error, of course, does not imply unbiasedness, but it at least rules out systematic biases, for otherwise a time-series model would incorporate those and improve upon professional forecasts. To quote Faust and Wright again (p. 21), “A useful way of assessing models [thus] is by their ability to match survey measures of inflation expectations.”

The RE assumption gives us an easy way to link  $\pi_{t,h}^{SPF}$  to the hidden components of the BNSW model. The BNSW model predicts  $E_t\pi_{t+h} = \tau_t + \rho^h\epsilon_t$ . Substituting this expression and an implication of equation (1),  $\tau_t = \pi_t - \epsilon_t$ , into  $\pi_{t,h}^{SPF} = E_t\pi_{t+h} + \sigma_{\psi,h}\psi_{t,h}$  leads to the RE-SPF observation equation

$$\Pi_{t,h}^{SPF} \equiv \pi_{t,h}^{SPF} - \pi_t = (\rho^h - 1)\epsilon_t + \sigma_{\psi,h}\psi_{h,t}, \quad (6)$$

for  $h = 1, \dots, \mathcal{H}$ . Trend inflation is annihilated in the observation equation (6) because  $\pi_{t,h}^{SPF}$  and  $\pi_t$  share the permanent component of the BNSW model in common. We interpret  $\Pi_t^{SPF}$  as the accumulated growth of inflation  $h$ -quarters ahead that is anticipated by the SPF on average at quarter  $t$ .

The RE-SPF state space model consists of the AR(1) of  $\epsilon_t$ , the random walk of  $\ln \xi_{v,t}^2$ , and  $\mathcal{H}$  RE-SPF observation equations. The RE-SPF observation equation (6) is

formed into a stack from  $h = 1, \dots, \mathcal{H}$ , where the vector of dependent variables  $\Pi_t^{SPF} = [\Pi_{t,1}^{SPF} \dots \Pi_{t,\mathcal{H}}^{SPF}]'$ . The system of observation equations constitutes a RE term structure of inflation growth anticipated by the mean of the SPF. The RE term structure is explained by a single factor and classical measurement errors. The single factor is gap inflation,  $\epsilon_t$ . Thus, its AR(1) is a state equation. This is equation (3) of the BNSW model. Since  $\epsilon_t$  is affected by SV,  $\xi_{v,t-1}$ , through its impact on the innovation,  $v_t$ , to gap inflation, the random walk of its SV,  $\ln \xi_{v,t}^2$  is the second state equation. This is equation (5). The top of table 1 and section A.1.a of the appendix summarize the RE-SPF state space model.

## 2.2 Sticky Information

Notwithstanding citations to earlier research showing professional forecasts are unbiased, a number of recent, statistical studies have found forecast errors contain predictable components. Several of these studies argue that finding a specific pattern of predictability grounded in mean forecast revisions leads to an alternative, parametric model of observed, mean forecasts.

We work with the sticky-information (SI) model introduced by Mankiw and Reis (2002). The SI model is applied to professional forecasters by Mankiw, Reis, and Wolfers (2004) and Coibion and Gorodnichenko (2012, 2015). Suppose SI forecasters update their information with probability  $1 - \lambda$ ,  $\lambda \in (0, 1)$ . In this case,  $\lambda$  measures the degree of stickiness in information. Let  $F_t \pi_{t+h}$  denote the SI forecasts, which is the cross-forecaster mean forecast at quarter  $t$  for inflation  $h$ -quarters ahead. Coibion and Gorodnichenko (2015) show this average forecast is a weighted average of the previous period's mean, reported forecast and the  $h$ -quarter ahead RE inflation forecast:

$$F_t \pi_{t+h} = \lambda F_{t-1} \pi_{t+h} + (1 - \lambda) E_t \pi_{t+h}, \quad (7)$$

for  $h = 1, \dots, \mathcal{H}$ . We also note  $1/(1 - \lambda)$  is the frequency at which SI forecasts are updated on average.

The SI-BNSW state space model is built with the SI law of motion (7), the BNSW model, and an assumption that connects  $\pi_{t,h}^{SPF}$  to  $F_t \pi_{t+h}$ . The assumption drops the RE inflation forecast for the SI inflation forecast in  $\pi_{t,h}^{SPF} = F_t \pi_{t+h} + \sigma_{\psi,h} \psi_{t,h}$ . Combining the SI law of motion (7) and the RE inflation forecast produced by the BNSW model gives

$(1 - \lambda L)F_t \pi_{t+h} = (1 - \lambda)[\tau_t + \rho^h \epsilon_t]$ . Once again, eliminate  $\tau_t$  using equation (1) of the BNSW model to find  $(1 - \lambda L)F_t \pi_{t+h} = (1 - \lambda)[\pi_t + (\rho^h - 1)\epsilon_t]$ . Next, rearranging the previous expression and noticing the BNSW model implies  $\Delta \pi_t = \xi_{\eta,t-1} \eta_t + (\epsilon_t - \epsilon_{t-1})$  produces the SI-SPF observation equation:

$$\Pi_{t,h}^{SPF} = \lambda \Pi_{t-1,h}^{SPF} + [(1 - \lambda)\rho^h - 1]\epsilon_t + \lambda \epsilon_{t-1} - \lambda \xi_{\eta,t-1} \eta_t + \sigma_{\psi,h} \psi_{t,h} - \lambda \sigma_{\psi,h} \psi_{t-1,h}, \quad (8)$$

where  $\pi_{t,h}^{SPF} = F_t \pi_{t+h} + \sigma_{\psi,h} \psi_{t,h}$ , is used to eliminate the SI inflation forecasts  $F_t \pi_{t+h}$  and  $F_{t-1} \pi_{t+h}$ .

Compared with the RE-SPF observation equation (6), SI imposes costs on inflation forecast updating for the average SPF participant that appear as AR(1), MA(1), and SV elements in the SI-SPF observation equation (8). The AR(1) component is the pre-determined explanatory variable,  $\Pi_{t-1,h}^{SPF} \equiv \pi_{t-1,h}^{SPF} - \pi_{t-1}$ . There are also MA(1)s in gap inflation,  $[(1 - \lambda)\rho^h - 1](1 + \lambda[(1 - \lambda)\rho^h - 1]^{-1}L)\epsilon_t$ , and the measurement error,  $\sigma_{\psi,h}(1 - \lambda L)\psi_{t,h}$ . The AR(1) and MA(1) terms reflect the sluggishness that SI inflation forecast updating creates in the responses of  $\pi_{t,h}^{SPF}$  to shocks that cause it to deviate from  $\pi_t$ . Lagged SV enters the SI-SPF observation equation (8) for the same reason. The SI-SPF observation equation (8) also has a forward-looking element,  $(1 - \lambda)\rho^h \epsilon_t$ , tied to the RE inflation forecast produced by our BNSW model. The SI restrictions leave the RE-SPF observation equation (6) nested by the SI-SPF observation equation (8). The nesting relies only on  $\lambda = 0$ .

The SI-SPF state space model begins by creating a stack of the SI-SPF observation equation (8) from  $h = 1, \dots, \mathcal{H}$ , where the dependent variable is  $\Pi_t^{SPF}$ . However, the SI-term structure of  $\Pi_t^{SPF}$  is driven by several state variables. Three of the states are  $\epsilon_t$ ,  $\xi_{\eta,t}$ , and  $\xi_{v,t}$ . Thus, equations (3)–(5), which are the AR(1) of gap inflation and the SVs of trend and gap inflation, are part of the system of state equations of the SI-SPF state space model. Part of this system of state equations consists of laws of motion for  $\psi_{1,t}$ ,  $\psi_{2,t}$ , and  $\psi_{3,t}$ , to capture the impact of the MA(1) of these measurement errors on the term structure of average SPF anticipated  $h$ -step ahead accumulated inflation growth. The bottom panel of table 1 and section A.1.b of the appendix give details about the SI-SPF state space model.

### 3. Related Research

The BNSW decomposition has been fruitful in studies of several aspects of inflation dynamics. Ireland (2007) estimates the Federal Reserve’s implicit, time-varying inflation target with a Beveridge-Nelson trend. Cogley and Sbordone (2008) use a similar, stochastic trend around which to estimate a New Keynesian Phillips curve while Ascari and Sbordone (2013) survey this approach and outline its implications for monetary policy. Stock and Watson (2007) interpret the changing persistence and forecastability of US inflation with the UC model with changes in shock variances. Cogley, Primiceri, and Sargent (2010) use a related model to identify changes in the persistence of the inflation gap,  $\epsilon_t$ . Stock and Watson (2016) extend the UC model with stochastic volatility to model outliers and to jointly track the dynamics of the components of the PCE deflator.

Estimation and forecasting with the UC model require one to use the Kalman filter to extract the unobserved components. The filter is applied beginning with orthogonality assumptions (for example, a zero covariance between  $\eta_t$  and  $\epsilon_t$ ) and a set of covariates in observation equations. Examples of studies that apply the Kalman filter to this model include Stock and Watson (2007), Shephard (2013), Cogley and Sargent (2015), and Mertens (2016). Mertens documents the roles of a wide range of covariates, whereas we outsource much of the work to the participants in the SPF.

Our use of forecast surveys means that our study can be thought of as a sequel to that of Kozicki and Tinsley (2012), who estimate the parameters of a UC model of CPI inflation using actual inflation and the long span of observations from the Livingston survey, allowing for higher-order dynamics to fit seasonally unadjusted data, and under the assumption that forecasts are conditional expectations. They provide a detailed discussion of the interpretation and need for a shifting endpoint,  $\tau_t$ , for inflation forecasts. Henzel (2013) similarly combines SPF forecasts with the UC model to estimate inflation expectations. He also contrasts the speed of adjustment (or Kalman gain) in SPF forecasts with that estimated for the UC model alone. Jain (2019) applies a state-space model with a persistent but stationary unobserved component to the forecasts of individual forecasters in the *SPF*. She uses the properties of forecast revisions to deduce the persistence implied in these forecasters’ views of the underlying state variables and finds that this persistence has

declined over time for many forecasters.

We use the SPF with multiple horizons and quarterly observations but also allowing for SV, a feature Grassi and Proietti (2010) and Creal (2012) find to be important for modelling US inflation using the BNSW model. They find that the volatility of CPI inflation has increased recently, with the increased volatility attributed to the transitory rather than the permanent component. We also consider the possibility that forecasts are sticky. This extension potentially reconciles the method with the bias in mean SPF forecasts yet still allows us to use those forecasts to estimate the UC model and measure inflation expectations. We work with the SI model because of its tractability in this application and also because it has been used to close and estimate macroeconomic models, not just to describe forecasts. For example, Kiley (2007), and Coibion (2010) test versions of the SI model applied to price-setting and hence to aggregate inflation. Mertens and Nason (2019) study a more complex UC model with time-variation in gap-persistence  $\rho$  and stickiness  $\lambda$ . They detect and interpret variation in both parameters.

#### **4. A Particle MH-MCMC Algorithm**

We estimate the RE- and SI-BNSW state space models using Bayesian sequential Monte Carlo methods. These methods consist of a PMH-MCMC (particle Metropolis-Hastings, Markov chain Monte Carlo), which wraps a random walk MH-MCMC around a RB-APF (Rao-Blackwellized auxiliary particle filter). The RB-APF is the source of estimates of the state variables of the state space models while the random walk MH-MCMC produces posterior distributions of the fixed parameters. The RB procedure and APF are discussed by Chen and Liu (2000), Creal (2012), Shephard (2013), and Pitt and Shephard (1999, 2001). We implement a PMH-MCMC algorithm that relies on results in Andrieu, Doucet, and Holenstein (2010) and the advice of Pitt, dos Santos Silva, Giordani, and Kohn (2012), Doucet, Pitt, Deligiannidis, and Kohn (2015), and Martino, Elvira, and Camps-Valls (2018). Sections A.2 and A.3 of the appendix discuss in detail our Bayesian framework for estimating the state space models.

##### **4.1 The Auxiliary Particle Filter**

Pitt and Shephard (1999, 2001) develop the APF to achieve greater efficiency com-

pared with a bootstrap particle filter. A canonical bootstrap particle filter is a sequential importance sampler with a resampling step. Let  $\{\mathcal{S}_{t-1|t-1}^{(j)}\}_{j=1}^{\mathcal{J}}$  denote  $\mathcal{J}$  date  $t-1$  filtered particles of the state vector of an arbitrary state space model. The resampling step draws  $\mathcal{J}$  particles from  $\{\mathcal{S}_{t-1|t-1}^{(j)}\}_{j=1}^{\mathcal{J}}$  with replacement. These draws rely on  $\mathcal{J}$  particle weights that are the normalized likelihood particles of the state space model. Thus, this step copies into the ensemble of resampled states,  $\{\tilde{\mathcal{S}}_{t-1|t-1}^{(j)}\}_{j=1}^{\mathcal{J}}$ , the particles of  $\{\mathcal{S}_{t-1|t-1}^{(j)}\}_{j=1}^{\mathcal{J}}$  associated with the largest likelihoods. Resampling aims to rule out all the weight of the normalized likelihoods falling on a few or even one particle. As Creal (2012) notes, this solves the problem of the degeneracy of particles that are used to estimate states at future dates. Finally, using the state dynamics, the stream  $\{\tilde{\mathcal{S}}_{t-1|t-1}^{(j)}\}_{j=1}^{\mathcal{J}}$  is propagated into date  $t$ , which yields  $\{\mathcal{S}_{t|t}^{(j)}\}_{j=1}^{\mathcal{J}}$ .

The efficiency gains of the APF come in part from it reversing the order of the resampling and propagation steps. Since the APF is a propagation-resampling algorithm, the particle weights used to resample the date  $t-1$  states are built on sample information available at date  $t$ . This results in two ensembles of particles. If the stream of resampled particles,  $\{\tilde{\mathcal{S}}_{t-1|t-1}^{(j)}\}_{j=1}^{\mathcal{J}}$  produce a distribution of likelihoods with fatter tails compared with the likelihoods conditioned on the prior ensemble of particles,  $\{\mathcal{S}_{t-1|t-1}^{(j)}\}_{j=1}^{\mathcal{J}}$ , the APF is more efficient; see Johansen and Doucet (2008) and Creal (2012).

## 4.2 Rao-Blackwellization and the APF

Rao-Blackwellization boosts the efficiency of the APF when applied to a nonlinear state space model. The improvement in efficiency rests on a state space model having states that are linear and Gaussian conditional on its nonlinear states. For example, the top panel of table 1 shows that gap inflation is the lone conditionally linear and Gaussian state,  $\mathcal{S}_{RE,t}$ , of the RE-SPF state space model. We label this state space model  $\mathcal{M}_{RE}$ . Its single nonlinear state is the SV of gap inflation,  $\xi_{v,t}$  as pointed out in table 1. Since  $\epsilon_t$  evolves according to the AR(1) of equation (3), it produces estimates of  $\mathcal{S}_{RE,t}$  that are linear and Gaussian conditional on a lag of gap inflation and on realizations of  $\xi_{v,t-1}$ . Realizations of  $\xi_{v,t}$  are retrieved from its random walk, which is equation (5).

These states are also part of the SI-SPF state space model that is referred to as  $\mathcal{M}_{SI}$ . However, as discussed in the lower panel of table 1,  $\mathcal{M}_{SI}$  has more elements in its

conditional linear and Gaussian state vector,  $\mathcal{S}_{SI,t}$ , than gap inflation. Six of the additional states are linear in the measurement errors  $\psi_{1,t}$ ,  $\psi_{2,t}$ , and  $\psi_{3,t}$ . Also, included in  $\mathcal{S}_{SI,t}$  is the fixed SI parameter,  $\lambda$ , that is in effect a ‘pseudo’ state. Table 1 describes a system of conditional linear and Gaussian state equations for  $\mathcal{M}_{SI}$ , given realizations of  $\xi_{v,t-1}$ . Similarly, forecasts of  $\Pi_{t,h}$  are produced by the linear and Gaussian observation equation (8) conditional on a lag of the second nonlinear state of  $\mathcal{M}_{SI}$ ,  $\xi_{\eta,t}$ .

A Rao-Blackwellized APF applied to  $\mathcal{M}_m$  employs the Kalman filter to construct an analytic distribution of  $\mathcal{S}_{m,t}$ ,  $m = RE, SI$ . Since these state space models are linear and Gaussian, running the Kalman filter produces  $\mathcal{S}_{m,t} \sim N(\mathcal{S}_{m,t|t}^{(j)}, \Sigma_{m,t|t}^{(j)})$ , where  $\mathcal{S}_{m,t|t}^{(j)}$  is the  $j$ th particle of the filtered state(s) at date  $t$  and its mean square error (MSE) is  $\Sigma_{m,t|t}^{(j)}$  for  $j = 1, \dots, \mathcal{J}$ . Drawing estimates of  $\mathcal{S}_{m,t}$  analytically makes the Rao-Blackwellized (RB-)APF more efficient; see Chen and Liu (2000), Creal (2012), and Särkkä (2013). The distribution of  $\xi_{v,t}$  is approximated by simulation in  $\mathcal{M}_{RE}$ . For  $\mathcal{M}_{SI}$ , simulation is also used to approximate the distribution of  $\xi_{\eta,t}$ . Section A.2.a of the appendix has more about RB of the APF.

We build a RB-APF by combining algorithm 3 of Creal (2012) and algorithm 1 of Pitt, dos Santos Silva, Giordani, and Kohn (2012). We sketch the RB-APF as it is applied to  $\mathcal{M}_{RE}$  in the rest of this section while leaving details about implementing the algorithm for  $\mathcal{M}_{SI}$  to section A.2.b of the appendix.

Initializing the RB-APF for  $\mathcal{M}_{RE}$  relies on priors for  $\mathcal{S}_{RE,0}$  and  $\ln \xi_{v,0}^2$  to draw  $\mathcal{J}$  particles at date 0 and on fixing the second stage weight  $\varpi_{SI,0}^{(j)} = \mathcal{J}^{-1}$  for  $j = 1, \dots, \mathcal{J}$ . Given these initial conditions,  $\Pi_1^{SPF}$ , and the parameter vector of  $\mathcal{M}_{RE}$ ,  $\Theta_{RE}$ , at date 1 engage the Kalman filter prediction formulas to compute the stream of first-stage predictive likelihoods,  $\{\exp\{l_{RE,1}^{(j)}\}\}_{j=1}^{\mathcal{J}}$  and  $j$ th first-stage weight  $\varpi_{RE,0|1}^{(j)} = \omega_{RE,0|1}^{(j)} / \sum_{j=1}^{\mathcal{J}} \omega_{RE,0|1}^{(j)}$ , which incorporates information in  $\Pi_1^{SPF}$ , where  $\omega_{RE,0|1}^{(j)} = (\varpi_{RE,0}^{(j)} / \sum_{j=1}^{\mathcal{J}} \varpi_{RE,0}^{(j)}) \exp\{l_{RE,1}^{(j)}\}$ . We input  $\{\varpi_{SI,0|1}^{(j)}\}_{j=1}^{\mathcal{J}}$  into a stratified resampling (with replacement) scheme described by Hol, Schön, and Gustafsson (2006) and Li, Bolic, and Djuric (2015) to obtain indexes that applied to  $\{\mathcal{S}_{RE,0|0}^{(j)}, \Sigma_{RE,0|0}^{(j)}, \xi_{\eta,0}^{(j)}\}_{j=1}^{\mathcal{J}}$  yield resampled linear and nonlinear states,  $\{\tilde{\mathcal{S}}_{RE,0|0}^{(j)}, \tilde{\Sigma}_{RE,0|0}^{(j)}, \tilde{\xi}_{\eta,0}^{(j)}\}_{j=1}^{\mathcal{J}}$ . In a second stage, pass  $\{\tilde{\mathcal{S}}_{RE,0|0}^{(j)}, \tilde{\Sigma}_{RE,0|0}^{(j)}, \tilde{\xi}_{v,0}^{(j)}\}_{j=1}^{\mathcal{J}}$  through the Kalman filter predictive and updating formulas to create  $\{\mathcal{S}_{RE,1|1}^{(j)}, \Sigma_{RE,1|1}^{(j)}\}_{j=1}^{\mathcal{J}}$ . These

operations also produce the second-stage predictive log likelihood,  $l_{RE,1}^{(j)}$  that is used to update the second-stage weights  $\varpi_{RE,1}^{(j)} = \omega_{RE,1}^{(j)} / \sum_{j=1}^{\mathcal{J}} \omega_{RE,1}^{(j)}$  for  $j = 1, \dots, \mathcal{J}$ , where  $\omega_{RE,1}^{(j)} = \exp\{l_{RE,1}^{(j)} - \tilde{l}_{RE,1}^{(j)}\}$ , and  $\tilde{l}_{RE,1}^{(j)}$  is resampled using  $\{\varpi_{RE,0|1}^{(j)}\}_{j=1}^{\mathcal{J}}$ . The last step propagates the SV of gap inflation,  $\{\ln[\xi_{v,1}^{(j)}]^2\}_{j=1}^{\mathcal{J}}$ , by entering  $\{\ln[\tilde{\xi}_{v,0}^{(j)}]^2\}_{j=1}^{\mathcal{J}}$  and  $\mathcal{J}$  draws from  $\phi_{\nu,1} \sim N(0, 1)$  into the random walk (5).

The last three steps gives us priors and second-stage weights to operate the RB-APF at date 2. Subsequently, we run the RB-APF at dates  $t = 3, \dots, T$ . Operating the RB-APF on the entire sample creates  $\{\{\mathcal{S}_{RE,t|t}^{(j)}, \Sigma_{RE,t|t}^{(j)}, \xi_{v,t}^{(j)}\}_{j=1}^{\mathcal{J}}\}_{t=1}^T$  and  $\{\{\varpi_{RE,t}^{(j)}\}_{j=1}^{\mathcal{J}}\}_{t=1}^T$ . These particle streams produce the filtered means  $\mathcal{S}_{RE,t|t} = \sum_{j=1}^{\mathcal{J}} \varpi_{RE,t}^{(j)} \mathcal{S}_{RE,t|t}^{(j)}$ ,  $\Sigma_{RE,t|t} = \sum_{j=1}^{\mathcal{J}} \varpi_{RE,t}^{(j)} \Sigma_{RE,t|t}^{(j)}$ , and  $\xi_{v,t} = \sum_{j=1}^{\mathcal{J}} \varpi_{RE,t}^{(j)} \xi_{v,t}^{(j)}$ .

The RB-APF algorithm also provides an estimate of the likelihood of  $\mathcal{M}_{RE}$ . Pitt, dos Santos Silva, Giordani, and Kohn (2012) recommend estimating the likelihood at date  $t$  with:

$$\mathcal{P}(\Pi_t^{SPF} | \Pi_{1:t-1}^{SPF}; \mathcal{M}_{RE}, \Theta_{RE}) = \left( \frac{1}{\mathcal{J}} \sum_{j=1}^{\mathcal{J}} \omega_{RE,t}^{(j)} \right) \sum_{j=1}^{\mathcal{J}} \omega_{RE,t-1|t}^{(j)}. \quad (9)$$

We generate  $\mathcal{P}(\Pi_t^{SPF} | \Pi_{1:t-1}^{SPF}; \mathcal{M}_{RE}, \Theta_{RE})$  from the start to the end of the sample using equation (9).

### 4.3 A Random Walk PMH-MCMC

A necessary input into a PMH-MCMC sampler is an unbiased estimate of the likelihood. Equation (9) supplies unbiased estimates of  $\mathcal{P}(\Pi_t^{SPF} | \Pi_{1:t-1}^{SPF}; \mathcal{M}_m, \Theta_m)$  for  $m = RE, SI$  at  $t = 1, \dots, T$ . However, this leaves open the question of whether errors in these estimates of the likelihood impair posterior distributions of  $\mathcal{M}_{RE}$  and  $\mathcal{M}_{SI}$ . We depend on Andrieu, Doucet, and Holenstein (2010) for proof the posterior distribution  $\mathcal{P}(\mathcal{S}_{m,t}, \xi_{v,t}, \Theta_m | \Pi_{1:T}^{SPF}, \mathcal{M}_m)$  is unaffected by estimation error that results from computing the likelihood (9) assuming this estimate is unbiased. They also lay out the conditions under which the RB-APF produces unbiased estimates of  $L(\Pi_{1:T}^{SPF} | \Theta_m, \mathcal{M}_m) = \exp[\sum_{t=1}^T \ln \mathcal{P}(\Pi_t^{SPF} | \Pi_{1:t-1}^{SPF}; \mathcal{M}_m, \Theta_m)]$  for  $m = RE, SI$ ; also see Creal (2012) and Pitt, et al (2012).

These results tells us to generate posterior distributions of  $\Theta_{RE}$  and  $\Theta_{SI}$  by placing the RB-APF inside a random walk MH-MCMC simulator. Our PMH-MCMC sampler employs

a multivariate random walk proposal with standard normal innovations. The scale matrix of the random walk proposal is calculated using the robust adaptive Metropolis (RAM) algorithm of Vihola (2012).

We start the PMH-MCMC sampler with an initial draw,  $\widehat{\Theta}_{m,0}$ , of the parameter vector of  $\mathcal{M}_m$ ,  $m = RE, SI$ . At the initial step of the PMH-MCMC sampler, the multivariate MH random walk is  $\Theta_{m,1} = \widehat{\Theta}_{m,0} + \Gamma_{\Theta_{m,0}}^{1/2} \vartheta_{m,1}$ ,  $\vartheta_{m,1} \sim N(0_{d_m \times 1}, I_{d_m})$ , where  $\Gamma_{\Theta_{m,0}}^{1/2}$  is the Cholesky decomposition of  $\Gamma_{\Theta_{m,0}}$ , which is the initial covariance matrix of  $\Theta_m$ , and  $d_m = \dim(\Theta_m)$ . This covariance matrix is drawn from an inverse Wishart distribution,  $\Gamma_{\Theta_{m,0}} = \left(2.4/\sqrt{d_m}\right)^2 \underline{\Gamma}_{\Theta_m}$ , where  $\underline{\Gamma}_{\Theta_m} \sim \mathcal{IW}(100.0, 0.01I_{d_m})$ . Next, set the counter  $r$  to zero and draw a uniform random variable  $\nu_1 \sim \mathcal{U}(0, 1)$  variable to evaluate the MH criterion:

$$\alpha_1 = \min \left\{ \frac{\widehat{L}(\Pi_{1:T}^{SPF} | \Theta_{m,1}, \mathcal{M}_m) q(\Theta_{m,0}, \Theta_{m,1}) \mathcal{P}(\Theta_{m,1})}{\widehat{L}(\Pi_{1:T}^{SPF} | \Theta_{m,0}, \mathcal{M}_m) q(\Theta_{m,1}, \Theta_{m,0}) \mathcal{P}(\Theta_{m,0})}, 1 \right\},$$

where the RB-APF is the source of  $\widehat{L}(\Pi_{1:T}^{SPF} | \Theta_{m,1}, \mathcal{M}_m)$ ,  $q(\cdot, \cdot)$  is the kernel of the proposal distribution, and  $\mathcal{P}(\Theta_{m,1})$  and  $\mathcal{P}(\widehat{\Theta}_{m,0})$  are priors conditional on  $\Theta_{m,1}$  and  $\widehat{\Theta}_{m,0}$ . If  $\nu_1 \leq \alpha_1$ , set  $\widehat{\Theta}_{m,1} = \Theta_{m,1}$  and  $r = 1$ . Otherwise,  $\widehat{\Theta}_{m,1} = \widehat{\Theta}_{m,0}$  and  $r = 0$ . When the support of the proposal and the target differ, the kernel of the proposal distribution is not symmetric. In this case, we calculate  $q(\cdot, \cdot)$  as Lindström (2017) suggests. Section A.2.d of the appendix has further details.

The PMH-MCMC sampler repeats this process for  $k = 2, \dots, \mathcal{K}$  with one exception. The scale or Cholesky matrix of the covariance matrix of  $\Theta_m$  in the multivariate MH random walk:

$$\Theta_{m,k} = \widehat{\Theta}_{m,k-1} + \Gamma_{\Theta_{m,k-1}}^{1/2} \vartheta_{m,k}, \quad \vartheta_{m,k} \sim N(0_{d_m \times 1}, I_{d_m}), \quad (10)$$

is computed using the RAM updating algorithm

$$\Gamma_{\Theta_{m,k}} = \Gamma_{\Theta_{m,k-1}}^{1/2} \Gamma_{\Theta_{m,k-1}}^{1/2'} + \Gamma_{\Theta_{m,k-1}}^{1/2} \left( \min((1, d_m k^\iota) \times (\alpha_k - \alpha^*) \frac{\vartheta_k \vartheta_{k'}}{\|\vartheta_k\|^2} \right) \Gamma_{\Theta_{m,k-1}}^{1/2}, \quad (11)$$

where  $\iota$  is the step size to adapt new proposals and  $\alpha^*$  is the desired acceptance rate of the PMH-MCMC. We follow the advice of Vihola (2012) to set  $\iota = -0.65$ . Our goal is

to achieve the optimal acceptance rate of a MH-MCMC with a multivariate target, which requires  $\alpha^* = 0.234$ . Although neither the empirical distribution (*i.e.*, outer product) of  $\Theta_m$  nor the Hessian of  $\mathcal{M}_m$  is used to estimate  $\Gamma_{\Theta, m, k}$ , it is guaranteed to be positive definite and is a function of  $\alpha^*$ . Section A.3 of the appendix has more information about our implementation of the RAM algorithm while section A.2.d does the same for the MH-MCMC sampler.

#### 4.4 Choosing the Number of Particles and MCMC Steps

The PMH-MCMC encloses the RB-APF inside the MH-MCMC sampler. We need to choose  $\mathcal{J}$  to generate the likelihood using the RB-APF at every one of the  $\mathcal{K}$  steps of the MH-MCMC to produce the posterior distributions of the states and parameters. This can be computationally costly. We adopt several strategies to reduce these costs.

Our template for setting  $\mathcal{J}$  is motivated by theory developed by Pitt et al (2012) and a stopping rule suggested by Doucet, Pitt, Deligiannidis, and Kohn (2015). Pitt et al build a theory to find the optimal number of particles by equating the error variance of the estimated log likelihood to the theoretical optimum. The problem for us is their theory relies on assuming the proposal of the PMH-MCMC is ideal (*i.e.*, is an exact match for the posterior distribution). Moving off of this assumption lets Doucet et al produce a theoretical upper bound on computing time. Doucet et al ground a rule for choosing  $\mathcal{J}$  that trades-off computing time and the (error) variance of the estimated log likelihood of a particle filter. The rule says to stop increasing  $\mathcal{J}$  if the error variance of the estimated log likelihood  $\hat{\sigma}_{L, m}^2$  is  $1.2^2$ .

A problem is an estimate of the posterior mean,  $\bar{\Theta}_m$ , is needed to operate the stopping rule. Pitt et al propose to estimate  $\bar{\Theta}_m$  by running the PMH-MCMC on a large  $\mathcal{J}$  and a small number of MH-MCMC steps. Our approach inputs  $\mathcal{J} = 100,000$  particles into the RB-APF while running the MH-MCMC for 30,000 steps. The first 10,000 steps serves as a burn-in sample that is discarded. The remaining 20,000 steps yield an estimate of the posterior mean,  $\bar{\Theta}_m$  for  $\mathcal{M}_m$ ,  $m = RE, SI$ , on the SPF-CPI and SPF-PGNP/PGDP inflation samples. The distributed parallel PMH-MCMC algorithm of Martino, Elvira, and Camps-Valls (2018) is used in this case. Details are found in section A.3 of the appendix.

We run the RB-APF beginning at  $\mathcal{J} = 26$  particles to estimate  $\hat{L}(\Pi_{1:T}^{SPF} | \Theta_m, \mathcal{M}_m)$ ,

$m = RE, SI$ . If  $\hat{\sigma}_{L,m}^2 \leq 1.2^2$ , stop. Otherwise, increase  $\mathcal{J}$  by one particle and repeat. On the CPI-SPF sample, the stopping rule yields  $\mathcal{J} = 50$  for  $\mathcal{M}_{RE}$  and  $\mathcal{J} = 112$  for  $\mathcal{M}_{SI}$ . The numbers of particles increase to  $\mathcal{J} = 63$  and 171 for  $\mathcal{M}_{RE}$  and  $\mathcal{M}_{SI}$  on the PGNP/PGDP-SPF sample.

The PMH-MCMC is run on a single thread for  $\mathcal{M}_{RE}$  and  $\mathcal{M}_{SI}$  using these values of  $\mathcal{J}$ . On a single thread, the PMH-MCMC is run on 2,250,000 steps. The first third of these steps are discarded in a burn-in. The remaining 1,500,000 draws are employed to assess the convergence of the posterior distributions.

There is substantial correlation in the posterior distributions of  $\bar{\Theta}_{RE}$  and  $\bar{\Theta}_{SI}$  on  $\mathcal{M}_{RE}$  and  $\mathcal{M}_{SI}$  and the two samples. We engage the statistically efficient algorithm of Owen (2017) to thin these posterior distributions. His algorithm trades-off the cost of updating the PMH-MCMC against the cost of obtaining an estimate of the likelihood. Since the maximal first-order autocorrelation coefficient is 0.963 across the posterior distributions of  $\bar{\Theta}_{RE}$  and  $\bar{\Theta}_{SI}$  and we assume the cost of updating the PMH-MCMC sampler is 33 percent of the cost of estimating the likelihood, Owen’s algorithm produces a factor of 12 to thin the posterior distributions leaving 125,000 draws in the posterior. The thinned posterior distributions have a maximal first-order auto-correlation coefficient less than 0.642. The next sections reports estimates of  $\mathcal{M}_{RE}$  and  $\mathcal{M}_{SI}$  on the SPF-CPI inflation and SPF-PGDP/GNP inflation samples using these thinned posterior distributions.

## 5. Estimating the RE- and SI-BNSW State Space Models

This section describes the CPI-SPF and PGNP/PGDP-SPF samples and the prior distributions.

### 5.1 The Data

The RE- and SI-BNSW state space models are estimated on two samples. One sample consists of realized CPI inflation and associated average SPF inflation predictions. The CPI-SPF samples runs from 1981Q4 to 2018Q4. The other sample is realized GNP/GDP deflator inflation and related average SPF inflation predictions. This sample begins in 1969Q1 and ends with 2018Q4. In both samples, the SPF’s fifth definition of real time inflation is realized inflation,  $\pi_t$ . The average SPF predictions are a 0-quarter ahead

nowcast and 1-, . . . , 4-quarter ahead horizons. The nowcast is dropped because of potential timing conflicts with  $\pi_t$ . We also do not use  $\pi_{t,4}^{SPF}$  in the dependent variable of the RE- and SI-BNSW state space models. The reason is  $\Pi_{t-1,h}^{SPF}$  is a predetermined variable in the SI-BNSW observation equation (8). Since the SPF does not compile 5-quarter ahead inflation predictions necessary for  $h = 4$ , this bounds our RE- and SI-term structures of average SPF anticipated accumulated inflation growth to  $\mathcal{H} = 3$ . However, a lag of the 4-quarter ahead average SPF inflation predictions does appear in the dependent variable  $\Pi_{t-1,3}^{SPF}$  of the SI-BNSW observation equation (8) for  $h = 3$ . More information about the data and links to the data sources are in section A.2.a of the appendix.

Figure 1 plots the CPI-SPF data on the 1981 Q4–2018 Q4 sample. The top left window displays realized inflation,  $\pi_t$ . Average SPF inflation predictions appear in the top right window for  $h = 1, 2$ , and 3. These are for common date of origin (rather than for a common target). As the horizon rises, the volatility of the forecast decreases. The bottom window presents  $\Pi_{t,1}$ ,  $\Pi_{t,2}$ , and  $\Pi_{t,3}$ , which appear in the observations equations (6) and (8). The CPI-term structure of anticipated accumulated growth of inflation by the average SPF shows troughs during the 1981 Q3–1982 Q4, 1990 Q3–1991 Q1, 2001 Q1–2001 Q4, and 2007 Q4–2009 Q2 recessions as dated by the NBER. At the end of all but the last recession,  $\Pi_{t,1}$ ,  $\Pi_{t,2}$ , and  $\Pi_{t,3}$  peak. During the 2007 Q4–2009 Q2 recession, the peak occurs in 2008 Q4. Note this peak matches the trough in  $\pi_t$  at the same date. The CPI-term structure of anticipated accumulated growth of inflation by the average SPF is often greater than zero between the first two recessions, the second and third recession, and subsequent to the last recession, but this term structure is often less than zero from 2002 Q1 to 2007 Q3.

Forecasts for inflation are available for a longer time span, beginning in 1968:4, if we study the inflation rate in the GDP deflator rather than the CPI. These forecasts are for seasonally adjusted levels of the deflator, defined as (a) the GNP deflator prior to 1992, (b) the GDP deflator from 1992 to 1995, and (c) the chain-weighted price index for GDP from 1996 to the present. Then implicit mean forecasts for the annualized growth rate in the deflator are from `mean_PGDP_Growth.xls`.

Plots of the PGNP/PGDP-SPF data on the 1969 Q1–2018 Q4 sample are found in figure 3. The span of years, of course, now includes the high-inflation years of the 1970s.

The structure of the plots matches that in figure 1. There is less volatility and less evidence of a business-cycle pattern in  $\Pi_{t,1}$ ,  $\Pi_{t,2}$ , and  $\Pi_{t,3}$  for this measure of inflation than for CPI inflation in figure 1. However, the term structure of PGNP/PGDP anticipated accumulated growth of inflation exhibits a sharp trough in the 1973 *Q4*–1975 *Q1* recession. As was true for CPI inflation, there is a peak late in the 2007 *Q4*–2009 *Q2* recession, though in this case it occurs exactly at 2009 *Q4*. Again this peak matches the trough in  $\pi_t$  at the same date. The PGNP/PGDP-term structure of anticipated accumulated growth of inflation by the average SPF is often negative in the 1970s, then positive in the 1980s and 1990s.

The SPF also contains data on long-term inflation forecasts, specifically over the next year and the next ten years. The one-year forecast is the average of the median forecasts for  $h = 1$  to  $h = 4$ . The ten-year forecast is the annual average inflation rate predicted for this period. However, this survey information has been collected only since 1991. Forecasts for the PCE deflator have been collected only since 2007.

## 5.2 Priors for the RE- and SI-SPF State Space Models

The parameter vector of  $\mathcal{M}_{RE}$  is  $\Theta_{RE} = [\rho \ \sigma_v \ \sigma_{\psi,1} \ \sigma_{\psi,2} \ \sigma_{\psi,3}]'$ . The parameter vector  $\Theta_{SI} = [\Theta'_{RE} \ \sigma_\eta \ \lambda]'$  is tied to  $\mathcal{M}_{SI}$ . Table 2 lists our priors for the elements of  $\Theta_{RE}$  and  $\Theta_{SI}$ . The PMH-MCMC simulator draws from these priors to generate posterior distributions of  $\mathcal{M}_{RE}$  and  $\mathcal{M}_{SI}$ .

We select a truncated-normal ( $\mathcal{TN}$ ) prior for  $\rho$  to reflect our uncertainty about inflation gap persistence. The AR1 coefficient of  $\epsilon_t$  is drawn from a standard normal distribution with support restricted to  $\rho \in (-1, 1)$ . This prior yields five and 95 percent prior quantiles of  $-0.87$  and  $0.87$ .

The prior on  $\sigma_v$  is a  $\chi(3)$  distribution with location and scale parameters of zero and 0.2. We invoke the same prior for  $\sigma_\eta$  in  $\Theta_{SI}$ . Table 2 shows these priors share a median of 0.31. The 5 and 95 percent quantiles are 0.12 and 0.56, which contain the standard calibration that sets  $\sigma_\eta$  and  $\sigma_v$  to  $\sqrt{0.2}$ ; see Stock and Watson (2007) and Creal (2012). Shephard (2013) places a prior on these parameters in his version of the BNSW model that is similar to ours. As is his aim, our intent is to squash incredible draws of  $\sigma_\eta$  and  $\sigma_v$ .

We select the inverse-gamma ( $\mathcal{IG}$ ) distribution for the prior on the fixed volatility

parameters,  $\sigma_{\psi,1}$ ,  $\sigma_{\psi,2}$ , and  $\sigma_{\psi,3}$ , on the measurement errors  $\psi_{1,t}$ ,  $\psi_{2,t}$ , and  $\psi_{3,t}$ . Table 2 shows the dependence of the shape and scale coefficients,  $0.1T$  and  $0.045T$ , of these  $\mathcal{IG}$  priors on the sizes of the CPI-SPF and PGNP/PGDP-SPF inflation samples,  $T = 149$  and  $200$ . Our priors imply a median of  $0.45$  and  $0.46$  for the CPI-SPF and PGNP/PGDP-SPF inflation samples. Prior uncertainty is greater for the shorter CPI-SPF inflation sample compared with the PGNP/PGDP-SPF inflation sample as shown by 5 and 95 percent quantiles of  $[0.31, 0.74]$  and  $[0.32, 0.68]$  in table 2. A similar  $\mathcal{IG}$  prior on volatility parameters is used by Berger, Everaert, and Hauke (2016).

The SI parameter  $\lambda$  is given a  $\mathcal{TN}$  prior. Table 2 displays a prior mean equal to  $0.5$  with a unit variance. The prior median of  $\lambda$  is  $0.44$ , which approximates the estimates reported by Coibion and Gorodnichenko (2015), and its 5 and 95 percent quantiles are  $0.05$  and  $0.95$ . These quantiles indicate the prior frequency of SI updating of inflation forecasts runs from one to 20 quarters on average. Hence, our prior on  $\lambda$  is weakly informative on its support of  $\lambda \in (0, 1)$ .

Priors also are needed to initialize gap inflation SV,  $\ln \xi_{v,0}^2$  and trend inflation SV,  $\ln \xi_{\eta,0}^2$ . We endow  $\ln \xi_{v,0}^2$  and  $\ln \xi_{\eta,0}^2$  with log normal ( $\mathcal{LN}$ ) priors. The prior mean equals the log of the square of the share of the pre-sample variance of the first difference of realized CPI or PGNP/PGDP inflation attributed to trend or gap inflation minus a half. The pre-sample covers the 40 quarters before the start of either sample. Table 2 lists the prior means of  $\ln \xi_{v,0}^2$  and  $\ln \xi_{\eta,0}^2$  for the CPI and PGNP/PGDP samples. The prior means are larger for the CPI sample or  $\ln \xi_{v,0}^2$ . We set unit variances for these priors. These parameterizations suggest the priors are uninformative because the 5, 50, and 95 percent quantiles are  $[0.21, 1.08, 5.62]$ ,  $[0.06, 0.32, 1.66]$ ,  $[0.06, 0.31, 2.23]$ , and  $[0.02, 0.09, 0.48]$  for  $\ln \xi_{v,0}^2$  and  $\ln \xi_{\eta,0}^2$  on the CPI and PGNP/PGDP samples, respectively.

## 6. Findings

We next describe the findings, in two formats. Tables 3 and 4 describe posterior distributions of parameters (including  $\rho$ , describing persistence, and  $\lambda$ , describing stickiness) and of innovations to trend inflation. Then figures 2 and 4 present the posterior medians of the unobserved states at each date:  $\tau_{t|t}$ ,  $\epsilon_{t|t}$ ,  $\xi_{\eta,t|t}$ , and  $\xi_{v,t|t}$ .

## 6.1 Posterior Distributions of Parameters

There are four key features of the posterior densities for the model parameters, which are shown in table 3. First, inflation-gap persistence is quite modest. The posterior median value of  $\hat{\rho}$  ranges from 0.099 to 0.254 depending on the price measure (and hence time period) and RE or SI version of the model. And the 95% credible sets are quite narrow and quite different from the prior sets in table 2 (which have a median of 0), so that this parameter is well-identified in each case.

Second,  $\hat{\sigma}_\eta$  and  $\hat{\sigma}_v$ —the standard deviations of innovations to the volatilities in trend and gap inflation respectively—are also well-identified. For CPI inflation these two standard deviations are of comparable scale, while for GNP/GDP deflator inflation the posterior median of  $\hat{\sigma}_\eta$  is roughly 2.5 times larger than that of  $\hat{\sigma}_v$ . However, for both measures of inflation both SV terms are clearly necessary to reconstruct the history of US inflation. Table 3 also shows the standard deviations of measurement error ( $\hat{\sigma}_{\psi,1}$ ,  $\hat{\sigma}_{\psi,2}$ , and  $\hat{\sigma}_{\psi,3}$ ) which are similar across horizons.

Third, the stickiness parameter  $\hat{\lambda}$  has a posterior median of 0.438 (with a 95% credible set [0.334, 0.548]) for the CPI inflation rate and a posterior median of 0.366 (with a 95% credible set of [0.290, 0.450]) for the GDP deflator inflation rate. The prior median in table 2 is 0.44. Thus for the CPI sample the posterior median barely moves from the prior. In contrast, for the longer GDP deflator sample both the median and the 95% credible set are shifted down relative to the prior median. This shift does not occur if we also limit the GDP deflator sample to the period after 1981, so it is an effect of the 1970s not of the measure of inflation. Intuitively, the surge in inflation in the 1970s followed by the disinflation of the 1980s identifies  $\lambda$ , as observed forecasts are slow to catch up to realized inflation. This pattern is evident in the lower panel of figure 3, where  $\Pi_{t,h}^{SPF} \equiv \pi_{t,h}^{SPF} - \pi_t$  is negative for much of the 1970s then positive for the early 1980s.

Fourth, then, there is a noteworthy difference between the findings for the CPI inflation rate and those for the GDP deflator inflation rate, as shown in the last row of table 3. For the CPI inflation rate the log marginal data densities favour the RE model over the SI one, while the reverse is true for the GDP deflator inflation rate. These last two observations show that data from the 1970s are highly informative and important for distinguishing

between the RE and SI models. The rest of our reporting thus focuses on the RE model for the CPI and the SI model for the GDP deflator.

## 6.2 Inflation Components

To study the slowly evolving trend, table 4 presents posterior moments of filtered, trend inflation  $\eta_{t|t} = (\tau_{t|t} - \tau_{t-1|t-1}) / \xi_{\eta,t-1}$ . The first row gives the median of the posterior standard deviation of  $\eta_{t|t}$ . The SI model finds larger values for this volatility measure than the RE models do.

The UC model assumes that  $\eta_t$  is a martingale difference series. Our approach allows for tests of this assumption, in the form of tests for persistence in the extracted measure  $\eta_{t|t}$ . Table 4 presents posterior median autocorrelations and mean Ljung-Box statistics (along with their 95% credible sets). For the most part, these show little remaining persistence. However, there is some evidence of higher-order persistence in the innovations to trend inflation measured by the GDP deflator.

For CPI inflation figure 2 shows posterior medians for the unobserved states at each date, along with 68% uncertainty bands. (For comparison, a similar format is used by Cogley and Sargent, 2015.) The first panel shows realized CPI inflation along with the slowly evolving trend  $\tau_{t|t}$ . The trend component converges towards 2% and hovers around that value after 2009. But 2% lies in the 68% band for most quarters since 2001. The median trend also is little affected by recessions. The second panel shows the posterior median of the inflation gap  $\epsilon_{t|t}$ . Its low estimated persistence is evident from the plot. Gap inflation falls sharply part way through each recession.

The third panel shows the stochastic volatility in trend inflation,  $\xi_{\eta,t|t}$ . There is a clear moderation in trend inflation over time, as the posterior median of  $\xi_{\eta,t|t}$  falls over time. Cecchetti et al (2017) model several measures of US inflation using a UC model with stochastic volatility (though not using survey data) and also note a downward trend in the innovation variance of  $\tau_t$ . Because  $\tau_t$  coincides with long-run inflation expectations under the BNSW model, the combination of the convergence of  $\tau_{t|t}$  to 2% and the decline in  $\xi_{\eta,t|t}$  can be seen as a success for monetary policy.

The fourth panel shows the stochastic volatility in gap inflation,  $\xi_{v,t|t}$ . This is on a much larger scale than trend SV. Grassi and Proietti (2010) and Creal (2012) find that the

volatility of CPI inflation has increased recently, with the increased volatility attributed to the transitory rather than the permanent component. We do not find this upward trend. Notice, though, that gap uncertainty tends to rise late in recessions. Though there are only four recessions in the CPI sample there thus is a clear pattern: As recessions unfold gap inflation falls and its uncertainty rises.

Figure 4 shows posterior medians for GDP deflator inflation. The first panel shows the median and 68% uncertainty band for  $\tau_{t|t}$ . The median trend estimate peaks in 1981 at a value near 9%. It then declines to hover around 2%. It is notable that the 68% uncertainty band includes the 2% value as early as 1996.

The second panel shows the inflation gap  $\epsilon_{t|t}$ . Again its low persistence is apparent from the figure. The largest spike in gap inflation occurs in the 1973–1975 recession. Gap inflation is mostly negative from the end of the 1981–1982 recession to 2000. And it has negative spikes at the end of the 2007–2009 recession and at several point thereafter. According to the logic of the BN model, this suggests that forecasters expected actual inflation to rise and return to trend at these times.

The third panel shows the stochastic volatility in trend inflation,  $\xi_{\eta,t|t}$ . Prior to the 1990s there is much uncertainty associated with this measure, as shown by the wide 68% uncertainty band. But the median shows a considerable decline by the mid 1990s. Cogley and Sargent (2015) find a moderation in GDP deflator volatility in the 1990s using a long span of annual data up to 2012.

The fourth panel shows the SV in gap inflation,  $\xi_{v,t|t}$ . For GDP deflator inflation, this is now comparable in scale to that of trend volatility. Stock and Watson (2007) found that the moderation in inflation after 1990 was due to a fall in the innovation variance in the permanent component not the temporary one. Here there is some evidence of moderations in both components, though the decline in gap volatility appears to have been complete by the 1990s. In contrast with estimates for CPI inflation, there is little evidence that gap volatility rises in recessions, of which there are now seven in the sample.

The key finding is common to both measures of inflation. Trend inflation has converged near 2% and its volatility has declined over time. It is noteworthy that both tendencies date from the 1990s. But the behavior of the inflation gap differs depending on the inflation

measure. For CPI inflation, the inflation gap declines in recessions, while its volatility increases. And the stochastic volatility in the inflation gap has not declined over time. For GDP deflator inflation there is little evidence of a cyclical pattern in the gap or its volatility, and more evidence of a decline in that volatility at least until the 1990s.

## 7. Conclusion

This paper studies the unobserved components model of US inflation. The trend component is identified with long-term inflation expectations. The deviation of actual inflation from its trend, which is inflation's cyclical component, is interpreted as the inflation gap. We show how to measure these components using professional forecasts under either rational expectations or sticky information. The procedure involves only realized inflation and mean forecasts, as the use of other covariates is implicitly outsourced to professional forecasters.

Consistent with the assumptions of the model, the resulting innovations to trend and gap inflation are serially uncorrelated as are the innovations to the stochastic volatilities. We find that inflation-gap persistence is quite low. Information stickiness can be identified and is especially important to describe the 1970s. Whether inflation is measured by the CPI or by the GDP deflator, the trend converges to near 2% and its volatility declines over time, two tendencies largely complete by the late 1990s.

## References

- Andrieu, A., Doucet, A., & Holenstein, R. 2010. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society, Series B* 72(3), 269–342.
- Ang A., Bekaert G, & Wei M. 2007. Do macro variables, asset markets, or surveys forecast inflation better? *Journal of Monetary Economics* 54(4), 1163–1212.
- Ascari G., & Sbordone A. 2014. The macroeconomics of trend inflation. *Journal of Economic Literature* 52(3), 679–739.
- Berger, T., Everaert, G., & Hauke, V. 2016. Testing for time variation in an unobserved components model for the U.S. economy. *Journal of Economic Dynamics and Control* 69(1), 179–208.
- Beveridge, S., & Nelson, C.R. 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(2), 151–174.
- Cecchetti, S.G., Feroli, M., Hooper, P., Kashyap, A.K., & Schoenholtz, K., 2017. Deflating inflation expectations: The implications of inflation’s simple dynamics. CEPR Discussion Paper 11925.
- Chen, R., & Liu, J.S. 2000. Mixture Kalman filters. *Journal of the Royal Statistical Society, Series B* 62(3), 493–508.
- Cogley, T., Primiceri, G.E. & Sargent, T.J. 2010. Inflation-gap persistence in the U.S. *American Economic Journal: Macroeconomics* 2(1), 43–69.
- Cogley, T., & Sargent, T.J. 2015. Measuring price-level uncertainty and instability in the U.S., 1850–2012. *Review of Economics and Statistics* 97(4), 827–838.
- Cogley, T. & Sbordone, A. 2008. Trend inflation, indexation, and inflation persistence in the New Keynesian Phillips curve. *American Economic Review* 98(5), 2101–2126.
- Coibion, O. 2010. Testing the sticky information Phillips curve. *Review of Economics and Statistics* 92(1), 87–101.
- Coibion O., & Gorodnichenko Y. 2012. What can survey forecasts tell us about informational rigidities? *Journal of Political Economy* 120(1), 116–159.
- Coibion, O., & Gorodnichenko, Y. 2015. Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review* 105(8), 2644–2678.
- Creal D. 2012. A survey of sequential Monte Carlo methods for economics and finance. *Econometric Reviews* 31(3) 245–296.
- Croushore D. 2010. An evaluation of inflation forecasts from surveys using real-time data. *B.E. Journal of Macroeconomics, Contributions* 10(1) 1–32.
- Doucet, A., Pitt, M., Deligiannidis, G., & Kohn, R. 2015. Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *Biometrika* 102(2), 295–313.

- Faust, J. & Wright, J.H. (2013). Forecasting inflation. In G. Elliot and A. Timmermann (eds.), *Handbook of Economic Forecasting*, vol. 2, 2–56. New York, NY: Elsevier.
- Geweke, J. 2005. *Contemporary Bayesian Econometrics and Statistics*, J. Wiley and Sons, Inc.: Hoboken, NJ.
- Gil-Alana L., Moreno A., & Pérez de Gracia, F. 2012. Exploring survey-based inflation forecasts. *Journal of Forecasting* 31(6), 524–539.
- Grassi, S., & Proietti, T. 2010. Has the volatility of U.S. inflation changed and how? *Journal of Time Series Econometrics* 2(1), article 6.
- Henzel S.R. 2013. Fitting survey expectations and uncertainty about trend inflation. *Journal of Macroeconomics* 35, 172–185.
- Hol, J.D., Schn, T.B., & Gustafsson, F. (2006). On resampling algorithms for particle filters. In *IEEE Nonlinear Statistical Signal Processing Workshop*, Ng, W. (ed.), 79–82. Red Hook, NY: Curran Associates, Inc.
- Ireland P.N. 2007. Changes in the Federal Reserve’s inflation target: Causes and consequences. *Journal of Money, Credit and Banking* 39(8), 1851–1882.
- Jain M. 2017. Perceived inflation persistence. *Journal of Business and Economic Statistics* 37(1), 110–120.
- Jeffreys, H. 1998. *Theory of Probability*, Oxford University Press: Oxford, UK.
- Keane M.P., & Runkle D.E. 1990. Testing the rationality of price forecasts: New evidence from panel data. *American Economic Review* 80(4), 714–735.
- Kiley, M.T. 2007. A quantitative comparison of sticky price and sticky information models of price setting. *Journal of Money, Credit and Banking* 39(1), 101–125.
- Kozicki S., & Tinsley, P.A. 2012. Effective use of survey information in estimating the evolution of expected inflation. *Journal of Money, Credit and Banking* 44(1), 145–169.
- Li, T., Bolic, M., & Djuric, P. 2015. Resampling methods for particle filtering: Classification, implementation, and strategies. *IEEE Signal Processing Magazine* 32(3), 70–86.
- Lindström, J. 2017. Transformed proposal distributions. Manuscript, Centre for Mathematical Sciences, Lund University.
- Mankiw N.G., & Reis R. 2002. Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve. *Quarterly Journal of Economics* 117(4), 1295–1328.
- Mankiw N.G., Reis R., & Wolfers, J. 2004. Disagreement about inflation expectations. In *NBER Macroeconomics Annual 2003*, Volume 18, 208–248. Cambridge, MA: The MIT Press.

- Martino, L., Elvira, V., & Camps-Valls, G. 2018. Group importance sampling for particle filtering and MCMC. *Digital Signal Processing* 82(1), 133–151.
- Mertens, E. 2016. Measuring the level and uncertainty of trend inflation. *Review of Economics and Statistics* 98(5), 950–967.
- Mertens, E., & Nason, J.M. 2019. Inflation and professional forecast dynamics: An evaluation of stickiness, persistence, and volatility. Manuscript. Department of Economics, NC State University.
- Morley, J.C., Nelson, C.R., & Zivot E. 2004. Why are the Beveridge-Nelson and unobserved-components decompositions of GDP so different? *Review of Economics and Statistics* 85(2) 235–243.
- Owen, A.B. 2017. Statistically efficient thinning of a Markov chain sampler. *Journal of Computational and Graphical Statistics* 26(3), 738–744.
- Pitt, M.K., dos Santos Silva, R., Giordani, P., & Kohn, R. 2012. On some properties of Markov chain Monte Carlo simulation methods based on the particle filter. *Journal of Econometrics* 171(2), 134–151.
- Pitt, M.K., & Shephard, N. 2001. Auxiliary Variable Based Particle Filters. In *Sequential Monte Carlo Methods in Practice*, Doucet, A., de Freitas, N., & Gordon, N. (eds.), 273–293. New York, NY: Springer-Verlag.
- Pitt, M.K., & Shephard, N. 1999. Filtering via Simulation: Auxiliary Particle Filters. *Journal of the American Statistical Association* 94(446), 590–599.
- Särkkä, S. 2013. *Bayesian Filtering and Smoothing*, Cambridge University Press: Cambridge, UK.
- Shephard N. 2013. Martingale unobserved component models. Economics Series Working Papers 644, Department of Economics, University of Oxford.
- Stock J.H., & Watson, M.W. 2007. Why has US inflation become harder to forecast? *Journal of Money, Credit and Banking* 39(S1), 3–33.
- Stock, J.H., & Watson, M.W. 2010. Modeling inflation after the crisis. In *Macroeconomic Challenges: The Decade Ahead*, Kansas City, MO: Federal Reserve Bank of Kansas City.
- Stock, J.H., & Watson, M.W. 2016. Core inflation and trend inflation. *Review of Economics and Statistics* 98(4), 770–784.
- Vihola, M. 2012. Robust adaptive Metropolis algorithm with coerced acceptance rate. *Statistics and Computing* 22(5), 997–1008.
- Watson, M.W. 1986. Univariate detrending methods with stochastic trends. *Journal of Monetary Economics* 18(1) 49–75.

**Table 1. RE- and SI-SPF State Space Models**

**RE-SPF State Space Model,  $\mathcal{M}_{RE}$**

RE-Observation Equations:  $\Pi_t^{SPF} = \mathcal{C}_{RE}\mathcal{S}_{RE,t} + \mathcal{D}_{RE}\Psi_{RE,t}$ ,  $\Psi_{RE,t} \sim N(0_{3 \times 1}, I_3)$ , where  $\Pi_t^{SPF} = [\pi_{t,1}^{SPF} - \pi_t, \dots, \pi_{t,3}^{SPF} - \pi_t]'$ ,  $\mathcal{C}_{RE} = [\rho - 1, \dots, \rho^3 - 1]'$ ,  $\mathcal{S}_{RE,t+1} = \epsilon_{t+1}$ ,  $\mathcal{D}_{RE}$  is a diagonal matrix with its nonzero elements the column vector  $[\sigma_{\psi,1}, \dots, \sigma_{\psi,3}]'$ , and  $\Psi_{RE,t} = [\psi_{t,1}, \dots, \psi_{t,3}]'$ .

Conditionally Linear State Equation:  $\mathcal{S}_{RE,t} = \mathcal{A}_{RE}\mathcal{S}_{RE,t-1} + \mathcal{B}_{RE,t-1}\mathcal{E}_{RE,t}$ , where  $\mathcal{A}_{RE} = \rho$ ,  $\mathcal{B}_{RE,t} = \xi_{v,t}$ ,  $\mathcal{E}_{RE,t} = v_t$ , and  $v_t \sim N(0, 1)$ .

Nonlinear State Equation, Log Random Walk of Gap Inflation SV:  $\ln \xi_{v,t}^2 = \ln \xi_{v,t-1}^2 + \sigma_v \phi_{v,t}$ , where  $\phi_{v,t} \sim N(0, 1)$ .

**SI-SPF State Space Model,  $\mathcal{M}_{SI}$**

SI-Observation Equations:  $\Pi_t^{SPF} = \mathcal{C}_{SI,t}\mathcal{S}_{SI,t} - \lambda \xi_{\eta,t-1} \eta_t I_3$ , where the vector of predetermined lagged dependent variables,  $\Pi_{t-1}^{SPF} = [\pi_{t-1,1}^{SPF} - \pi_{t-1}, \dots, \pi_{t-1,3}^{SPF} - \pi_{t-1}]'$ , creates time variation in  $\mathcal{C}_{SI,t} = [(1 - \lambda)\mathcal{C}_{RE} I_3 \ 0_{3 \times 3} \ \Pi_{t-1}^{SPF}]$ ,  $\mathcal{S}_{SI,t} = [\epsilon_t \ \delta_{t,1} \dots \delta_{t,3} \ -\lambda \zeta_{t,1} \ \dots \ -\lambda \zeta_{t,3} \ \lambda]'$ ,  $\delta_{t,h} = \zeta_{t,h} - \lambda \zeta_{t-1,h}$ , and  $\zeta_{t,h} = \sigma_{\psi,h} \psi_{t,h}$  for  $h = 1, 2$ , and  $3$ .

Conditionally Linear State Equations:  $\mathcal{S}_{SI,t} = \mathcal{A}_{SI}\mathcal{S}_{SI,t-1} + \mathcal{B}_{SI,t-1}\mathcal{E}_{SI,t}$ , where the first row of  $\mathcal{A}_{SI}$  has  $\rho$  in the (1,1) position and zeros elsewhere, the next three rows of  $\mathcal{A}_{SI}$  are  $[\lambda \ I_3 \ 0_{3 \times 3} \ I_3 \ 0_{3 \times 3}]$  and the last four rows are full of zeros except the bottom diagonal position that is  $\lambda$ . Similarly,  $\mathcal{B}_{SI,t-1}$  has  $\xi_{v,t-1}$  as its (1,1) element and zeros in the rest of the first row. The next six rows are  $[0_{3 \times 1} \ I_3]$  and  $[0_{3 \times 1} \ -\lambda I_3]$ . Zeros fill the the last row of  $\mathcal{B}_{SI,t-1}$ . The SI-state impulse vector is  $\mathcal{E}_{SI,t} = [v_t \ \zeta_{t,1} \ \zeta_{t,2} \ \zeta_{t,3}]'$ .

Nonlinear State Equations, Log Random Walks of Trend and Gap Inflation SVs:  $\ln \xi_{\eta,t}^2 = \ln \xi_{\eta,t-1}^2 + \sigma_{\eta} \phi_{\eta,t}$  and  $\ln \xi_{v,t}^2 = \ln \xi_{v,t-1}^2 + \sigma_v \phi_{v,t}$ , where  $\phi_{\eta,t} \sim N(0, 1)$ ,  $\phi_{v,t} \sim N(0, 1)$ , and  $E[\phi_{\eta,t} \phi_{v,\ell}] = 0$  for all dates  $t$  and  $\ell$ .

**Table 2. Priors of the RE- and SI-SPF Models**

Model Parameter	Prior Distribution	Parameters		Quantiles		
		$\theta_1$	$\theta_2$	5%	50%	95%
Inflation Gap AR(1), $\rho$	$\mathcal{TN}$	0.00	1.00	[-0.87, 0.00, 0.87]		
Trend Inflation SV, $\sigma_\eta$	$\chi(3)$	0.00	0.20	[0.12, 0.31, 0.56]		
Inflation Gap SV, $\sigma_v$	$\chi(3)$	0.00	0.20	[0.12, 0.31, 0.56]		
Measurement Error, $\sigma_{\psi,h}^{CPI}$	$\mathcal{IG}$	14.90	6.70	[0.31, 0.45, 0.77]		
Measurement Error, $\sigma_{\psi,h}^{PGDP}$	$\mathcal{IG}$	20.00	9.00	[0.32, 0.46, 0.68]		
Sticky Information, $\lambda$	$\mathcal{TN}$	0.50	1.00	[0.05, 0.44, 0.95]		
Initial Gap SV, $\ln \xi_{v,0}^{2,CPI}$	$\mathcal{LN}$	0.08	1.00	[0.21, 1.08, 5.62]		
Initial Gap SV, $\ln \xi_{v,0}^{2,PGDP}$	$\mathcal{LN}$	-1.14	1.00	[0.06, 0.32, 1.66]		
Initial Trend SV, $\ln \xi_{\eta,0}^{2,CPI}$	$\mathcal{LN}$	-1.16	1.00	[0.06, 0.31, 2.23]		
Initial Trend SV, $\ln \xi_{\eta,0}^{2,PGDP}$	$\mathcal{LN}$	-2.38	1.00	[0.02, 0.09, 0.48]		

Columns under  $\theta_1$  and  $\theta_2$  are parameters of the prior distributions. The AR1 parameter,  $\rho$ , of gap inflation has a prior that is distributed truncated-normal ( $\mathcal{TN}$  and parameterized by a zero mean,  $\theta_1=0$ , and unit standard deviation,  $\theta_2=14$ ). The bounds on the prior are  $\rho \in (-1,1)$ . The scale volatilities  $\sigma_\eta$  and  $\sigma_v$  on the SVs of trend and gap inflation have priors distributed chi ( $\chi$ ) with three degrees of freedom, located at  $\theta_1=0$ , and the scale parameter  $\theta_2=0.20$ . The inverse-gamma ( $\mathcal{IG}$ ) distribution describes the priors on the scale volatility parameters,  $\sigma_{\psi,h}$ , of the measurement errors on the SPF inflation predictions net of realized inflation,  $h=1,\dots,4$ . The shape and scale coefficients are  $\theta_1$  and  $\theta_2$  for the priors on  $\sigma_{\psi,h}$ ,  $h=1,\dots,\mathcal{H}$ . Since the CPI-SPF and PGNP/PGDP-SPF inflation sample sizes differs,  $T_{CPI}=149$  and  $T_{PGDP}=200$ ,  $\theta_1=0.1 \times T$  and  $\theta_2=0.045 \times T$  for the former and latter samples. The SI parameter  $\lambda \in (0,1)$  and has a prior distributed  $\mathcal{TN}$  with mean and standard deviation of  $\theta_1=0.5$  and  $\theta_2=1.0$ . Inflation trend and gap stochastic volatilities are initialized by drawing from the log normal ( $\mathcal{LN}$ ) distribution. The prior means are computed as the log of the square of share of the variance of the difference of inflation attributed to trend or gap inflation minus one-half. The shares are 0.28 and 0.52, respectively, with the remaining 0.20 assigned to measurement error. The variance is 1.4 for  $\pi_t^{PGNP}$  from 1958Q1 to 1967Q4 and 2.57 for  $\pi_t^{CPI}$  from 1961Q to 1980Q4. The prior standard deviation is one for these initial conditions.

**Table 3. Summary of the Posterior Distributions**

Parameter	CPI Inflation Sample: 1981Q4–2018Q4		GNP/GDP Deflator Inflation Sample: 1969Q1–2018Q4	
	$\mathcal{M}_{RE}$	$\mathcal{M}_{SI}$	$\mathcal{M}_{RE}$	$\mathcal{M}_{SI}$
$\hat{\rho}$	0.099 [0.078, 0.120]	0.162 [0.126, 0.212]	0.234 [0.193, 0.276]	0.254 [0.198, 0.319]
$\hat{\sigma}_\eta$	0.447 [0.407, 0.514]	0.391 [0.221, 0.612]	0.400 [0.352, 0.475]	0.406 [0.266, 0.592]
$\hat{\sigma}_v$	0.352 [0.225, 0.511]	0.356 [0.229, 0.517]	0.175 [0.100, 0.282]	0.149 [0.088, 0.243]
$\hat{\sigma}_{\psi,1}$	0.220 [0.197, 0.247]	0.191 [0.163, 0.219]	0.289 [0.261, 0.319]	0.189 [0.161, 0.219]
$\hat{\sigma}_{\psi,2}$	0.132 [0.113, 0.153]	0.115 [0.099, 0.134]	0.195 [0.169, 0.225]	0.140 [0.122, 0.162]
$\hat{\sigma}_{\psi,3}$	0.180 [0.157, 0.204]	0.156 [0.132, 0.181]	0.261 [0.230, 0.293]	0.166 [0.142, 0.193]
$\hat{\lambda}$	—	0.438 [0.334, 0.548]	—	0.366 [0.290, 0.450]
ln MDD	−129.932	−150.516	−184.480	−125.837

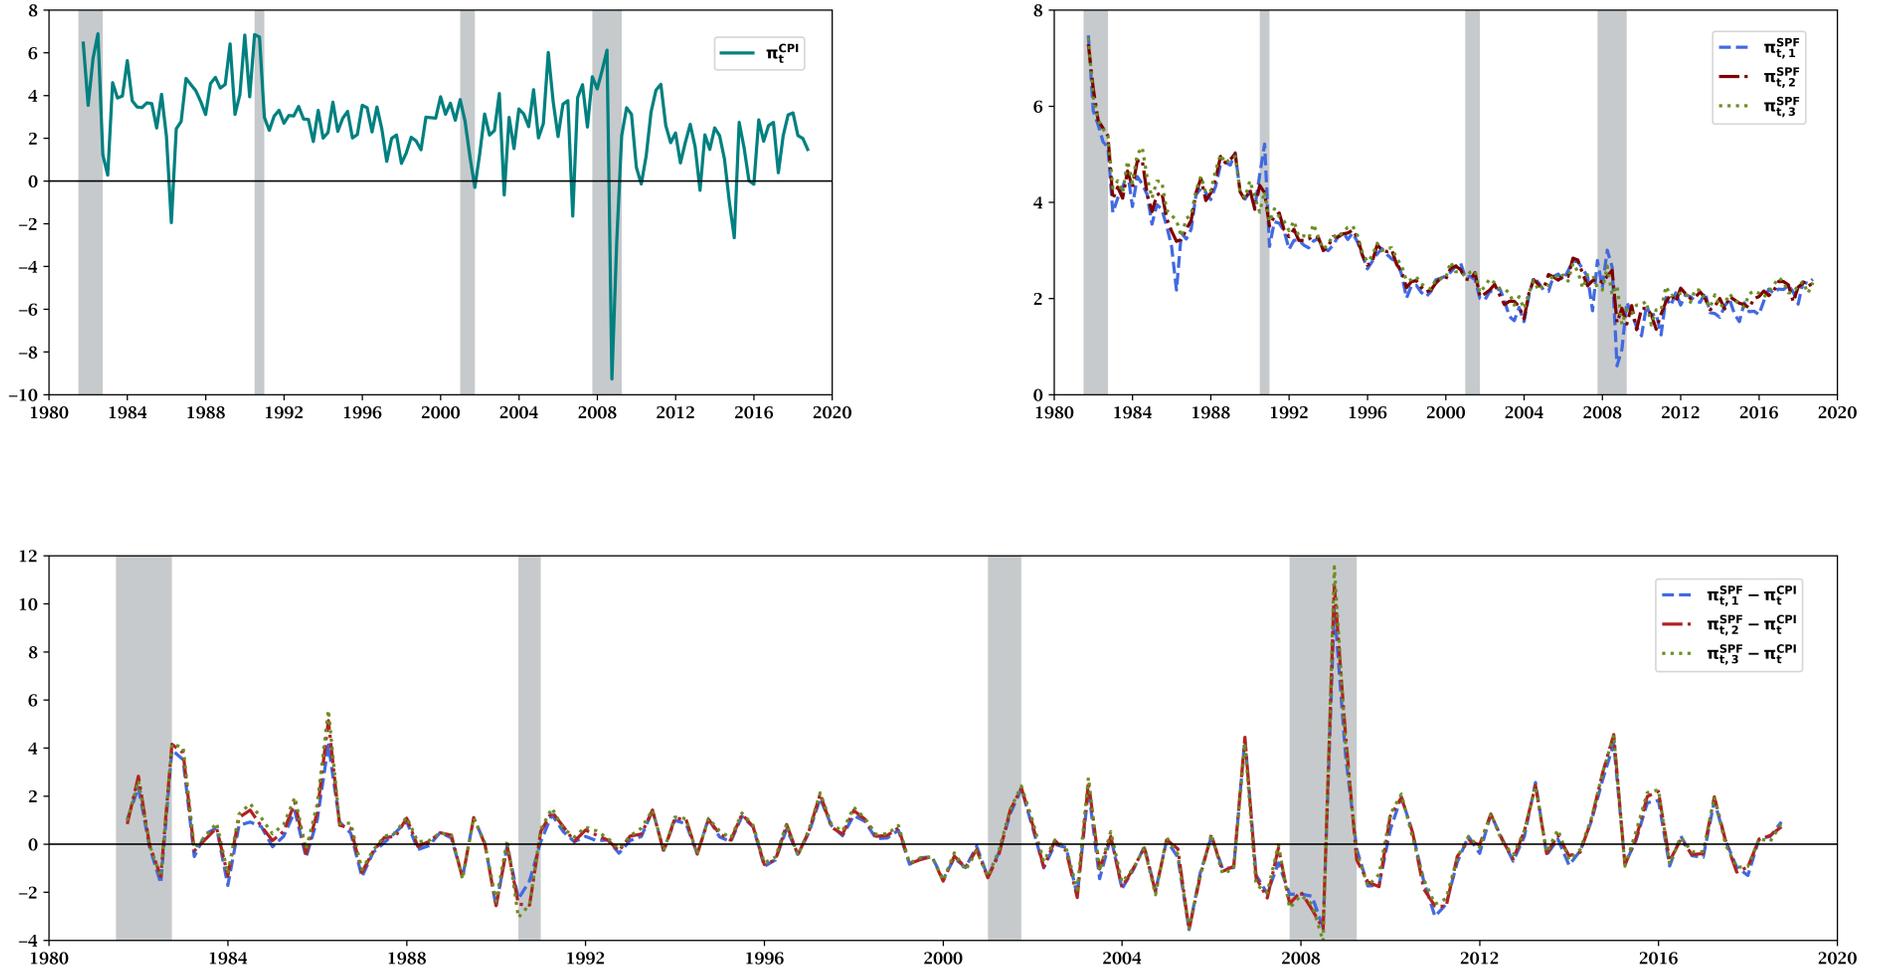
Posterior median estimates of RE- and SI-SPF model parameters are displayed in the table. Below these estimates, the brackets contain five and 95 percent quantiles extracted from the posterior distributions. The PMH-MCMC produces the posterior distribution of the scale volatility parameter,  $\sigma_\eta$ , on innovations to trend SV for the SI-SPF model,  $\mathcal{M}_{SI}$ . This parameter is backed out of the  $i$ th draw of the posterior distribution of trend SV,  $\xi_{\eta,t}$ , as  $\hat{\sigma}_{\eta,i} = \sqrt{T^{-1} \sum_{t=1}^T (\Delta \ln \xi_{\eta,t,i}^2)}$ ,  $i=1, \dots, \mathcal{K}$ , for the RE model,  $\mathcal{M}_{RE}$ ; see section 3 of the appendix for details. Log marginal data densities, ln MDDs, are listed in the last line of the table. We compute the ln MDDs using the modified harmonic mean estimator; see Geweke (2005) and section 3 of the appendix for details.

**Table 4. Posterior Moments of Innovations to Trend Inflation**

Parameter	CPI Inflation Sample: 1981Q4–2018Q4		GNP/GDP Deflator Inflation Sample: 1969Q1–2018Q4	
	$\mathcal{M}_{RE}$	$\mathcal{M}_{SI}$	$\mathcal{M}_{RE}$	$\mathcal{M}_{SI}$
$s(\eta_{t t})$	1.218 [1.140, 1.329 ]	2.381 [1.604, 3.676 ]	1.022 [0.955, 1.101]	1.303 [1.035, 1.655]
ACF(1)	-0.002 [-0.042, 0.034]	0.027 [-0.025, 0.088]	0.030 [-0.005, 0.074]	0.071 [0.006, 0.136]
ACF(2)	-0.200 [-0.230, -0.151]	-0.120 [-0.191, -0.061]	-0.192 [-0.243, -0.125]	-0.123 [-0.204, -0.047]
ACF(3)	-0.010 [-0.043, 0.012]	0.013 [-0.049, 0.072]	0.007 [-0.026, 0.048]	0.019 [-0.040, 0.063]
ACF(4)	0.045 [0.015, 0.084]	-0.003 [-0.055, 0.042]	0.186 [0.159, 0.203]	0.198 [0.157, 0.247]
ACF(8)	0.071 [0.044, 0.101]	0.156 [0.049, 0.129]	0.096 [0.055, 0.129]	0.101 [0.037, 0.189]
$Q(4)$	6.608 (0.177)	3.263 (0.568)	15.077 (0.011)	13.488 (0.022)
$Q(8)$	17.862 (0.032)	12.215 (0.196)	20.106 (0.017)	21.331 (0.025)

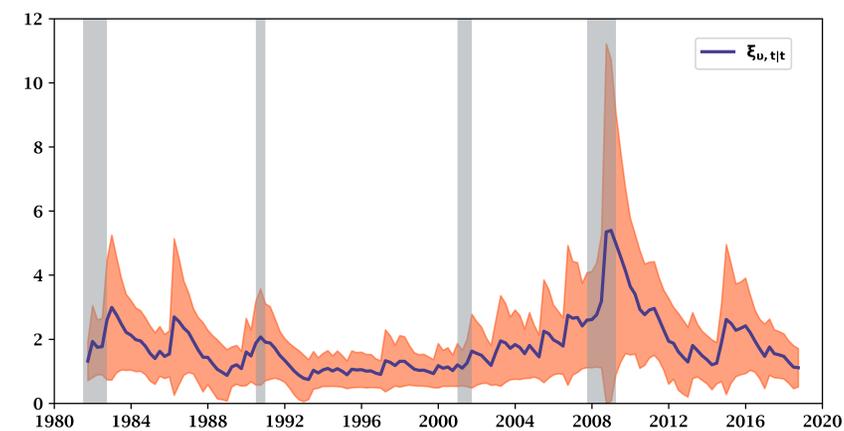
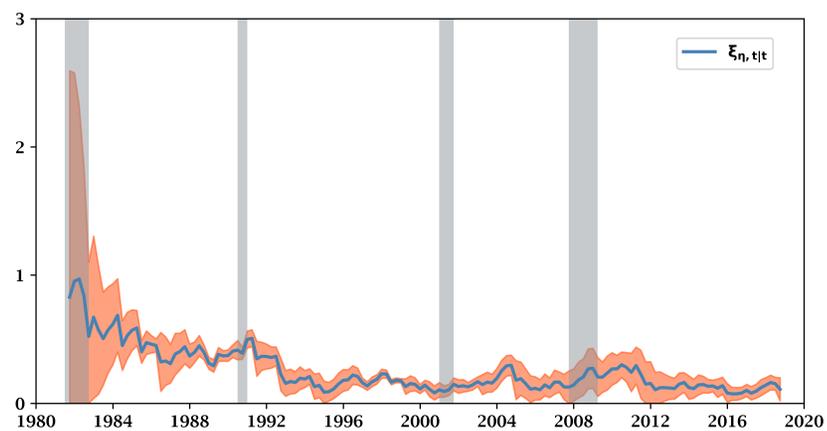
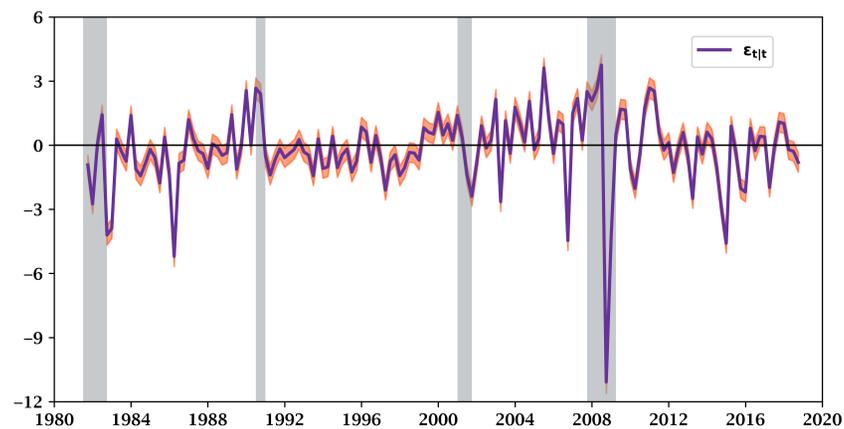
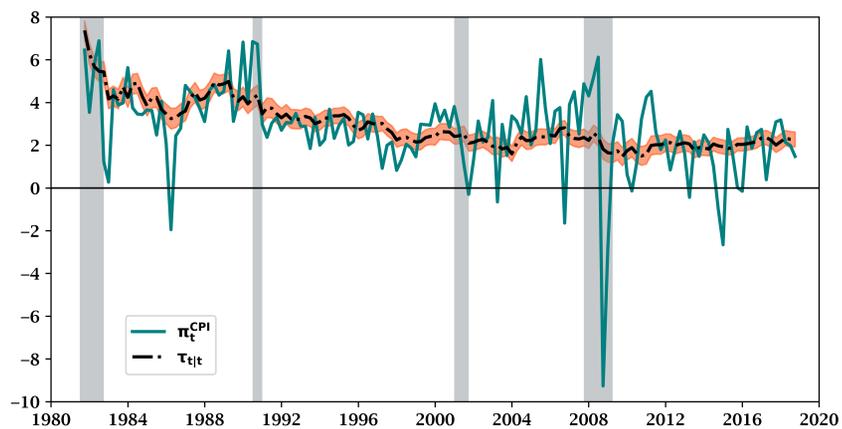
The table presents posterior moments of innovations to filtered trend inflation,  $\eta_{t|t}=(\tau_{t|t}-\tau_{t-1|t-1})/\xi_{\eta,t-1}$ , which imply a loss of one observation from  $T_{CPI}=149$  and  $T_{PGDP}=200$ . The first row is the median of the posterior standard deviation of  $\eta_{t|t}$ ,  $s(\eta_{t|t})$ . The median autocorrelation function at lag  $j$  is denoted  $ACF(j)$ . Five and 90 percent quantiles appear in brackets. The mean posterior Ljung-Box statistic with  $q$  lags is in the row denoted  $Q(q)$ . The rows below  $Q(4)$  and  $Q(8)$  display Bayesian  $p$ -values in parentheses.

Figure 1: CPI Inflation and SPF Predictions, 1981Q4–2018Q4



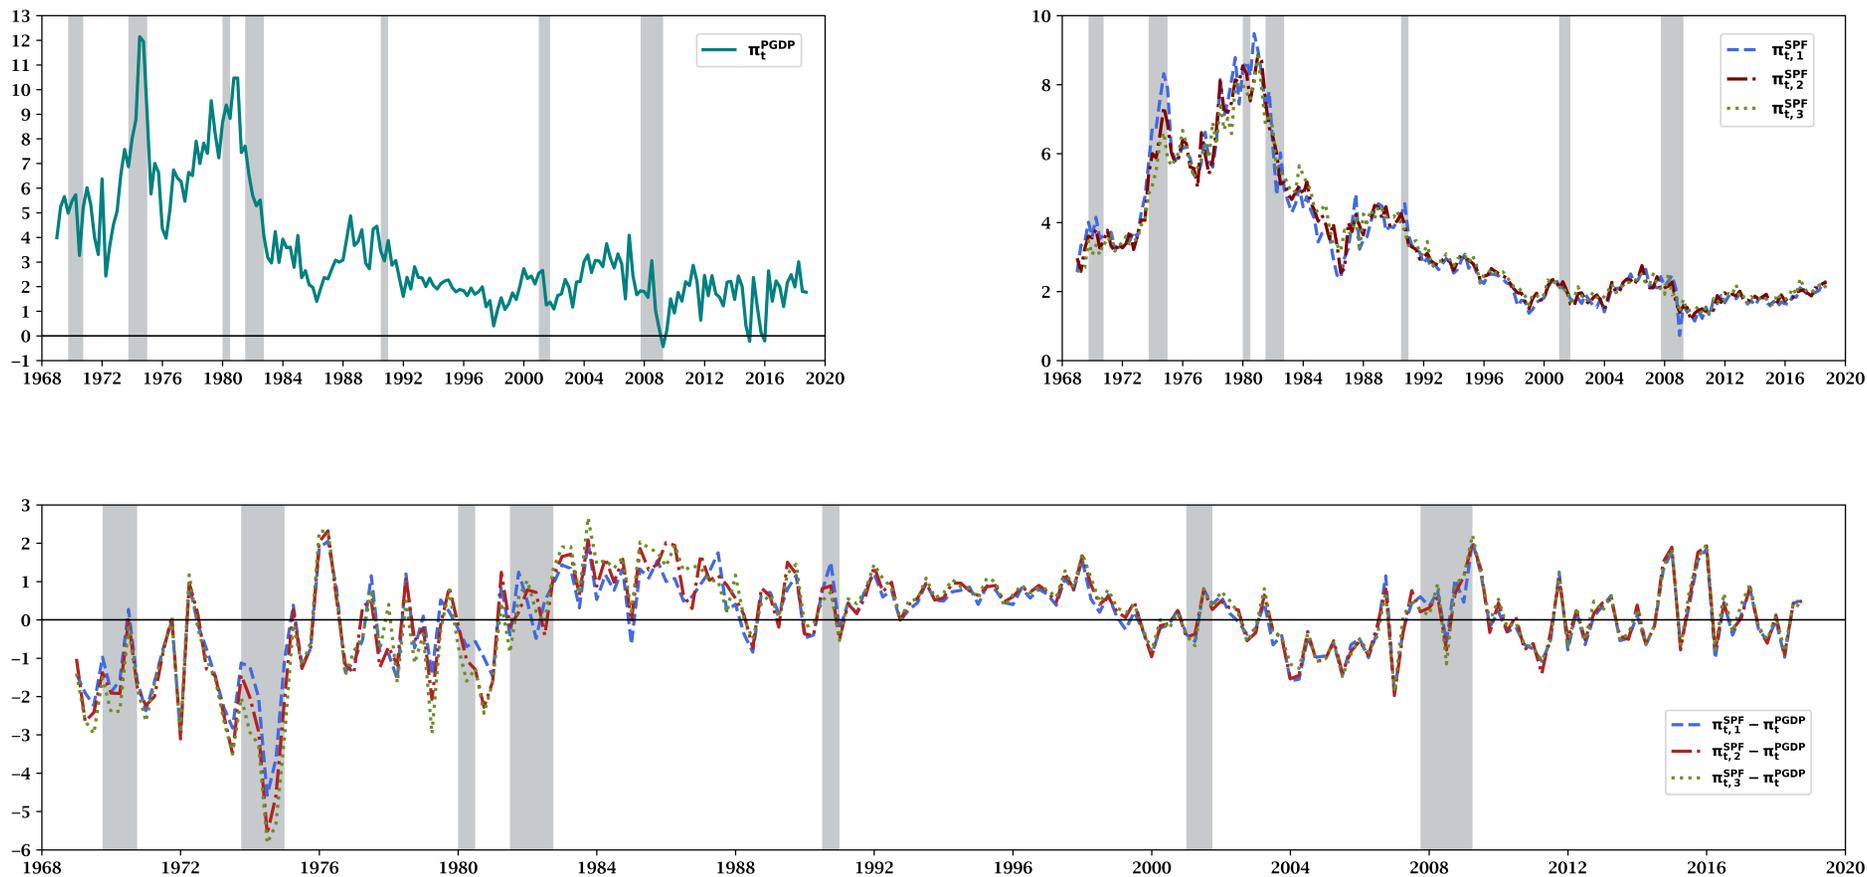
Notes: The first panel shows realized CPI inflation. The second panel shows SPF forecasts at a common date, at  $h = 1, 2, 3$  quarters ahead. The third panel shows the differences  $\Pi_{t,h}^{SPF} \equiv \pi_{t,h}^{SPF} - \pi_t$  used in the observation equations. Vertical gray bands denote NBER dated recessions.

Figure 2: Unobserved States in CPI Inflation, 1981Q4–2018Q4



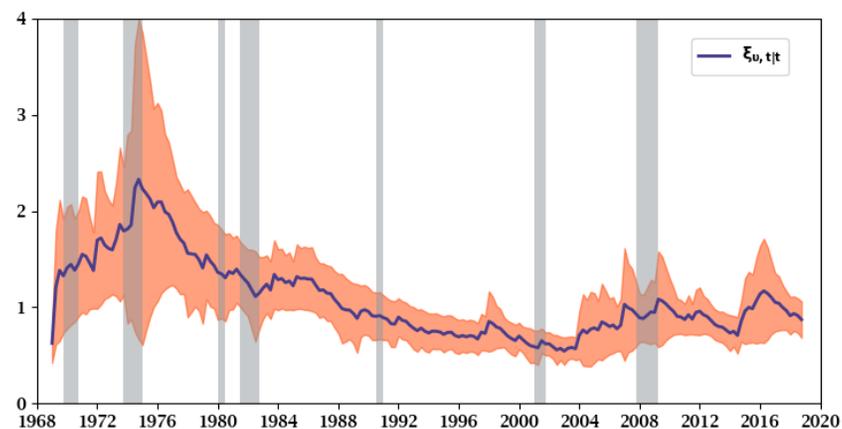
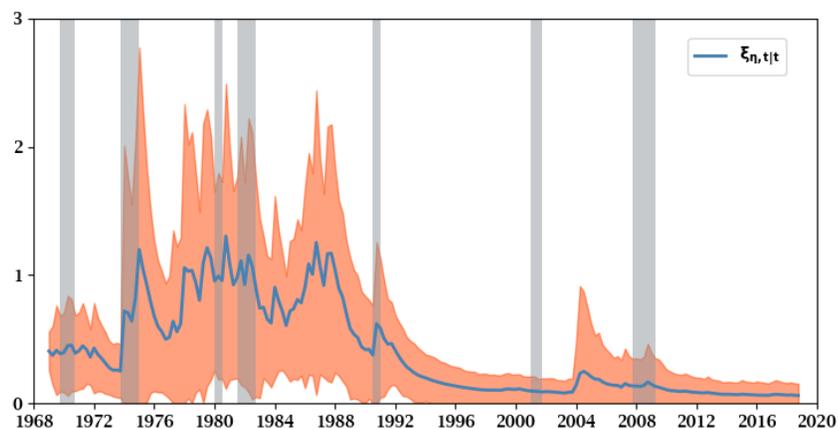
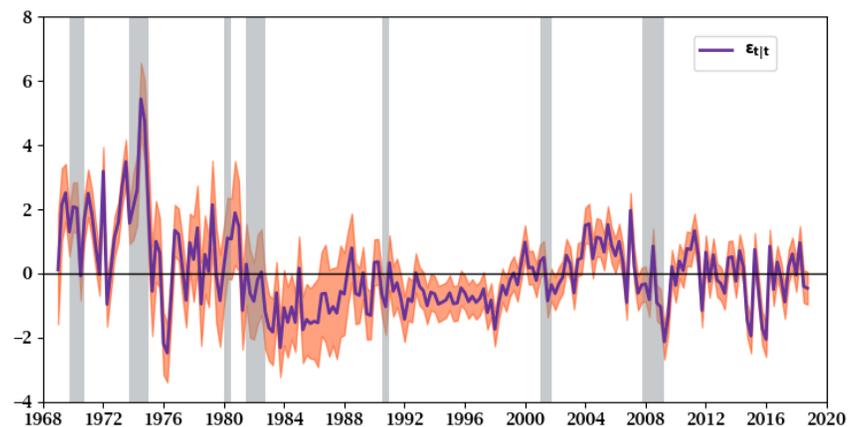
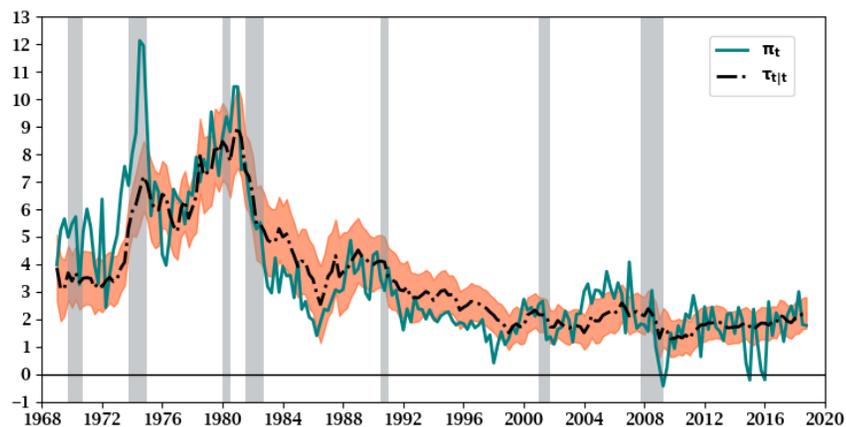
Notes: The first panel shows realized CPI inflation with the median of the posterior distribution of trend inflation  $\tau_{t|t}$ . The second panel shows the median of the posterior distribution of gap inflation  $\epsilon_{t|t}$ . The lower panels shows the medians of the posterior distributions of SV in trend and gap inflation,  $\xi_{\eta,t}$  and  $\xi_{\nu,t}$ , respectively. The panels also contain 68% uncertainty bands. Vertical gray bands denote NBER dated recessions.

Figure 3: PGNP/PGDP Inflation and SPF Predictions, 1969Q1–2018Q4



Notes: The first panel shows realized PGNP/PGDP inflation. The second panel shows SPF forecasts at a common date, at  $h = 1, 2, 3$  quarters ahead. The third panel shows the differences  $\Pi_{t,h}^{\text{SPF}} \equiv \pi_{t,h}^{\text{SPF}} - \pi_t^{\text{PGDP}}$  used in the observation equations. Vertical gray bands denote NBER dated recessions.

Figure 4: Unobserved States in PGNP/PGDP Inflation, 1969Q1–2018Q4



Notes: The first panel shows realized PGNP/PGDP inflation with the median of the posterior density of trend inflation  $\tau_{t|t}$ . The second panel shows the median for gap inflation  $\epsilon_{t|t}$ . The lower panels shows the medians for SV in trend and gap inflation,  $\xi_{\eta,t}$  and  $\xi_{\nu,t}$  respectively. The panels also contain 68% uncertainty bands. Vertical gray bands denote NBER dated recessions.