Federal Reserve Bank of Cleveland Working Paper Series

A New Model of Trend Inflation Using Disaggregates, Survey Expectations, and Uncertainty

Ellis Tallman and Saeed Zaman

Working Paper No. 26-08R

November 2026

Suggested citation: Tallman, Ellis and Saeed Zaman. 2026. "A New Model of Trend Inflation Using Disaggregates, Survey Expectations, and Uncertainty." Working Paper No. 26-08R. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202608r.

Federal Reserve Bank of Cleveland Working Paper Series
ISSN: 2573-7953

Working papers of the Federal Reserve Bank of Cleveland are preliminary materials circulated to stimulate discussion and critical comment on research in progress. They may not have been subject to the formal editorial review accorded official Federal Reserve Bank of Cleveland publications.

See more working papers at: www.clevelandfed.org/research. Subscribe to email alerts to be notified when a new working paper is posted at: https://www.clevelandfed.org/subscriptions.

This work is licensed under Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International. To view a copy of this license, visit https://creativecommons.org/licenses/by-nc-nd/4.0/

A New Model of Trend Inflation Using Disaggregates, Survey Expectations, and Uncertainty*

Ellis Tallman† Saeed Zaman‡

This version: October 27, 2026

Abstract

This paper develops a new empirical model that estimates trend inflation by combining modeling features that have advanced the literature on trend inflation over the past two decades. These features include incorporating information about long-term inflation expectations from surveys in a flexible way, modeling aggregate inflation via sectoral data (goods and services), allowing for stochastic volatility (SV) in the shocks to the trend and transitory components of inflation, allowing for a time-varying price Phillips curve, and allowing for time-varying uncertainty effects on the level of inflation. We estimate the model using Bayesian methods. We assess the relative importance of each feature for trend estimation and document the new model's competitive properties. The new model provides a more interpretable historical decomposition of inflation data than the models it extends. The decomposition suggests that uncertainty effects play a greater role than business cycle effects in explaining inflation fluctuations.

Keywords: disaggregates of inflation, inflation uncertainty, trend inflation, inflation expectations, nonlinear state space, Bayesian methods

JEL Codes: C11, C32, E31

*We thank editors of Advances in Econometrics, two anonymous referees, Todd Clark, Christian Matthes, James Mitchell, and participants at the Advances in Econometrics Bayesian Macroeconometrics Workshop, IAAE 2026, and Workshop on Uncertainty, Economic Activity and Policy for helpful conversations and comments. The views expressed herein are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Federal Reserve System.

†Economist Emeritus, Federal Reserve Bank of Cleveland, 1455 E. 6th St., Cleveland, OH 44114. Email: ellistallman@gmail.com

‡Senior Research Economist, Federal Reserve Bank of Cleveland, 1455 E. 6th St., Cleveland, OH 44114. Email: saeed.zaman@clev.frb.org

1 Introduction

Over the past two decades, econometric models to estimate trend inflation, a latent variable, have advanced significantly. This progress partly reflects the successful practice of modeling inflation dynamics using a general formulation in which inflation is thought to fluctuate around a time-varying *trend*.[1] Trend inflation is considered a long-run estimate of the inflation rate, conditional on prices and/or other economic data available up to the present, consistent with the notion of the Beveridge and Nelson (1981) trend.[2] Accordingly, it is common to estimate trend inflation as a (driftless) random walk. In an influential paper on trend inflation, reflecting the general formulation to model inflation discussed above, Stock and Watson (2007) developed a univariate unobserved components (specifically, local-level) model with stochastic volatility (UCSV) in both the trend and the transitory components of inflation. At a high level, the model can be viewed as a time-varying time-series smoothing method that adapts to changes in the signal-to-noise ratio of the trend and transitory components. This model has been shown to perform reasonably well in forecasting inflation. It is often considered the benchmark model in various studies on trend inflation estimation and point and density forecasting of inflation. Accordingly, since 2007, contributions to the trend inflation literature have generally involved extending Stock and Watson's model with empirically important features. However, as detailed below, empirical extensions of the Stock and Watson model have generally been explored in isolation. In this paper, we combine these extensions into a single model and assess whether doing so benefits trend inflation estimation.

Stock and Watson (2007) treat the cyclical component as noise (idiosyncratic), whose variance is time-varying. Stella and Stock (2015) and Chan et al. (2016), among others, extend

[1] See "Inflation, Trends, and Long-Run Expectations: Perspectives from Forecasting Research" by Todd E. Clark, https://www.federalreserve.gov/monetarypolicy/files/FOMC20180118memo01.pdf.

[2] The term "trend inflation" is also widely used in the literature on estimation of medium-term inflation developments as captured by measures such as trimmed-mean and median measures of inflation that are computed using the cross-sectional (sectoral) data on inflation. In contrast to long-run trend inflation, medium-term trend inflation is strongly associated with the business cycle, and thus measures such as trimmed-mean inflation exhibit statistically significant correlations with variables such as the unemployment rate and/or the output gap (see Ball et al. (2024); Garciga et al. (2025)).

the Stock and Watson model by modeling the cyclical component of inflation as a function of the unemployment rate and/or output gap via the Phillips curve (PC) relationship. Importantly, considering the ample evidence of instability in the Phillips curve relationship in US data, these studies permit a time-varying Phillips curve, which, together with SV, implies a time-varying predictability of inflation. These researchers demonstrate that, by more effectively isolating cyclical variations in inflation data through joint inference of movements in inflation and the unemployment rate, the resulting inflation trend estimates are less volatile than those obtained from the UCSV of Stock and Watson.

Chan et al. (2013) extend the UCSV of Stock and Watson by imposing bounds on the random walk process of trend inflation (i.e., the trend inflation is constrained to lie in an interval) and show that doing so marginally helps improve forecasts of inflation and yields less volatile estimates of trend inflation. Adding bounds, motivated by central bank inflation target ranges, imposes inequality constraints on the latent components, which, combined with SV, make the model a nonlinear state-space model. Accordingly, they developed a Bayesian estimation algorithm based on the methods of Chan and Jeliazkov (2009) and Chan and Strachan (2012), which we use in this paper.

Chan (2017) extends the Stock and Watson model to allow uncertainty (as measured by an estimate of stochastic volatility) to have a direct, time-varying effect on the level of inflation, known as the stochastic volatility in mean model (SVM). Chan's model remains univariate in that it uses only information from inflation's past and includes estimated stochastic volatility as a covariate, thereby permitting volatility feedback on the level of inflation. Following much of the literature on uncertainty, the estimated volatility is interpreted as a model-based measure of inflation uncertainty. Chan finds strong evidence of time variation in the parameter governing volatility feedback. Although Chan (2017) did not focus on trend inflation, his results indicated the important role of uncertainty in explaining inflation fluctuations.[3]

[3]Following widespread usage, Stock and Watson used the RW assumption to define the law of motion for the evolution of stochastic volatility. Eisenstat and Strachan (2016) explored various formulations of the stochastic volatility process in the context of the Stock and Watson model and found that all SV formula-

Chan et al. (2018) extend Stock and Watson by bringing information from long-run survey-based inflation expectations. Specifically, they allow long-run expectations to influence the estimates of trend inflation by modeling a direct link between the two, which is allowed to be time-varying. They find that doing so leads to substantially more stable and reasonable estimates of trend inflation. Similarly, using a multivariate extension of the Stock and Watson unobserved components model, Mertens (2016) shows that adding long-run survey forecasts of inflation, alongside other indicators, delivers crucial additional information in refining trend inflation estimates. The results of Zaman (2024) reinforce the findings of Chan et al. (2018) and Mertens (2016). It would be expected that when inflation expectations are anchored to the central bank's inflation target, estimates of the long-run trend in inflation are closely associated with the central bank's inflation objective. It is also a common practice to equate trend inflation with long-run survey expectations of inflation in macroeconomic models. However, simply equating trend inflation with survey expectations is not recommended, as shown by Chan et al. (2018), Garcia and Poon (2022), and Zaman (2024). This is because the survey data can sometimes be biased measures of trend inflation for various reasons, including information rigidities, which leads to sluggish adjustment of expectations, as emphasized by Coibion and Gorodnichenko (2015).

More recently, Zaman (2025) shows, using a bivariate model of inflation and inflation expectations, that including data on long-run survey-based expectations significantly reduces the empirical importance of imposing bounds on trend inflation, as proposed by Chan et al. (2013). Put differently, once inflation expectations are incorporated into the model, imposing bounds on the random-walk process that defines trend inflation is empirically unimportant.

Finally, a strand of the literature has extended Stock and Watson (2007) model to multivariate extensions that estimate the trend in aggregate inflation using cross-sectional disaggregates or components of inflation. Tallman and Zaman (2017) estimate the trend in

tions yield similar posterior inferences about volatility, suggesting that Stock and Watson's use of the RW assumption is reasonable.

aggregate inflation by separately modeling the goods and services components of aggregate inflation. Specifically, they apply the univariate Stock and Watson method to inflation in the goods sector. For services inflation, they extend the univariate Stock and Watson model to a bivariate UCSV model by jointly modeling services inflation and the unemployment rate. In modeling services inflation, they allow for a time-varying Phillips curve relationship. The aggregate trend is constructed as a weighted composite of the estimated goods inflation trend and the estimated services inflation trend, with weights equal to the respective nominal expenditure shares of each component in the overall consumer expenditure basket.

Eo et al. (2023) also model the aggregate inflation trend by modeling trends in the goods and services components of inflation, which they denote as a two-sector UCSV. They apply the UCSV Stock and Watson model to both services and goods inflation and directly link these two inflation measures via the error covariance structure, allowing a time-varying correlation between the innovations in the measurement equations of services and goods inflation and in the state equations describing the respective trends. By doing so, they explain the source of changes in the volatility of the aggregate inflation trend.[4] Stock and Watson (2016) model the trend in aggregate inflation by jointly modeling the co-movements in the 17 components of aggregate inflation using a factor model. In their setup, the 17 sectors are linked via a common factor, with time-varying loadings on the factor that reflect each sector's changing influence on trend inflation. One insightful finding of Stock and Watson (2016) is that food components have carried a nontrivial weight in determining trend inflation since the 1990s, suggesting that approaches that exclude food and energy (i.e., core inflation) for trend estimation may not be optimal. Compared with the univariate Stock and Watson

[4]Because Eo et al. (2023) applied the univariate UCSV of Stock and Watson separately to goods and services inflation and then combined the resulting trend estimates of goods and services via expenditure weights, they obtained posterior mean estimates of aggregate trend quite similar to those of the univariate SW model, but with improved precision, i.e., tighter credible bands. They allow for a time-varying covariance structure among the innovations to the various components of the models, including the trends, which allowed them to decompose the conditional volatility of aggregate trend inflation into the volatility of the trend in the goods sector, the volatility of the trend in services inflation, and the comovement between the trends in goods and services, facilitating a narrative about the movements in the aggregate trend (documented elsewhere in the literature, e.g., Stock and Watson (2007)).

model, both Stock and Watson (2016) and Eo et al. (2023) report comparable point forecast performance for their multivariate models.

Although the models discussed above yield comparable average forecast accuracy, the resulting estimates of trend inflation differ substantially at times, in turn implying contrasting inflation forecasts. Figure 1 plots the trend inflation estimates from some of the models discussed above. As is evident, different modeling assumptions can sometimes yield very different inflation trend estimates. Given the comparable forecast accuracy of these models, yet at times producing different estimates of the inflation trend (and implied inflation forecasts), Clark and Doh (2014) suggest using a suite of models to track the inflation trend. During periods when most models yield similar estimates of trend inflation, a policymaker might be confident in the inference; however, models indicating notable divergences, i.e., different signals about the trend inflation, would necessitate robust policy deliberations.

The empirical features discussed above have not yet been incorporated jointly into a single model. To address this gap, in this paper, we extend the univariate UCSV of Stock and Watson in the following ways: we (1) model the trend in aggregate inflation by separately modeling the dynamics in goods and services inflation (two-sector UCSV model); (2) allow for a time-varying Phillips curve to isolate business cycle variations in services inflation; (3) allow for uncertainty effects on the levels of goods and services inflation; and (4) introduce information from the long-run survey expectations to anchor the model-based estimate of the (implied) aggregate inflation trend, which is constructed as the weighted average of the trends in goods and services inflation. We not only combine the features into a single model but also thoroughly investigate their relative importance for inflation dynamics.

Overall, our results indicate that modeling choices influence the estimated inflation trend and the inference about model parameters; some choices matter more than others. In particular, including survey expectations data has the greatest impact, whereas allowing for the Phillips curve relationship has the least. The modeling of the aggregate inflation trend using sectoral data is useful when other features are jointly modeled. Furthermore, sectoral data

improve both the precision of the aggregate trend and the accuracy of the density forecast, especially for the sample period that includes the post-COVID inflation surge, supporting the findings of Watson (2026), who shows that the UCSV model with sectoral data tracked the post-COVID inflation surge better than the univariate model.

We summarize our findings as follows. First, inflation uncertainty plays an important role in explaining inflation fluctuations. Our results indicate economically and statistically significant evidence of time variation in the effects of uncertainty on the levels of goods and services inflation, especially goods inflation. This suggests the importance of jointly modeling sectoral data and uncertainty to tease out movements in inflation associated with the trend component. Importantly, we find that the inflation response to the estimated uncertainty is substantially smaller (or zero for services inflation) in a model without long-run survey expectations than in a model that includes survey data. Without survey data, the trend component explains a much larger share of inflation variation, reducing the share attributed to the uncertainty component. The inclusion of survey data anchors the trend estimate at stable levels, leaving a larger share of inflation variation to be explained by the uncertainty component and the Phillips curve.

Second, consistent with prior research, including long-run survey expectations to inform the model-based trend is key to obtaining a more reasonable and stable estimate of trend inflation.[5] The pseudo-out-of-sample forecasting exercises evaluating forecast performance at horizons of 1 to 3 years suggest comparable point performance but more accurate density forecasts for the model with survey data, reinforcing the findings of Chan et al. (2018). Moreover, survey expectations data played a crucial role in reliably inferring the inflation trend during the post-COVID inflation surge. Excluding survey data from the model yielded markedly different inferences, attributing most of the pandemic surge to trend inflation, similar to Stock and Watson (2007), which, in hindsight, seems unreasonable. From mid-

[5]As in prior research, in this paper, we use the terms "more reasonable" or "sensible" to characterize trend inflation estimates that are much less volatile (or erratic or variable), as would be expected of a long-run estimate.

2021 through 2023, inflation moderated from a high of 8 percent to a rate closer to 3 percent, suggesting most of the surge was transitory.

Third, the new model provides a more interpretable historical decomposition of inflation data than the models it extends. The decomposition suggests that uncertainty effects play a much greater role than business cycle effects in explaining inflation fluctuations.

Lastly, the baseline model that combines all features yields point and density forecasts that generally rival those of the univariate Stock and Watson benchmark model. Sectoral data, survey data, and the PC relationship improve medium-term forecast accuracy, whereas including uncertainty significantly improves near-term forecast accuracy but (conditional on including survey data) worsens medium-term forecast accuracy, offsetting the gains from other features. Although the baseline model's average forecast accuracy is comparable to that of the univariate Stock and Watson model, it provides a more detailed decomposition of inflation data; hence, we view this as a useful and practical contribution.

The paper is organized as follows. The next section describes the empirical model and its variants in detail. Section 3 describes the data, estimation, and model comparison metrics. Section 4 presents empirical results, and Section 5 concludes. This paper includes a supplementary online appendix that discusses Bayesian estimation steps, prior-posterior comparisons, and additional results.

2 Empirical Framework

2.1 Baseline: Two-Sector Empirical Model with All Features

Our empirical model substantially builds on the two-sector UCSV model of Eo et al. (2023), who extend the univariate UCSV Stock and Watson model to a bivariate setting. Specifically, they apply the Stock and Watson model separately to services and goods inflation and permit a direct link between the two inflation measures via a (time-varying) correlation be-

tween the innovations in the equations describing trends in the respective services and goods inflation. We extend their bivariate UCSV model (but without the correlations in the model innovations) by introducing the following: (1) inclusion of the cyclical unemployment rate to tease out movements in inflation associated with the business cycle; (2) inclusion of volatility feedback – permitting a direct (possibly time-varying) influence of volatility (uncertainty) shocks on the levels of goods and services inflation; and (3) bringing in information from the long-run survey expectations of inflation to inform the model-based estimate of trend inflation.[6]

The details of our model, which jointly estimates the dynamics of goods and services inflation, are as follows:

Services inflation is modeled as

$$\pi_t^S = \underbrace{\tau_t^S}_{\text{Trend}} + \underbrace{\lambda_t \ (U_t - U_t^*)}_{\text{Business cycle effect}} + \underbrace{\alpha_t^S \ (e^{h_t^{Ser}} - e^{\overline{h}_t^{Ser}})}_{\text{Uncertainty effect}} + \varepsilon_t^S, \ \ \varepsilon_t^S \sim N(0, e^{h_t^{Ser}}) \qquad (1)$$

where $\overline{h}_t^{Ser} = \frac{\sum_{i=1}^{t-1} h_i^{Ser}}{t-1}$; U_t^* is informed from outside the model

This formulation assumes services inflation fluctuates around its long-run time-varying trend (τ_t^S), and fluctuations respond to deviations of the unemployment rate from its natural rate (U^*) and deviations of uncertainty ($e^{h^{Ser}}$) from its historical average ($e^{\overline{h}^{Ser}}$) . Accordingly, in the long run, as the unemployment rate and the uncertainty gaps, assumed stationary, converge to zero, services inflation gravitates to τ^S.

[6]Unlike Eo et al. (2023), we do not model correlations between innovations to the trend in goods inflation and innovations to the trend in services inflation, nor between the innovations in the services and goods measurement equations. There are two main reasons for this modeling choice: (1) to make the estimation feasible because our model is richer than Eo et al.'s model, even though it lacks the explicit modeling of the error correlation structure; (2) because our first-order interest is in the trend estimates and, as Eo et al. show, allowing a direct link via the second moments does not affect the aggregate trend estimates; it helps explain the drivers of volatility in aggregate inflation, which in their case are similar to trend estimates of aggregate inflation from the Stock and Watson univariate model. That said, even though we do not allow for the "direct" link between the two inflation measures via the second moments (error-covariance structure), we do allow the first-moment link between the goods and services inflation measures via the respective trend measures in the inflation accounting equation (see equation 7).

Goods inflation is modeled as (measurement equation)

$$\pi_t^G = \underbrace{\tau_t^G}_{\text{Trend}} + \underbrace{\alpha_t^G \ (e^{h_t^{Goods}} - e^{\overline{h}_t^{Goods}})}_{\text{Uncertainty effect}} + \ \varepsilon_t^G, \ \ \varepsilon_t^G \sim N(0, e^{h_t^{Goods}}) \tag{2}$$

where $\overline{h}_t^{Goods} = \frac{\sum_{i=1}^{t-1} h_i^{Goods}}{t-1}$

Following prior research, which has found a weak to nonexistent PC relationship for goods inflation, we do not include the PC relationship in equation 2.

The time-varying parameters: λ, the slope of the price Phillips curve is constrained between (-1,0),

$$\lambda_t = \lambda_{t-1} \ + \ \varepsilon_t^{\lambda}, \ \ \varepsilon_t^{\lambda} \sim TN(-1 - \lambda_{t-1}, 0 - \lambda_{t-1}; 0, \sigma_{\lambda}^2) \tag{3}$$

The parameters, α^S and α^G capturing volatility feedback, evolve according to a RW,

$$\alpha_t^i = \alpha_{t-1}^i \ + \ \varepsilon_t^{\alpha i}, \ \ \varepsilon_t^{\alpha i} \sim N(0, \sigma_{\alpha i}^2), \quad i = \{S, G\} \tag{4}$$

The trend in services inflation, τ_t^S, is modeled as a driftless RW, with a shock term, whose variance is allowed to vary over time,

$$\tau_t^S = \tau_{t-1}^S \ + \ \varepsilon_t^{\tau S}, \ \ \varepsilon_t^{\tau S} \sim N(0, e^{h_t^{\tau S}}) \tag{5}$$

The trend in goods inflation, τ_t^G, is modeled as a driftless RW, with a shock term, whose variance is allowed to vary over time,

$$\tau_t^G = \tau_{t-1}^G \ + \ \varepsilon_t^{\tau G}, \ \ \varepsilon_t^{\tau G} \sim N(0, e^{h_t^{\tau G}}) \tag{6}$$

The trend in aggregate inflation is a weighted average of trends in services and goods inflation, where w_t^G and w_t^S are time-varying weights – reflecting the relative shares of services

inflation and goods inflation in the overall headline inflation,

$$\tau_t^* = w_t^S \ \tau_t^S \ + \ w_t^G \ \tau_t^G \tag{7}$$

The aggregate trend τ^* is linked to the long-run survey expectations Z^π,

$$Z_t^\pi = C_t^\pi \ + \ \beta^\pi \tau_t^* \ + \ \varepsilon_t^{z\pi}, \ \ \varepsilon_t^{z\pi} \sim N(0, \sigma_{z\pi}^2) \tag{8}$$

$$C_t^\pi = C_{t-1}^\pi \ + \ \varepsilon_t^{c\pi}, \ \ \varepsilon_t^{c\pi} \sim N(0, \sigma_{c\pi}^2) \tag{9}$$

where C_t^π is a time-varying intercept assumed to evolve as an RW process to possibly capture the permanent wedge between the survey estimate and the model-based estimate of trend inflation. This wedge can arise for several reasons, including that model trend inflation is assumed to be the infinite-horizon forecast, whereas the (SPF) survey estimate refers to the ten-year-ahead forecast.

Lastly, the SV processes are modeled as a driftless RW in the log-variance,[7]

$$h_t^{id} = h_{t-1}^{id} + \varepsilon_t^j, \ \ \varepsilon_t^j \sim N(0, \sigma_j^2) \quad id = \{Ser, Goods, \tau S, \tau G\}, \ \ j = \{hS, hG, h\tau S, h\tau G\} \tag{10}$$

2.2 Baseline Model Variants

The equations (1), (2)...(10) define our Baseline model (denoted **M1: Baseline**). We estimate additional model specifications to assess the usefulness and marginal contribution

[7]We explored modeling SV in the measurement equations via a stationary mean-reverting log-volatility process. This specification required estimating four additional parameters, mu and phi, separately, for goods and services inflation. To feasibly estimate the model, we had to employ very tight priors for the parameters mu and phi, and the resulting time-varying estimates of the SV resembled the SV estimates from the RW assumption, which was not surprising, because the estimated rho for both goods and services was above 0.9, i.e., highly persistent; hence, we decided to use the RW assumption, as it entailed fewer parameters. Our finding that rho is estimated to be close to 1 is consistent with other studies, including Chan (2017), who estimates 0.963 (defined as phi; see Table 2, page 22). Our finding of similarity in the SV estimates from RW and stationary specifications aligns with consensus in the literature; e.g., see Eisenstat and Strachan (2016), who document their finding based on examination of several SV processes: "The results suggest either a random walk or stationary model is a practically sensible specification to use for estimating inflation volatility."

of various empirical features incorporated into the Baseline model. These specifications are listed below. Additional model variants that are not central to the paper but useful to investigate a specific curiosity are reported in Appendix A6.

M2: Univariate, Single-Sector. It is a univariate UCSV model similar to the one estimated by Stock and Watson on aggregate inflation data.

M3: Two-Sector. It extends the univariate model to the two main sectors of aggregate inflation: goods and services. Comparing model M3 with M2 provides an assessment of the usefulness of the sectoral data in modeling aggregate inflation.

M4: Two-Sector + Survey. It extends the Two-Sector model by including long-run survey expectations data via equations (8) and (9). Comparing M4 with M3 assesses the value added of including survey data in trend estimation.

M5: Two-Sector + Survey + PC. It extends the Two-Sector model with survey data (M4) by allowing a Phillips curve with a possibly time-varying relationship to capture business cycle effects on the level of services inflation, equation (1). Comparing model M5 with M4 provides an assessment of the importance of adding the PC relationship in explaining inflation dynamics.

M6: Two-Sector + Survey + Uncertainty. It extends the Two-Sector model with survey data (M4) by allowing for a possibly time-varying effect of uncertainty on the level of inflation. For services inflation, the covariate is estimated uncertainty about services inflation; for goods inflation, it is estimated uncertainty about goods inflation. Comparing model M6 with M4 provides an assessment of the usefulness of allowing for the uncertainty (or volatility) feedback on the levels of goods and services inflation, conditional on including survey data.

M7: Two-Sector + Uncertainty. It extends the plain vanilla Two-Sector model (M3) (i.e., without the survey data) by allowing for a possibly time-varying effect of uncertainty on the level of inflation. For services inflation, the covariate is the estimated uncertainty about services inflation; for goods inflation, it is the estimated uncertainty about goods

inflation. Comparing model M7 with M3 provides an assessment of the usefulness of allowing for the uncertainty (or volatility) feedback on the level of goods and services inflation; and comparing model M7 with M6 provides a joint assessment of the usefulness of survey data and uncertainty.

3 Bayesian Estimation and Model Comparison

3.1 Data and Estimation

We estimate the Baseline model and its variants using the following quarterly data: (1) aggregate personal consumption expenditures (PCE) inflation; (2) services PCE inflation; (3) goods PCE inflation; (4) relative share of services in the PCE, that is, nominal share of the PCE of services over the nominal PCE; (5) similarly, relative share of goods in the PCE; (6) the overall unemployment rate; (7) Congressional Budget Office (CBO) estimate of the long-run unemployment rate (which we equate with U*); and (8) long-run inflation expectations of PCE inflation: from 2007 to the present, we use the Survey of Professional Forecasters (SPF) median forecast of annual average PCE inflation over the next 10 years, and prior to 2007, we use the Federal Reserve Board's PTR series. To demonstrate the robustness of the main results to alternative inflation measures, we also collect inflation as measured by the consumer price index (CPI): aggregate CPI inflation, services CPI inflation, goods CPI inflation, their corresponding relative weights, and survey long-run measures of CPI inflation expectations. For all series, the data span the period 1959Q2 to 2025Q2 and are downloaded from Haver Analytics, except for the inflation expectations data, which are downloaded from the Federal Reserve Board's website (PTR) and the Philadelphia Fed's website (SPF).

As the model description shows, we impose many inequality constraints on the latent parameters, and combining them with SV in the errors yields a nonlinear, non-Gaussian state-

space model that makes Kalman Filter based estimation infeasible. However, other methods can solve such nonlinear state-space models; in particular, we use the computationally efficient methods developed by Chan and Strachan (2012) and Chan (2017), which rely on band- and sparse-matrix routines. These routines speed up the Gaussian approximations for the nonstandard conditional posterior densities (of the states). They use a specific Accept-Reject Metropolis-Hastings algorithm (developed in Chan and Strachan (2012)), which better approximates the conditional densities compared to the standard MH algorithm, resulting in significantly more efficient simulation of states.

For the Baseline model and its variants, we simulate 1 million posterior draws from the MCMC sampler. We discard the first 500,000 draws and keep every 100th of the remaining draws. The reported results are based on 5000 retained draws.

3.2 Marginal Likelihood and the Posterior Odds Ratio

We use Bayesian model comparison based on the marginal likelihood criterion to compare the fit of the Baseline model and its variants, assessing the importance of the various empirical features in the Baseline model. Following in the footsteps of a number of papers (e.g., Amisano and Geweke (2010); Chan et al. (2018); Eo et al. (2023); Carriero et al. (2024); and Zaman (2024)), marginal likelihood or marginal data density (MDD) is computed as the product of one-step-ahead predictive likelihoods or equivalently as the sums of the log of one-step-ahead predictive likelihoods. This approach allows us to compute marginal data density separately for services and goods inflation, as well as separately for various subsamples.

Specifically, if we let $y^j_{1:t} = (y^j_1, ...y^j_t)'$ denote the data for the variable of interest, j, up to time t, then we can write the marginal data density of the variable of interest j as,

$$p(y^j|X^j_i, M_i) = \prod_{t=3}^{T} p(y^j_t|y^j_{1:t-1}, X^j_{1:t,i}, M_i) \tag{11}$$

where j = {Services inflation, goods inflation}, $p(y^j_t|y^j_{1:t-1}, X^j_{1:t,i}, M_i)$ is the predictive likeli-

hood for variable j, and X_i^j is a set of covariates that influence variable j in model M_i.

The MDD $p(y^j|X_i^j, M_i)$ can be seen as a density forecast of the data computed using model M_i and evaluated at the actual observed realization y^j. The higher the MDD, the more accurate the density forecast.

As discussed in Koop (2003), when comparing two models M_1 and M_2, if the MDD of M_1 is larger than that of M_2, then the observed actual realizations are more likely under model M_1 compared to model M_2. This is then viewed as evidence in favor of model M_1, and the strength of this evidence can be quantified by the **posterior odds ratio** between the two models,

$$\underbrace{\frac{\mathbb{P}(M_1|y)}{\mathbb{P}(M_2|y)}}_{\text{posterior odds ratio}} = \underbrace{\frac{\mathbb{P}(M_1)}{\mathbb{P}(M_2)}}_{\text{prior odds ratio}} \times \underbrace{\frac{p(y|M_1)}{p(y|M2)}}_{\text{Bayes factor}}$$

The posterior odds ratio equals the Bayes factor when the prior odds ratio is 1, implying that, a priori, both models are equally likely. A Bayes factor of 20 would mean that, given the data, model M_1 is 20 times more likely than model M_2.

3.3 Savage-Dickey Density Ratio and Dynamic Posterior Probabilities

The additional model specifications defined in Section 2.2 are restricted variants of the Baseline model; that is, they are nested models, which makes model comparison—assessing empirical support for the restriction defining the model—using the Savage-Dickey density ratio (SDDR) relatively straightforward. It is convenient because the SDDR is directly related to the Bayes factor (BF) and can be computed from the posterior output (densities) of the unrestricted model only (e.g., the Baseline model in our case), so it does not use the output from the restricted model. As shown by Koop et al. (2010), when working with models with time-varying parameters, the BF between two models, or the posterior odds ratio (PO),

whose computation is facilitated by the SDDR, conveniently permits the computation of (dynamic) posterior probabilities – the posterior probability at each time t – for the specific restrictions defining model specifications. Accordingly, we use dynamic posterior probabilities as an additional metric to assess the empirical support of various features implemented in the Baseline model.

To fix ideas, we illustrate the computation of the SDDR and the BF using general notation. Suppose the unrestricted model (e.g., the Baseline model in our case), Model_UR, has a parameter vector $\Theta = (\omega', \phi')'$, where ϕ collects parameter(s) that are unrestricted in both models being compared, and ω collects parameter(s) that are unrestricted in the unrestricted model but are restricted in the restricted model. The likelihood and prior for the unrestricted model are given by $p(y|\omega, \phi, Model_UR)$ and $p(\omega, \phi|Model_UR)$.

The restricted model, Model_R, has $\omega = \omega_0$, where ω_0 is a vector of constants (e.g., $\alpha_t^S = 0$ for all periods $t = 1$ to T). The parameters in vector ϕ are left unrestricted in each model. The likelihood and prior for this model are given by $p(y|\phi, Model_R)$ and $p(\phi|Model_R)$. Since ω is simply equal to ω_0 under Model_R, we do not need to specify a prior for it.

Suppose the priors in the two models satisfy:

$$p(\phi \mid \omega = \omega_0, Model_UR) = p(\phi \mid Model_R)$$

Then $BF_{Model_R_to_Model_UR}$, the Bayes factor comparing Model_R to Model_UR, has the form

$$BF_{Model_R_to_Model_UR} = \frac{p(\omega = \omega_0 \mid y, Model_UR)}{p(\omega = \omega_0 \mid Model_UR)}$$

where $p(\omega = \omega_0 \mid y, Model_UR)$ and $p(\omega = \omega_0|Model_UR)$ are the unrestricted posterior and prior for ω evaluated at the point ω_0. The numerator, i.e., the posterior density, can be

estimated and often is using the Monte Carlo average

$$\frac{1}{D}\sum_{d=1}^{D} p(\omega = \omega_0 \mid y, \phi^d, Model_UR)$$

where D represents the total number of posterior draws.

$BF_{Model_R_to_Model_UR}$ is also called the SDDR (see Verdinelli and Wasserman (1995), who prove that the BF equals the SDDR). Furthermore, assuming a prior odds ratio of one – a priori, both models, restricted and unrestricted, are equally probable – the $BF_{Model_R_to_Model_UR}$ equals PO (posterior odds ratio in favor of the restricted model).

Given the posterior odds ratios, the (dynamic) posterior probabilities of the restriction $\mathbb{P}(\omega_t \neq \omega_0 \mid y)$ for each time period t are computed as,

$$\mathbb{P}(\omega_t \neq \omega_0 \mid y) = \frac{1}{1 + PO_t}$$

For our empirical analysis, we compute the dynamic posterior probabilities for the following restrictions, $\mathbb{P}(\alpha_t^S \neq 0 \mid y)$, $\mathbb{P}(\alpha_t^G \neq 0 \mid y)$, $\mathbb{P}(-1 < \lambda_t < 0 \mid y)$, and $\mathbb{P}(\beta \ \& \ C_t \neq 0 \mid y)$

3.4 Pseudo-Out-of-Sample Forecasting

A common practice in the literature when introducing a new model for trend inflation estimation is to compare its out-of-sample forecasting performance with that of the univariate Stock and Watson model, which is widely regarded as "hard to beat."

Accordingly, we assess the forecasting performance of the Baseline model (model M1) and its variants for aggregate inflation and compare each to the univariate Stock and Watson benchmark model, denoted UCSV in the reported results.

We assess both point and density forecasts for each model. We evaluate point forecasts, computed as the posterior mean of the density forecasts, using the root mean squared forecast error (RMSFE) metric. The statistical significance of gains in the accuracy of point

forecasts between the candidate models is assessed using the Diebold-Mariano test (with the Newey–West correction), based on a two-sided standard normal test. We evaluate the density forecasts using the logarithmic predictive score and assess the statistical significance of gains using the likelihood-ratio test of Amisano and Giacomini (2007), with a two-sided t-test. The forecasting evaluation uses a recursively expanding estimation window. For all models considered, the estimation sample starts in 1959Q3, and the recursive forecast evaluation runs from 1999Q1 to 2025Q2. At each recursive run, forecasts of average inflation are produced up to three years out (i.e., the forecast horizon, h, ranges from h=1 to h=12 quarters ahead), but for brevity, we report results for h=2Q, h=4Q, h=8Q, and h=12Q. Following Stock and Watson (2016), Mertens (2016), and Eo et al. (2023), among others, we forecast average inflation over the forecast horizon. Specifically, (similar to Eo et al. (2023)) if we define the quarter-over-quarter annualized inflation rate as $\pi_t = ((\frac{P_t}{P_{t-1}})^4 - 1) * 100$, where P_t is the quarterly price level, then the h-step-ahead forecast generated at time t is for the object (average inflation over h periods): $\frac{1}{h}\sum_{i=1}^{h} \pi_{t+i}$. So, forecasts for $h = 4Q$, $h = 8Q$, and $h = 12Q$ are four-quarter growth rates one year ahead, two years ahead, and three years ahead, respectively.

4 Empirical Results

This section starts with the presentation of estimates of the trend in aggregate, services, and goods inflation inferred from the Baseline model. It then assesses the marginal contributions of sectoral data, survey information, uncertainty (SVM component), and the Phillips curve relationship in explaining model fit to the inflation data, out-of-sample forecasting of aggregate inflation, and the impact on trend estimates of inflation rates. The assessment relies on extending the univariate (Stock and Watson, UCSV) model one-by-one with the features of interest: Univariate –> Two-Sector –> Two-Sector+Survey –> Two-Sector+Survey+Uncertainty and Two-Sector+Survey+PC –> Two-Sector+Survey+Uncertainty+PC

(which is the Baseline). Table 1 reports the posterior estimates of the Baseline parameters, Table 2 reports the marginal data density (MDD) for the various models, and Tables 3 and 4 report the out-of-sample forecasting performance for the various models.

4.1 Trend Inflation: Baseline Model

Figure 2 presents the estimates of the trend in services inflation (panel a), in goods inflation (panel b), and in aggregate inflation (panel c) inferred from the (two-sector) Baseline model. The model indicates that the trend in aggregate inflation was low in the 1960s; high in the 1970s, peaking at 7.6 percent in early 1981, as the model inferred significant increases in inflation as mostly permanent; fell sharply in the 1980s; continued a steady deceleration in the 1990s; fluctuated in a narrow range between 1.8 percent and 2.1 percent in the 2000s; increased gradually during the post-pandemic inflation surge to reach 2.45 percent by the end of 2022; and since then moderating slightly, reaching 2.1 percent by 2025Q2.

The two periods, 2012-2019 and post-COVID, were of particular interest for inflation dynamics. In the first period, 2012 to 2019, when inflation remained stubbornly below 2 percent, trend inflation is estimated to be stable just below 2 percent. As shown by the historical decomposition in Figure 7, the model explains realizations of low inflation partly by labor-market slack over most of this period (a positive unemployment rate gap) and partly by lower uncertainty.[8] During the post-COVID inflation surge, when inflation peaked closer to 8 percent, trend inflation is estimated to rise only modestly, by about 0.5 ppts (from 1.9 percent to 2.4 percent). The model inferred that the sharp spike in inflation was mostly transitory, explained mainly by the uncertainty component associated with the increasing volatility of the (positive) shocks to both services and goods inflation (see Figure 5, panels a and b; and the blue bars in Figure 7), and a small portion of the spike is explained by a tight labor market (green bars in Figure 7, panels a and b).

[8] See Appendix A5, Figure 9, which plots the CBO unemployment rate gap.

4.2 Importance of Sectoral Data

A growing body of research has illustrated the usefulness of modeling aggregate trend inflation via sectoral data, as doing so leads to more precise inference about the trend and some evidence of improved gains in forecasting inflation; see, for example, Stock and Watson (2016), Tallman and Zaman (2017), and Eo et al. (2023), among others. Consistent with prior research, extending the univariate (single-sector) Stock and Watson (SW) model to a two-sector model of goods and services inflation, where each sector is modeled similarly to the univariate process of SW, we find improved fit of the model to inflation data, comparable point forecast accuracy, but more accurate density forecasts in the medium-term horizon.[9]

The improved model fit to aggregate inflation is evident when comparing the MDD of the models denoted "Univariate" and "Two-Sector" in Table 2, bottom panel. MDD based on the full sample is -398.8 for the Univariate model and -359.4 for the Two-Sector model. Table 3 shows evidence that the Two-Sector model has competitive point-forecast accuracy relative to the Univariate model.[10] However, across the two forecast evaluation samples, the Two-Sector model generates more accurate inflation forecasts than its Univariate counterpart in the post-COVID period, as the relative RMSFEs for the full sample (including post-COVID data) are smaller than those for the sample ending in 2019. Because the post-COVID sample is a much smaller fraction of the overall sample, i.e., contains fewer observations, the gains in accuracy are statistically not significant.[11]

Table 4 also shows that as the forecast horizon increases, the Two-Sector model produces more accurate density forecasts than the univariate model. Even though the estimates of the aggregate trend inflation across the two models are similar through most of the sample

[9] Stock and Watson (2016) and Eo et al. (2023) considered a point forecast accuracy comparison only, whereas Tallman and Zaman (2017) considered both point and density forecast performance.

[10] The finding of competitive point-forecast accuracy is consistent with previous research that found no consistent benefit of aggregating sector-level forecasts to forecast aggregate inflation compared with approaches that forecast aggregate inflation directly (see Clark et al. (2025) and references therein).

[11] The finding that the Two-Sector model does better than the univariate model is consistent with Watson (2026), who documents the superior forecast accuracy of the sectoral UCSV model relative to that of the univariate SW model in tracking the post-COVID inflation surge.

period (shown only for the Two-Sector model in Figure 3), the Two-Sector model produces more precise estimates, that is, tighter credible intervals (sharper inference), which translates into more precise density forecasts that fit the data better. This suggests that the sectoral model's increased parameter uncertainty is outweighed by the additional information it brings for signal extraction. This finding of more precise trend estimates from the use of sectoral data is consistent with Stock and Watson (2016) and Eo et al. (2023).

The advantages of using sectoral data become even more apparent when the model includes other elements, as is the case with our Baseline model. Given that the dynamics of goods inflation differ materially from those of services inflation, prior research (e.g., Tallman and Zaman (2017)) argues that key factors influencing goods inflation can and do differ from those affecting prices in the services sector. Specifically, services inflation is found to be correlated with labor-market conditions (i.e., the Phillips curve relationship), whereas goods inflation shows a weaker correlation. This suggests a weaker Phillips curve relationship in a model that does not use sectoral data and instead directly links aggregate inflation to labor-market conditions. As discussed later, consistent with prior research, our model provides evidence of a statistically significant (time-varying) Phillips curve relationship for services inflation.

Our results for the Baseline model also find that the estimated relationship between services inflation and uncertainty differs from that between goods inflation and uncertainty. A model that does not use sector-level data and incorporates uncertainty and Phillips curve relationships for aggregate inflation will likely yield different parametric inferences (e.g., a weaker Phillips curve because it combines services inflation and goods inflation) and, in turn, different inferences about trend inflation. In fact, this is confirmed by comparing the Baseline model's estimate of trend inflation to a single-sector model with all the bells and whistles but no sectoral data (Figure 13, Appendix 6.2): it is consistently a few tenths of a percentage point higher than the Baseline model's estimate from 1990 onward. The difference between the two estimates was notable between 2000 and 2007, as the Baseline model implied a

stable trend of 1.8 percent, whereas the single-sector model implied a trend of 2.3 percent. These trend estimates imply different policy prescriptions for a central bank with an inflation target of 2 percent: a trend inflation of 1.8% calls for accommodative policy, whereas a trend inflation of 2.3% calls for restrictive policy.[12]

4.3 Usefulness of Survey Expectations

Recent research has demonstrated the importance of incorporating survey-based long-run inflation expectations into unobserved components models used to estimate long-run trend. Including survey information, which reflects the credibility of policymakers' commitment to the inflation target, yields more reasonable estimates of the model-based trend (e.g., Mertens (2016); Chan et al. (2018); Zaman (2024)) and improves the accuracy of inflation density forecasts. These benefits are realized partly because surveys have been shown to adapt more quickly to perceived structural changes in the economy (e.g., Chan et al. (2018)). Furthermore, including survey data makes it tractable to estimate high-dimensional models during periods of high uncertainty (e.g., Zaman (2024)) or when working with small sample sizes, as discussed later.

Our results provide strong visual evidence supporting the assessment that models with survey data yield reasonable estimates of (long-run) trend inflation, confirming prior research findings. As shown in Figure 3, the model denoted "Two-Sector + Survey," which extends the Two-Sector model by including survey data, generates a trend that is much less volatile and less responsive to arguably large transitory spikes compared to the Two-Sector model without survey data. Comparing trend estimates from the model with and without survey data, trend inflation is similar in the 1960s, but beyond that the differences are sizable: the model without survey data implies trend inflation estimates (services, goods, and aggregate)

[12]Furthermore, the model comparison shows that the Baseline model fits the historical data better than the Single-Sector model, which has the same features as the Baseline but excludes sectoral data (single-sector model, M13): MDD of -398.0 for the Baseline model vs. -415.8 for the Single-Sector model, shown in Table 3 in Appendix A6.

that are highly volatile. During the post-COVID inflation surge, the model without the survey data attributes most of the increase in inflation to the trend component, as the aggregate trend is estimated to have increased quite abruptly and sharply by roughly 4.3 percentage points, from 1.5 percent to a level closer to 6.0 percent, driven by sharp increases in the trend in services inflation, which rises from 2.0 percent to 5.6 percent, and the trend in goods inflation, which rises from 0.4 percent to 6.5 percent.

The stabilizing role of survey data is also evident in the stable trend estimate from the Baseline model (Figure 1) and the less volatile estimates from model variants that include survey data (Figure 4).

The parametric inference corroborates the visual evidence of the survey data's strong influence on trend inflation estimation. The parameter β^{π}, which together with parameter C^{π} governs the relationship between the Baseline model's aggregate trend and survey expectations, is estimated to be 0.93 (posterior mean) with 90% credible intervals spanning 0.85 to 1.01 (reported in Table 1). The estimates of the time-varying intercept C^{π} (see Figure 6, panel d) indicate evidence of time variation in the relationship between the long-run survey expectations and the model's trend inflation. Specifically, between 1960 and 1990, there are sizable deviations between the estimated trend and survey expectations; however, since 1990, the gap between the two has diminished, as captured by the parameter C^{π}, such that it is not statistically different from zero, as evidenced by the wide credible bands around the posterior mean of C^{π}.

Interestingly, and somewhat consistent with prior research on fitting PCE inflation data, adding survey data worsens the model fit, as shown in Table 2 comparing the MDD of the "Two-Sector + Survey" model with the MDD of the "Two-Sector" model. A deeper examination of the time series of one-step-ahead predictive likelihoods (i.e., our measure of model fit) suggests that adding survey data improves the precision of the trend estimates, yielding predictive densities that are tighter than those of the model variant without survey data, which performs worse during periods of highly volatile inflation. During periods when

inflation is relatively stable (e.g., 1990 to 2000), the Baseline model with survey data generates more accurate predictive densities and therefore fits the inflation data better than model variants without survey data.[13]

Focusing beyond near-term predictive accuracy (used to compute MDD), the pseudo-out-of-sample forecasting results reported in Table 3 suggest that including survey data yields competitive point forecast accuracy, as evident by comparing the models denoted "Two-Sector" and "Two-Sector + Survey," which show relative ratios that are not statistically significant. However, for density forecast accuracy, including survey data yields significantly more accurate density forecasts as the forecast horizon increases. The more accurate density forecasts for the model with survey data stem from improved precision in the trend estimates and, in turn, sharper density forecasts that fit the data better.[14] The finding that adding survey data yields competitive (but not better) point forecasts and more accurate density forecasts is consistent with Chan et al. (2018).

Lastly, we end this sub-section by illustrating the importance of survey data when estimating the model(s) on shorter sample periods, which could be due to choice (e.g., 1980s and onward, as is sometimes done in papers focusing on US inflation given the well-documented structural change) or data limitations (as is the case for the Euro area, where data are not

[13]See Table 5, page 31 in Chan et al. (2018); their model M1 (with survey data) has an MDD of −367.28 compared to −366.35 for their model M4 (without survey data). However, it is worth pointing out that in the Chan, Clark, and Koop paper, when the models are estimated with data starting in 1980, which the authors do for CPI inflation, the fit of the model when measured over the period 1980 to 2016 is better than the model without survey data; see Table 1 on page 21 of their paper. In our supplementary exercises, we find that when models are estimated starting in 1980, the difference in fit between models with and without the survey data is much smaller. However, for models estimated with post-1980 data, if we implement tight priors similar to Chan et al. (2018) and add an AR(1) term, we find that adding survey data improves model fit. That said, the data prefer less tight priors rather than tighter priors. A detailed supplementary examination shows that the influence of survey data on model fit depends on several factors, including the estimation sample, model formulation, and prior specification (see Appendix A11). For more results, including a model fit comparison for univariate models extended to include an AR(1) term and/or survey data similar to Chan et al. (2018), refer to Appendix A11.

[14]As shown by previous research, and by this paper, bringing long-run survey expectations data, which could be viewed as a judgment-based proxy of the long-run trend, into the estimation yields more precise estimates of the model-based trend and, in turn, predictive densities of inflation with tighter credible intervals compared to the model without survey data. The sharper density estimates fit the inflation data better when inflation is relatively stable, as was the case for most of the sample period spanning the evaluation sample.

available prior to 1990). To do so, we re-estimate our Baseline model (M1) and additional model variants with and without survey data using data spanning 1980 through 2015Q2. The plots of the resulting trend inflation estimates are shown in Appendix A10. The results indicate that trend inflation estimates from model variants without survey data are highly volatile and imprecise, whereas those from models with survey data evolve smoothly, as expected for a long-horizon trend estimate, and thus appear sensible. Furthermore, as shown in Table 4 of Appendix A10, the differences in model fit between models with and without survey data, when estimated using a shorter sample, are smaller than those in Table 2, highlighting the role of the estimation period in the model-fit comparison.

If we were to estimate our model with an even shorter window, i.e., with data starting from 1990 onward, the variant of the Baseline model with all features but without survey data has estimation issues, highlighting the important role of survey data not only in providing sensible estimates of trend but also in helping feasibly estimate the richer specifications of unobserved components models when faced with data limitations.

4.4 Role of Uncertainty Effects

Extending the Two-Sector model with survey data to include the uncertainty (SVM) component increases the fit of the model to inflation data and substantially improves short-term out-of-sample forecasting accuracy in predicting inflation. The improved fit is evident when comparing MDD between the model specification denoted "Two-Sector + Survey" and the model specification "Two-Sector + Survey + Uncertainty" in Table 2. The inclusion of the uncertainty component also improves fit to inflation data in a model without survey data, as shown by comparing the models "Two-Sector" and "Two-Sector + Uncertainty," consistent with the results of Chan (2017).

Regarding out-of-sample forecast accuracy, adding SVM improves short-term predictive accuracy whether or not the survey data are included. This is evident in Table 3 by com-

paring the models "Two-Sector + Survey" and "Two-Sector + Survey + Uncertainty" or by comparing "Two-Sector" and "Two-Sector + Uncertainty." For example, the two-sector model augmented with survey data yields two-quarters-ahead forecasts that are 8 percent less accurate on average than the univariate model, but adding uncertainty (SVM term) to this model yields forecasts that are about 18 percent more accurate than the univariate model for h=2Q, and the gains are statistically significant.

However, the usefulness of uncertainty (SVM) for the medium-term forecast horizon depends on whether the model includes survey data. In a model that excludes survey data, adding the SVM component yields inflation forecasts as accurate as those of the univariate model and the basic Two-Sector model, as shown in Tables 3 and 4. In contrast, for a model that includes survey data, adding the SVM component (model, "Two-Sector+Survey+Uncertainty") worsens medium-term forecast accuracy, i.e., forecast horizons beyond four quarters, particularly density forecasts. This worsening is partly because inclusion of the SVM component widens the width of the predictive densities more than would be considered reasonable and partly because the model with survey data is estimated to have a much stronger response to spikes in uncertainty associated with goods inflation, resulting in a higher inflation (point) forecast trajectory for longer than is optimal. Without the survey data, the estimated trend moves around a lot, explaining a greater share of the variation in the inflation data, leaving less of the inflation variation attributed to the uncertainty component. Hence, the smaller response of inflation to changes in uncertainty than in the model with survey data, which pins down the trend component to reasonable and stable levels, leaving a greater share of variation to be explained by the SVM component.

As discussed in the next section, the addition of the PC relationship improves the forecast accuracy for the medium-term forecast horizons, so adding the PC relationship to a Two-Sector model that has uncertainty and survey data (which is the Baseline model) yields forecasts at the medium-term forecast horizon that are more accurate than those from the model "Two-Sector+Survey+Uncertainty," as can be seen in Table 2 comparing "Baseline"

and "Two-Sector+Survey+Uncertainty." Thus, adding the PC relationship dampens the uncertainty component's adverse influence on medium-term inflation forecasts.

Figure 5 (top row) presents uncertainty estimates, measured by the stochastic volatility estimated in the Baseline model. Panel (a) shows uncertainty estimates for services inflation (standard deviation of the shocks to services inflation) and panel (b) for goods inflation. Shown are the posterior median uncertainty estimate and the 90% credible bands, which represent the uncertainty around the uncertainty estimate.[15]

According to the estimates, for both services and goods inflation, uncertainty is estimated to be high during the Great Inflation period of the 1970s, recedes during the Great Moderation period, rises slightly for services inflation but more sharply for goods inflation following the Global Financial Crisis, then subsides until it rises again quite sharply during the COVID pandemic and the subsequent inflation surge. More recently, estimated uncertainty (stochastic volatility) for both services and goods inflation has moderated, but substantial uncertainty remains around the uncertainty estimates, as shown by very wide credible bands around the posterior median estimates of stochastic volatility.

The broad contours of the uncertainty estimates shown in Figure 5 are consistent with those reported in the literature on inflation; e.g., see Figure 1 in Eisenstat and Strachan (2016). However, a closer look at uncertainty estimates across baseline model variants shows that modeling choices affect uncertainty estimates (stochastic volatility), as shown in Appendix A7. The comparison suggests that survey data have the most impact, followed by the inclusion of the uncertainty channel (volatility feedback). These differing estimates of uncertainty, in turn, affect the parameter estimates that capture the effects of uncertainty on services and goods inflation; see, for example, Appendix A9 for a comparison of the parameters α^S and α^G between the Baseline model and the model without the survey data.

[15]It is worth mentioning that the ability to quantify uncertainty around the uncertainty estimates is one of the main reasons the literature arguably favors the approach of using stochastic volatility as an uncertainty measure for estimating the impact of uncertainty on variables of interest (see Carriero et al. (2018), Cross et al. (2023), and Canova and Forero (2024)).

Figure 6, panels (a) and (b), report the estimates of the uncertainty parameters α^S and α^G. These estimates indicate substantial time variation in the effects of uncertainty, particularly on the level of goods inflation. For example, in the case of goods inflation, prior to 1990, the estimates indicate a strong positive association between uncertainty and its effect on goods inflation, whereas this association becomes much smaller after 1990, goes to zero following the Global Financial Crisis through 2019, and then increases but remains modest compared to pre-1990. The smaller response of goods inflation to higher inflation uncertainty post-1990 could reflect anchored inflation expectations (captured via the long-run survey data). Overall, the evidence points to a positive impact of the estimated uncertainty on the levels of both goods and services inflation. Our finding that higher uncertainty positively affects inflation is consistent with the predictions of theoretically motivated models emphasizing the positive effects of cost-side uncertainty channels on inflation (e.g., Born and Pfeifer (2014)) and of monetary policy uncertainty on inflation (e.g., Ball (1992)), and empirical studies Chan (2017) and Acharya et al. (2026).

The dynamic posterior probabilities shown in Figure 8, panels (a) and (b), reinforce the importance of allowing for time variation in the parameters governing the uncertainty effects on goods and services inflation. For example, for the period 2003-2010, it is highly likely that $\alpha^S = 0$, whereas for 1965-1985 and 2012 to the present, the data strongly favor $\alpha^S \neq 0$. Similarly, for goods inflation, the data strongly favor $\alpha^G \neq 0$ for 1960-1990 and 2020Q4-2023Q1, whereas for 2005-2019, the data strongly favor $\alpha^G = 0$.

The inferences implied from dynamic posterior probabilities are in agreement with the posterior estimates of α^S and α^G plotted in Figure 6. For example, panel (a) of Figure 6 indicates that for the 1960s and for the period 2003-2010, the credible bands for the parameter α^S include zero, providing informal evidence in favor of the restriction $\alpha^S = 0$, as also confirmed formally by the dynamic posterior probability.

Figure 7 plots the model-based decomposition of the historical inflation data. As is evident, the uncertainty component plays an important role in explaining fluctuations in the

inflation data. According to the model, the uncertainty component accounts for a large share of fluctuations in the inflation data, particularly over the past five years. Interestingly, the model sees (estimated) uncertainty explaining a greater share of variation in services inflation than the business cycle component.[16]

4.5 Role of Phillips Curve

The addition of the PC relationship has a trivial impact on the estimate of aggregate trend inflation, as shown in Figure 4 by the solid green plot labeled "Two-Sector+Survey+PC," which builds on the "Two-Sector + Survey" model (plotted in red) by adding the PC relationship. However, under the hood, this trivial impact masks significant offsetting effects in the estimates of trends in services and goods inflation. For example, between the late 1960s and the early 1970s, the estimate of the services trend from the model with PC was 1.3 to 1.5 ppts lower than that inferred from the model without PC (Model "Two-Sector+Survey"), but the trend in goods inflation during this period was 1.0 to 1.3 ppts higher, leaving the estimated trend in aggregate inflation across the two models little changed. Similarly, between 2010 and 2012, the estimate of the services trend from the model with PC was about 0.6 ppt higher than the model without PC, while for goods inflation it was roughly 1.00 ppts lower, leaving the estimated trend in aggregate inflation nearly identical across the two models.

Bayesian model comparison provides support for the time-varying PC slope parameter. As reported in Table 2, the model denoted "Two-Sector + Survey + PC" fits inflation data better than the model without PC ("Two-Sector + Survey"). Furthermore, the data prefer a time-varying PC relationship rather than a time-invariant one; as shown in Appendix A6.2, the model that removes time variation in the PC slope has a worse fit to inflation data.

[16]It is worth mentioning that our estimates of uncertainty effects likely reflect effects beyond those of "true" inflation uncertainty because the measured uncertainty could reflect effects not captured directly by the model arising from factors like, for example, food and energy shocks during the 1970s, globalization and low import price inflation during the 1990s and 2000s, and a broken supply chain (during the post-pandemic inflation surge). However, to the extent the effects of these omitted factors on inflation are correlated with "true" inflation uncertainty, which they appear to be, as evidenced by the size of the blue bars during these specific periods, it is reasonable to assume the estimated effects are less biased.

The out-of-sample forecasting results reported in Tables 3 and 4 also support including the PC relationship, as shown by small gains in both point and density forecast accuracy when comparing the relative RMSFEs for the models "Two-Sector + Survey" and "Two-Sector + Survey + PC" in both evaluation samples. The usefulness of adding the PC relationship can also be seen when comparing the models "Two-Sector+Survey+Uncertainty" and "Baseline" (which extends the former model by bringing in the PC relationship). The relative RMSEs for the row denoted "Baseline" are smaller, and the relative log score for density accuracy is notably larger than the relative measures for "Two-Sector+Survey+Uncertainty." The forecast evaluation results for the sample ending in 2019 are consistent with those documented by Tallman and Zaman (2017).

Figure 6, panel (c), presents the Baseline model's posterior estimates of parameter λ, the slope of the price Phillips curve. The estimates indicate strong evidence of time variation in the slope of the Phillips curve. For example, the model infers a steeper Phillips curve in the 1960s and 1970s, which subsequently weakens (becomes less negative) through 2000. From 2000 through the Great Recession of 2007-2009, it steepens; thereafter, it weakens for the next few years until 2015, when it gradually steepens. As of 2025Q2, the posterior mean of the PC slope is estimated at -0.4, with considerable uncertainty, as the 68% credible interval spans -0.15 to -0.6.[17]

The model-based decomposition of inflation from the Baseline model shown in Figure 7 indicates that the business cycle effects (computed as a product of the PC slope parameter and the unemployment rate gap) on inflation are estimated to be small, except during periods when the unemployment gap is very positive (during deep recessions) or very negative (as between 1965 and 1970); see the plot of the unemployment rate gap in Appendix A5. The periods when business cycle effects played a nontrivial role in explaining inflation fluctuations have generally coincided with recessions and the first year or two of the expansions that

[17] A model variant that removes time variation in the PC slope is estimated with a posterior mean of -0.3; see Appendix A8, Figure 16, Panel (c).

followed them. In the past couple of years, the model shows that business cycle effects have positively influenced inflation, but this influence pales in comparison to uncertainty effects.

4.6 Summary

Overall, our examination reinforces the evidence also documented elsewhere that modeling choices notably influence inferences about the inflation trend. Consistent with the literature, which has established generally competitive forecast accuracy (in predicting inflation) for the proposed models (some of them referenced in Figure 1) relative to the univariate Stock and Watson model, we also find generally competitive forecasting accuracy for the Baseline model and many of its restricted variants. On the one hand, this is encouraging; on the other, it highlights the difficulty of choosing a model solely based on forecast evaluation from a pool of competitive candidates, a point Clark and Doh (2014) also emphasize. In our view, the Baseline model developed in this paper, which incorporates many of the elements shown to be empirically useful by prior research, should, in principle, better isolate the trend component from other fluctuations, providing a more credible read of the long-run trend.

5 Conclusion

This paper proposes a new empirical model to estimate the (long-run) trend in inflation. The new model draws on prior research that has advanced the popular Stock and Watson (2007) model of trend inflation by introducing several innovations. But these innovations have generally been explored in isolation. This paper brings those innovations into a single model. Specifically, the paper's proposed model extends the Stock and Watson UCSV model by incorporating the following four elements: uncertainty effects on inflation (volatility feedback), business cycle forces (the Phillips curve relationship), survey expectations of long-run inflation (forward-looking feature), and sectoral data (goods and services) in teasing out inflation fluctuations associated with the (long-run) trend component. Using various model

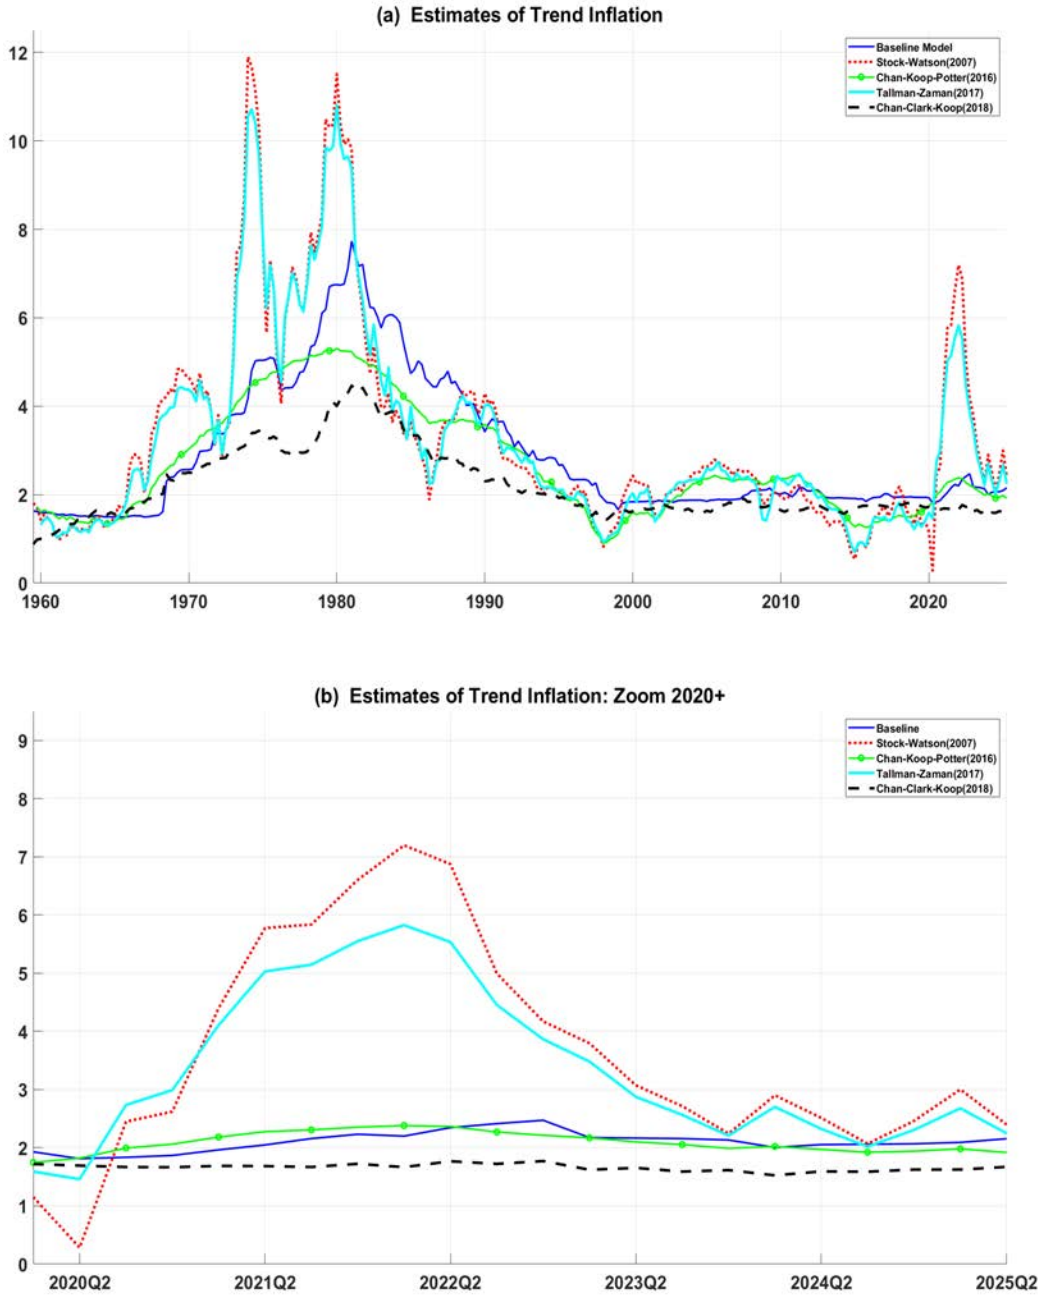
comparison metrics, the paper assesses the combined and individual roles of these elements in estimating trend inflation. This assessment compares the model with all four elements to the model variants that exclude one or more of them.

Our results indicate the importance of jointly modeling all four elements, with some elements being more informative for estimating trend inflation than others. Survey expectations have the greatest impact on trend estimates, followed by uncertainty and sector-level data, while the Phillips curve has the least impact.

Including survey data is important for obtaining sensible estimates of the trend, especially when estimating models with a shorter sample window. Results also indicate the importance of allowing for the time-varying effects of estimated uncertainty on the levels of goods and services inflation. The parametric inference suggests a strongly positive relationship between inflation and inflation uncertainty. The model comparison metric favors modeling trend inflation using sector-level data over using aggregate inflation data alone.

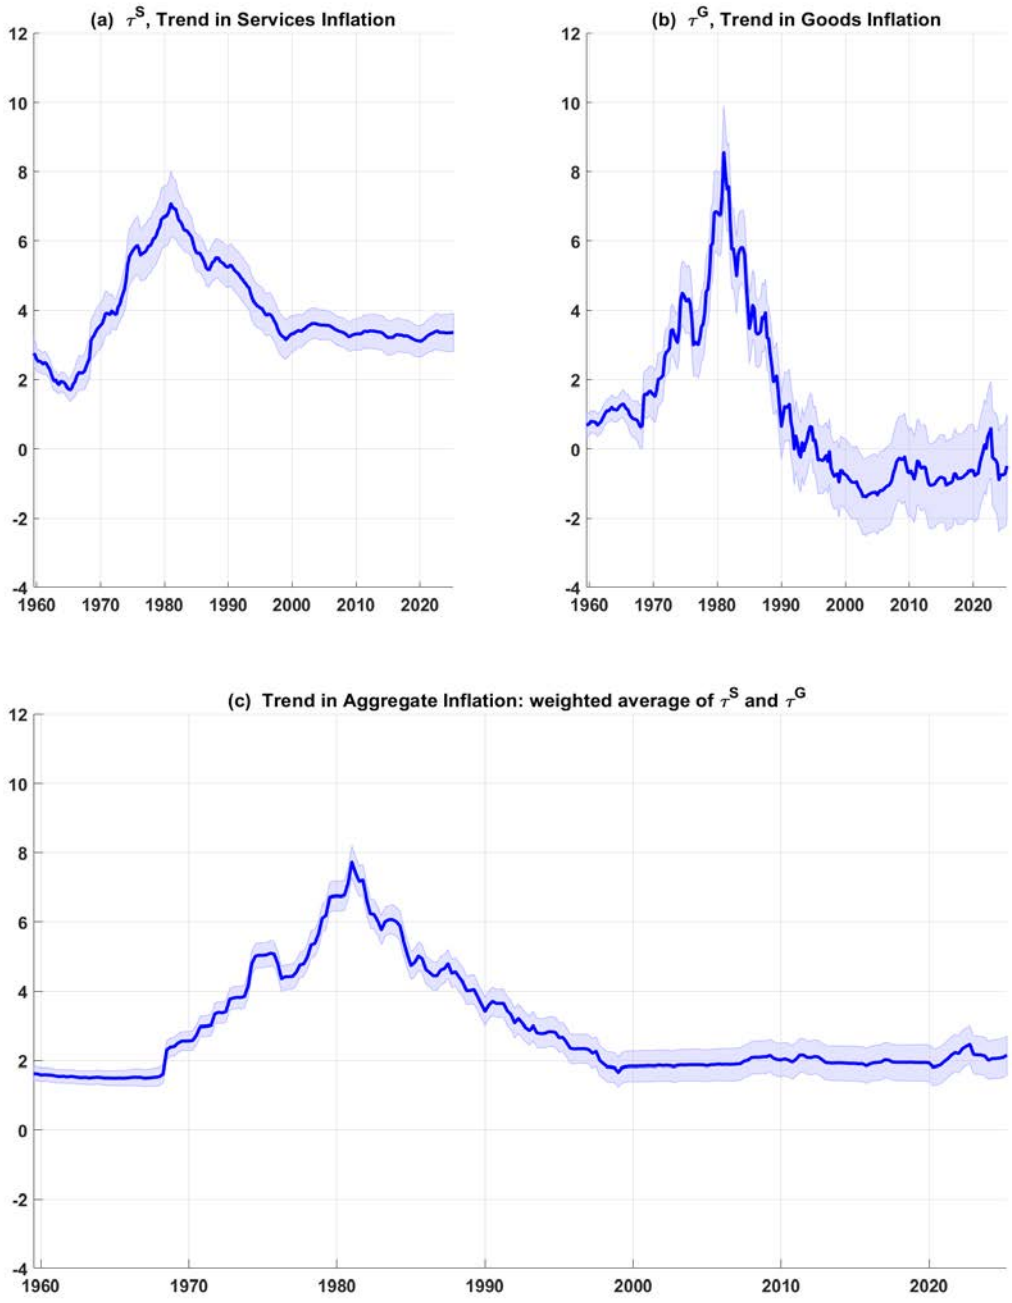
The out-of-sample forecast comparison indicates that the new model's forecast accuracy is competitive with but not superior to the univariate Stock and Watson model for predicting future inflation. Lastly, the new model provides a more interpretable historical decomposition of inflation data than the models it extends. The decomposition suggests that uncertainty effects play a greater role than cyclical effects in explaining inflation fluctuations.

Figure 1: Trend Inflation: Comparison with External Models



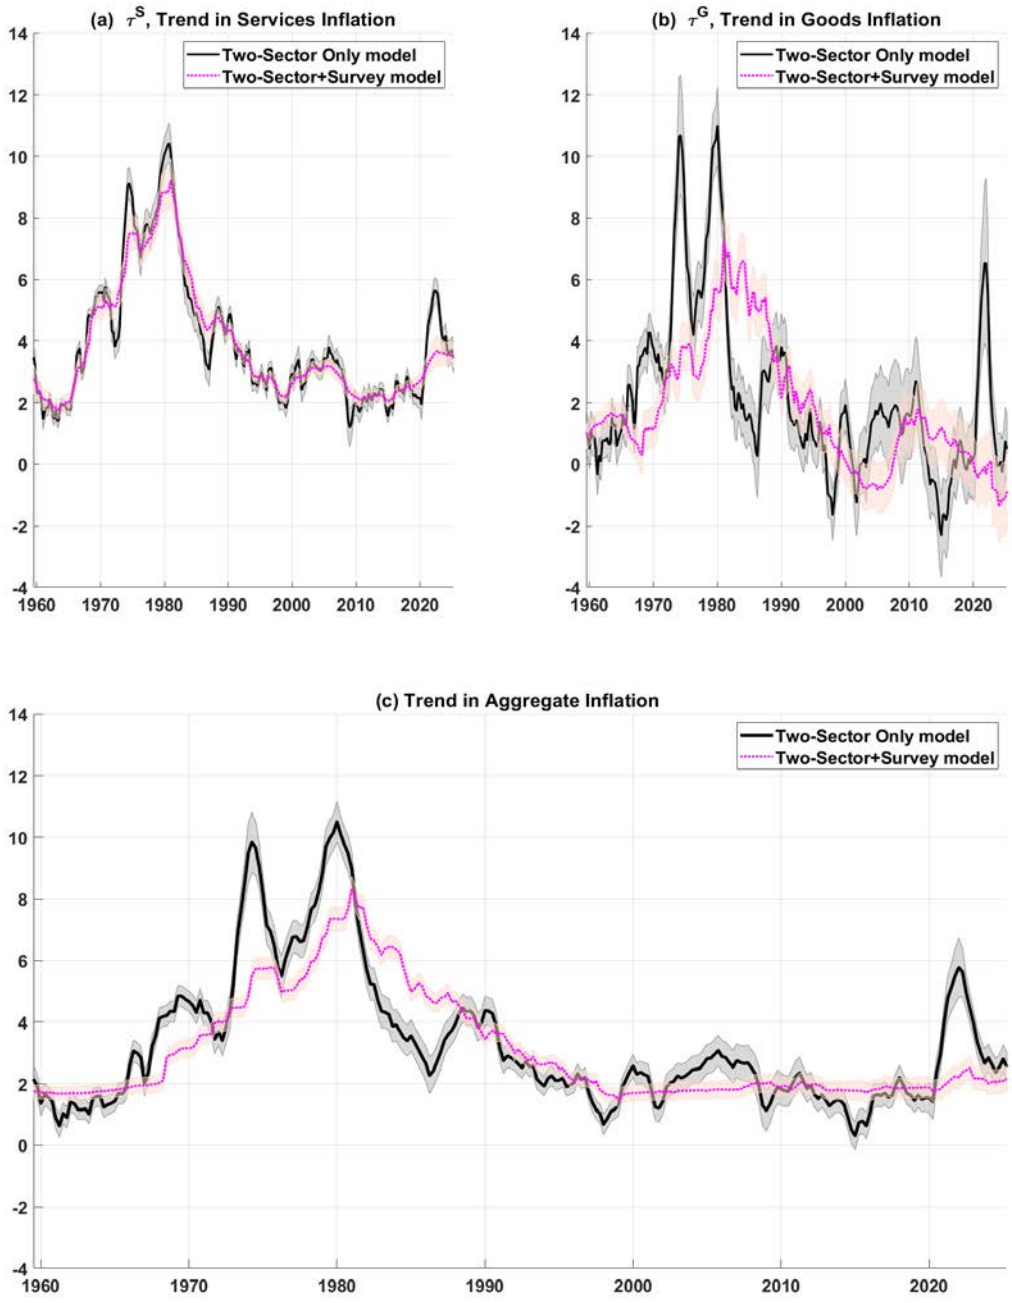
Note: Plotted are posterior mean estimates based on the full sample.

Figure 2: Trend Inflation, Baseline Model



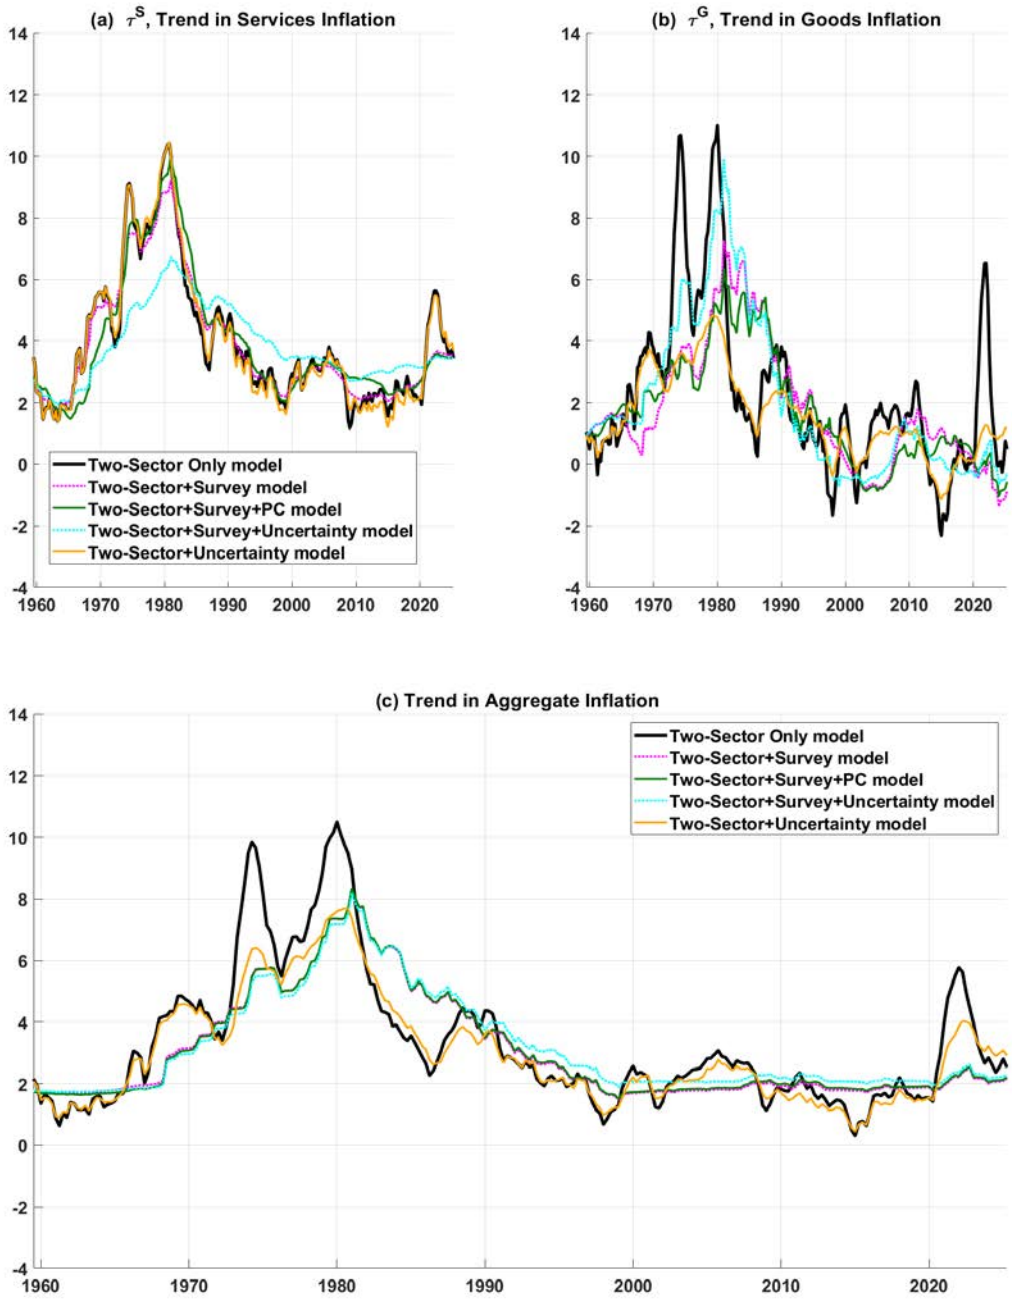
Note: The posterior estimates are based on the full sample (from 1959Q3 through 2025Q2). The shaded area represents the 68% posterior coverage interval. Panel (a), posterior estimates of the trend in services PCE inflation. Panel (b), posterior estimates of the trend in goods PCE inflation. Panel (c), posterior estimates of the trend in aggregate PCE inflation, constructed as a weighted average of the trends in services and goods PCE inflation.

Figure 3: Marginal Contribution of Survey Expectations Data on Trend Inflation



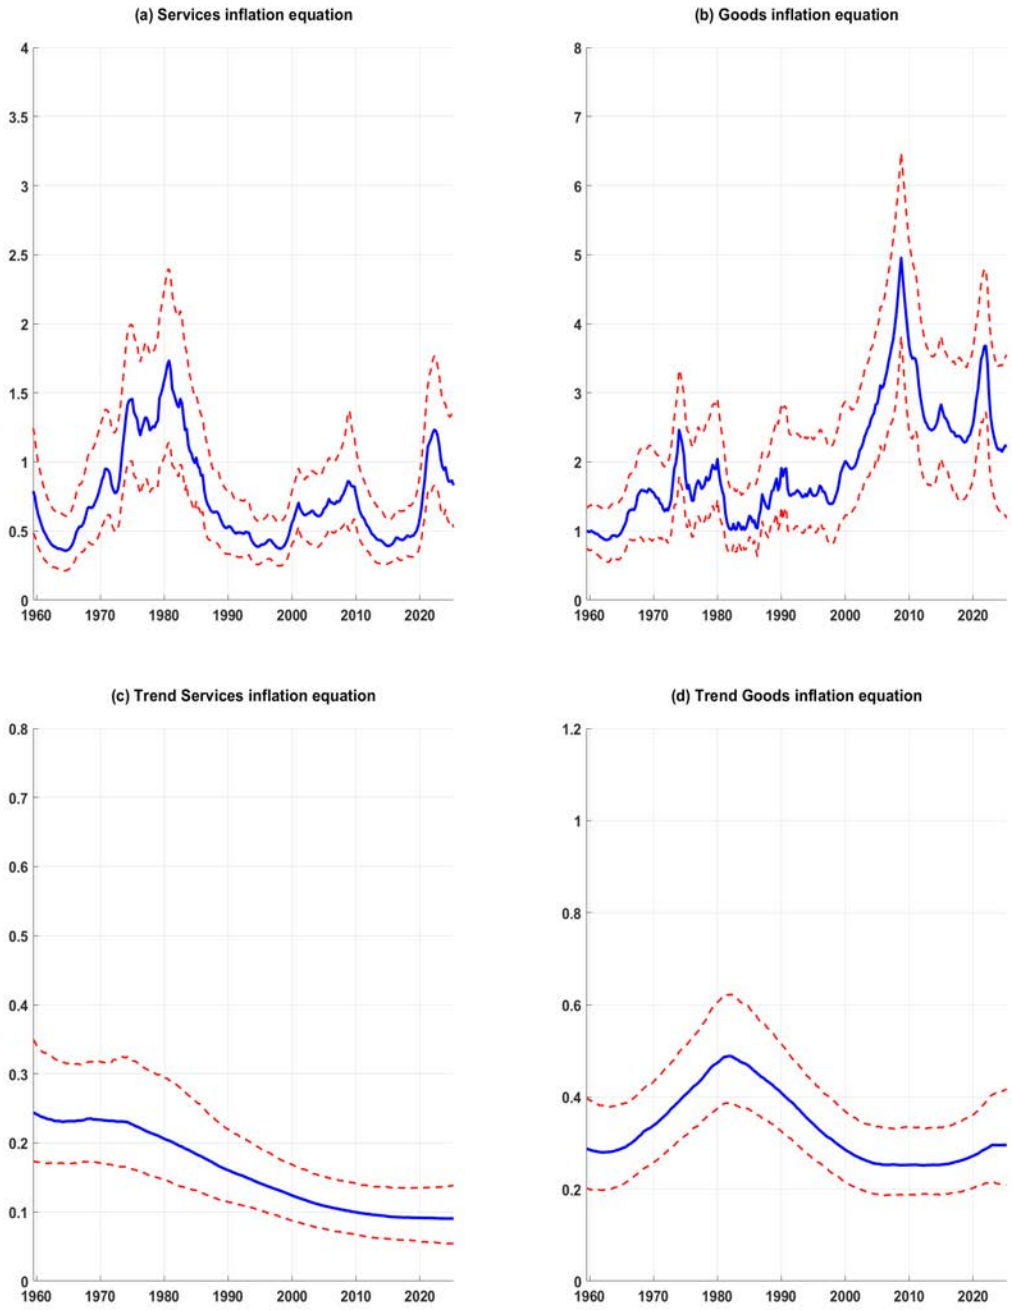
Note: The posterior estimates are based on the full sample (from 1959Q3 through 2025Q2). The shaded area represents the 68% posterior coverage interval. Panel (a), posterior estimates of the trend in services PCE inflation. Panel (b), posterior estimates of the trend in goods PCE inflation. Panel (c), posterior estimates of the trend in aggregate PCE inflation, constructed as a weighted average of the trends in services and goods PCE inflation.

Figure 4: Marginal Contributions of Uncertainty and the Phillips Curve



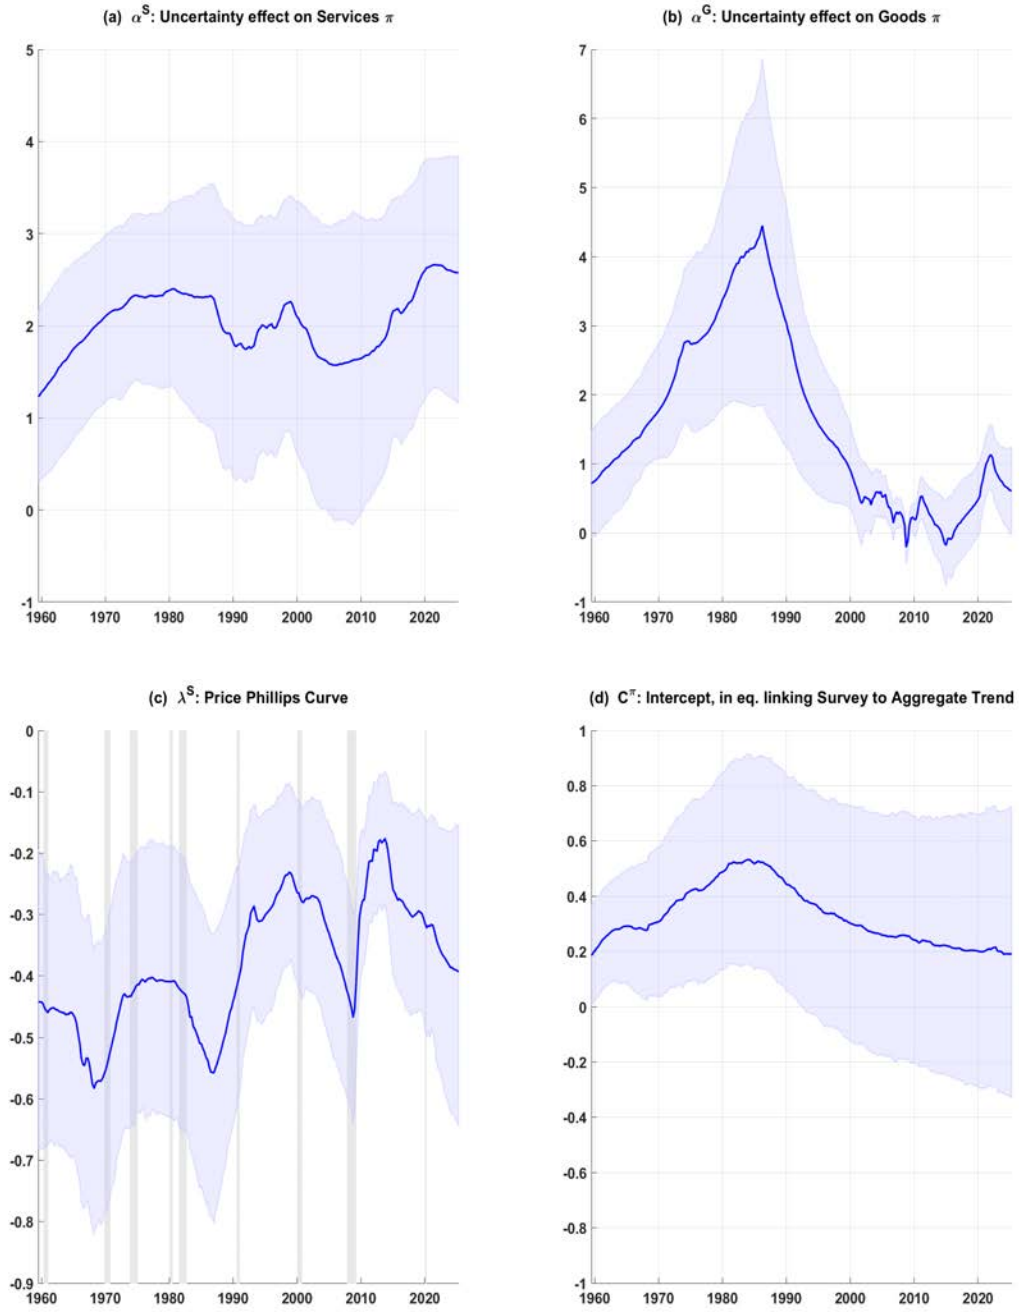
Note: The posterior estimates are based on the full sample (from 1959Q3 through 2025Q2). Panel (a), posterior estimates of the trend in services PCE inflation. Panel (b), posterior estimates of the trend in goods PCE inflation. Panel (c), posterior estimates of the trend in aggregate PCE inflation, constructed as a weighted average of the trends in services and goods PCE inflation.

Figure 5: Estimates of Stochastic Volatility



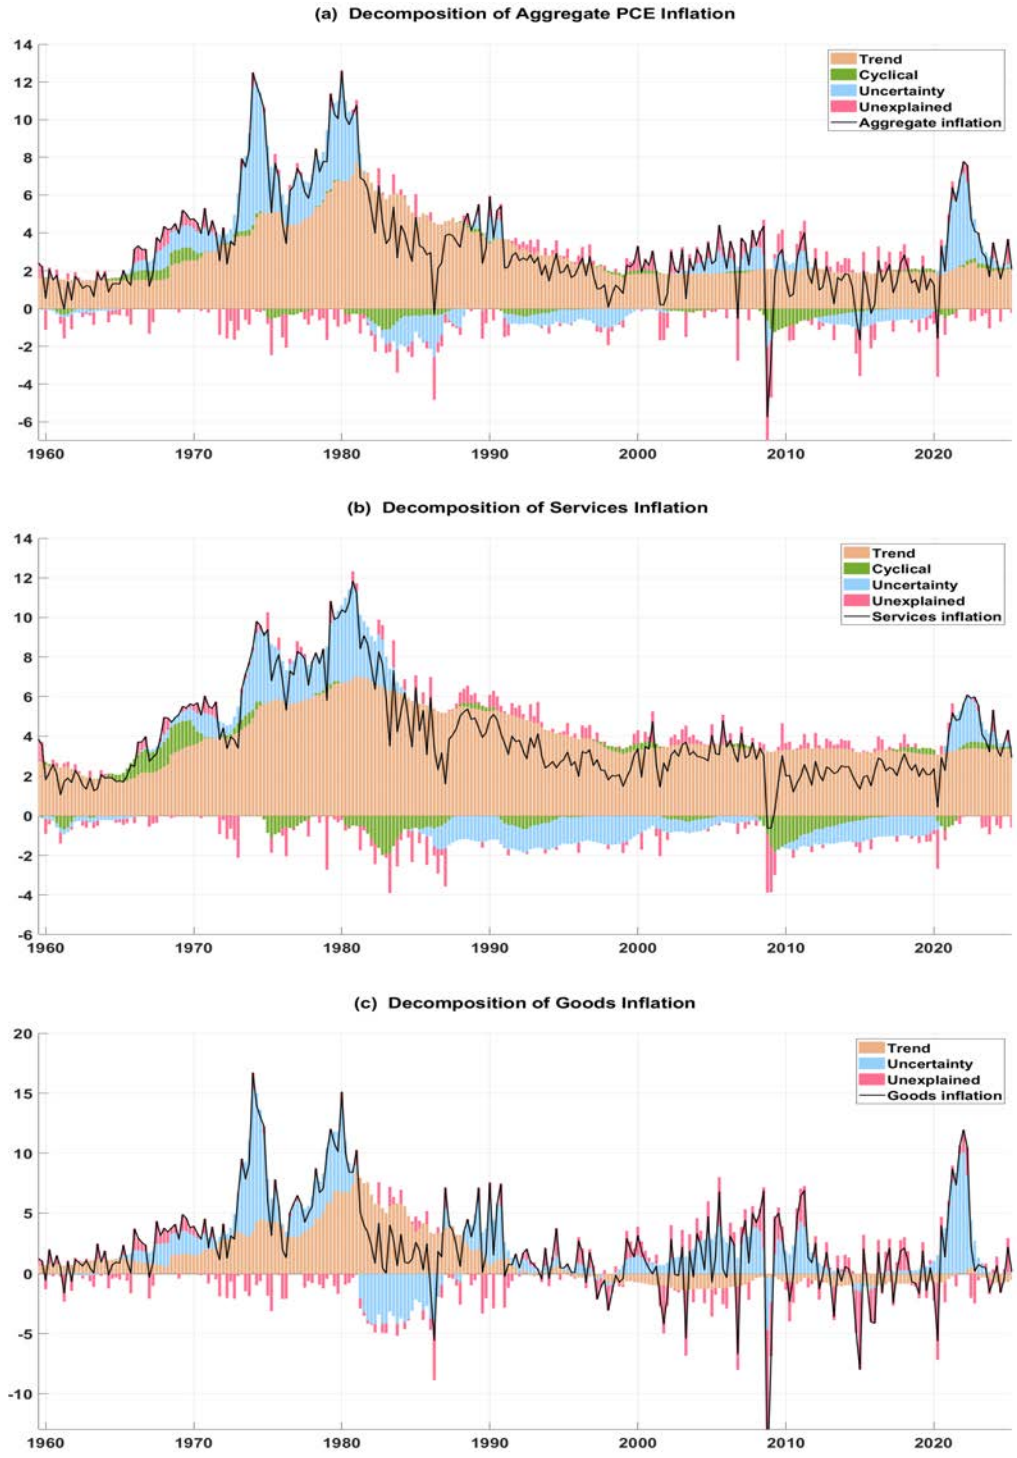
Note: Panel (a) and Panel (b) present the posterior estimates of the standard deviation (sd) of the shocks defining the services inflation and goods inflation measurement equations. Panel (c) and Panel (d) present the estimates of the sd of the shocks to the processes governing the evolution of the trend in services inflation and the trend in goods inflation. The solid lines represent the posterior median and the dashed lines represent the 90% credible intervals. The estimates are computed using the full sample (from 1959Q3 through 2025Q2).

Figure 6: Estimates of Time-varying Parameters



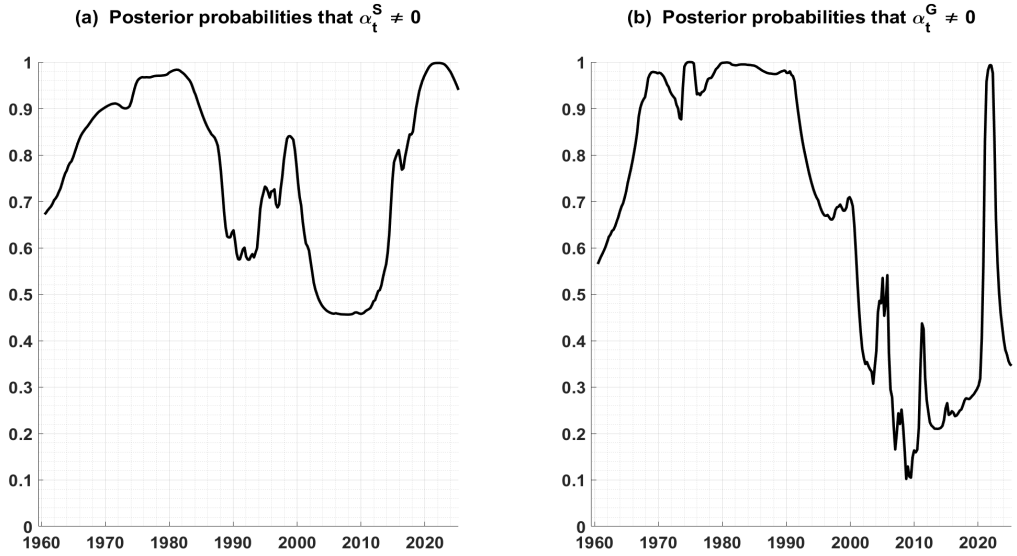
Note: Plotted are the posterior estimates of the time-varying parameters computed using the full sample (from 1959Q3 through 2025Q2). The solid lines represent the posterior mean, and the shaded area represents the 68% credible intervals.

Figure 7: Historical Data Decomposition of PCE Inflation



Note: The decomposition is based on the posterior mean estimates computed using the full sample (from 1959Q3 through 2025Q2).

Figure 8: Dynamic Posterior Probabilities



Note: The posterior estimates are based on the full sample (from 1959Q3 through 2025Q2). Panel (a), formal assessment of the time-varying uncertainty effect on the level of services inflation. Panel (b), formal assessment of the time-varying uncertainty effect on the level of goods inflation.

Table 1: Parameter Estimates Based on Estimation Sample, 1959Q3-2025Q2

Parameter	Parameter description	Posterior estimates M1: Baseline		
		Mean	**5%**	**95%**
σ^2_{hS}	Var. of the Volatility – Services π measurement eq.	0.243^2	0.191^2	0.299^2
σ^2_{hG}	Var. of the Volatility – Goods π measurement eq.	0.227^2	0.171^2	0.289^2
$\sigma^2_{h\tau S}$	Var. of the Volatility – Trend Services π eq.	0.058^2	0.042^2	0.076^2
$\sigma^2_{h\tau G}$	Var. of the Volatility – Trend Goods π eq.	0.064^2	0.045^2	0.087^2
σ^2_{λ}	Var. of the shocks to TVP λ (Phillips curve slope)	0.056^2	0.039^2	0.078^2
$\sigma^2_{\alpha S}$	Var. of the shocks to TVP α^S	0.194^2	0.084^2	0.188^2
$\sigma^2_{\alpha G}$	Var. of the shocks to TVP α^G	0.256^2	0.089^2	0.456^2
$\sigma^2_{c\pi}$	Var. of the shocks to TVP C^{π}	0.060^2	0.052^2	0.068^2
$\sigma^2_{z\pi}$	Var. of the shocks to the eq. linking survey to π^*	0.141^2	0.130^2	0.152^2
β^{π}	Link between π^* and survey	0.925	0.845	1.005

Table 2: Bayesian Model Comparison: Baseline Model vs. Variants

Panel A: MDD of Services and Goods Inflation

	1960-2025Q2, MDD		1960-2019, MDD	
Model	Services π	Goods π	Services π	Goods π
M3: Two-Sector only	-285.0	-582.7	-256.1	-527.9
M4: Two-Sector + Survey	-325.7	-677.4	-286.4	-613.4
M5: Two-Sector + Survey + PC	-319.5	-661.9	-280.6	-598.8
M6: Two-Sector + Survey + Uncertainty	-327.6	-623.3	-292.2	-568.9
Baseline M1: Two-Sector + Survey + PC + Uncertainty	-310.4	-592.6	-276.1	-536.1
M7: Two-Sector + Uncertainty	-282.7	-584.9	-254.4	-529.0

Panel B: MDD of Aggregate Inflation

	MDD Aggregate PCE π	
Model	1960-2025Q2 sample	1960-2019 sample
M2: Univariate	-398.8	-361.2
M3: Two-Sector only	-359.4	-325.1
M4: Two-Sector + Survey	-412.1	-367.6
M5: Two-Sector + Survey + PC	-403.6	-359.5
M6: Two-Sector + Survey + Uncertainty	-399.1	-360.0
Baseline M1: Two-Sector + Survey + PC + Uncertainty	-398.0	-358.4
M7: Two-Sector + Uncertainty	-357.7	-323.8

Note: MDD refers to the log of the marginal data density, which measures the model's fit to the services inflation data and to the goods inflation data. The top panel reports the model's fit to the data over the entire sample (1960-2025Q2) and the pre-COVID sample (1960-2019Q4); to reduce sensitivity to the prior, in computing the MDD, we discard the initial two predictive likelihoods corresponding to 1959Q3 and 1959Q4. The bottom panel reports the model fit to the aggregate PCE inflation for the Univariate model (which models directly the aggregate PCE inflation) and variants of the Baseline model.

Table 3: Out-of-Sample Point Forecasting Performance

	1999-2025Q2, sample				1999-2019Q4, sample			
Relative RMSFE, relative to Univariate UCSV	h=2Q	4Q	8Q	12Q	2Q	4Q	8Q	12Q
M3: Two-Sector only	1.05	0.96	0.82	0.77	1.03	1.03	1.02	0.99
M4: Two-Sector + Survey	1.08	1.00	0.85	0.81	0.99	0.97	0.96	0.95
M5: Two-Sector + Survey + PC	1.04	0.98	0.85	0.77	0.95	0.95	0.91	0.85
M6: Two-Sector + Survey + Uncertainty	0.82**	0.86	0.94	1.02	0.88	0.99	1.11	1.17
Baseline M1: Two-Sector + Survey + PC + Uncertainty	0.74**	0.82*	0.94	1.03	0.79*	0.90***	0.97	1.07
M7: Two-Sector + Uncertainty	0.78**	0.80**	0.82	0.87	0.82	0.91	0.98	1.06

Note: The forecasts and associated accuracy correspond to the average of the quarterly annualized rate of aggregate PCE inflation over h horizons. All the ratios reported are relative root mean squared errors; relative to forecast accuracy of the univariate unobserved-component stochastic volatility model, similar to Stock and Watson (2007). The table reports statistical significance based on the Diebold-Mariano and West test with the lag $h-1$ truncation parameter of the HAC variance estimator and adjusts the test statistic for the finite sample correction proposed by Harvey et al. (1997); *up to 10% significance level. The test statistics use two-sided standard normal critical values for horizons less than or equal to 8 quarters, and two-sided t-statistics for horizons greater than 8 quarters.

Table 4: Out-of-Sample Density Forecasting Performance

	1999-2025Q2, sample				1999-2019Q4, sample			
Mean Relative Log Score	h=2Q	4Q	8Q	12Q	2Q	4Q	8Q	12Q
M3: Two-Sector only	0.01	0.02	0.06	0.10*	0.01	0.03	0.06	0.12*
M4: Two-Sector + Survey	0.01	0.03	0.08**	0.10**	0.01	0.04	0.08*	0.12**
M5: Two-Sector + Survey + PC	0.01	0.03	0.08**	0.11***	0.01	0.04	0.09**	0.13***
M6: Two-Sector + Survey + Uncertainty	-0.01	-0.03*	-0.03**	-0.04***	-0.02	-0.03*	-0.05***	-0.06***
Baseline M1: Two-Sector + Survey + PC + Uncertainty	0.02	0.00	-0.01	-0.02**	0.01	0.00	-0.02	-0.04***
M7: Two-Sector + Uncertainty	0.04***	0.03	0.02	0.01	0.04***	0.04*	0.02	0.00

Note: The forecasts and associated accuracy correspond to the average of the quarterly annualized rate of aggregate PCE inflation over h horizons. All the numbers reported are the average relative log scores, i.e., the mean difference between the log predictive score of the model of interest and that of the benchmark model; relative to forecast accuracy of the univariate unobserved-component stochastic volatility model, similar to Stock and Watson (2007). The table reports statistical significance based on the likelihood-ratio test of Amisano and Giacomini (2007), with a two-sided t-test.

References

Acharya, Viral V., Sebastian Hillenbrand, Venky Venkateswaran, and Margaret Underwood (2026). "Inflation uncertainty: Measurement, causes, and consequences." *Working Paper Manuscript*.

Amisano, Gianni and John Geweke (2010). "Comparing and evaluating bayesian predictive distributions of asset returns." *International Journal of Forecasting*, 26(2), pp. 216–230. doi:https://doi.org/10.1016/j.ijforecast.2009.10.007.

Amisano, Gianni and Raffaella Giacomini (2007). "Comparing density forecasts via weighted likelihood ratio tests." *Journal of Business & Economic Statistics*, 25(2), pp. 177–190. doi:10.1198/073500106000000332.

Ball, Laurence (1992). "Why does high inflation raise inflation uncertainty?" *Journal of Monetary Economics*, 29, pp. 371–388. doi:https://doi.org/10.1016/0304-3932(92)90032-W.

Ball, Laurence, Carlos Carvalho, Christopher Evans, and Luca Antonio Ricci (2024). "Weighted median inflation around the world: A measure of core inflation." *Journal of International Money and Finance*, 142, p. 103,017. doi:10.1016/j.jimonfin.2024.103017.

Beveridge, Stephen and Charles 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(2), pp. 151–174. doi:10.1016/0304-3932(81)90040-4.

Born, Benjamin and Johannes Pfeifer (2014). "Policy risk and the business cycle." *Journal of Monetary Economics*, 68(1), pp. 68–85. doi:https://doi.org/10.1016/j.jmoneco.2014.07.012.

Canova, Fabio and Fernando Perez Forero (2024). "Does the transmission of monetary policy shocks change when inflation is high?" *CEPR Discussion Papers 18993*. URL `https://cepr.org/publications/dp18993`.

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino (2018). "Measuring uncertainty and its impact on the economy." *Review of Economics and Statistics*, 100(5), pp. 799—-815. doi:https://doi.org/10.1162/rest_a_00693.

Carriero, Andrea, Todd E. Clark, Massimiliano Marcellino, and Elmar Mertens (2024). "Addressing COVID-19 Outliers in BVARs with Stochastic Volatility." *The Review of Economics and Statistics*, 106(5), pp. 1403—-1417. doi:10.1162/rest_a_01213.

Chan, Joshua C. C. (2017). "The stochastic volatility in mean model with time-varying parameters: An application to inflation modeling." *Journal of Business & Economic Statistics*, 35(1), pp. 17–28. doi:10.1080/07350015.2015.1052459.

Chan, Joshua C. C., Todd E. Clark, and Gary Koop (2018). "A new model of inflation, trend inflation, and long-run inflation expectations." *Journal of Money, Credit and Banking*, 50(1), pp. 5–53. doi:10.1111/jmcb.12452.

Chan, Joshua C. C. and Ivan Jeliazkov (2009). "Efficient simulation and integrated likelihood estimation in state space models." *International Journal of Mathematical Modelling and Numerical Optimisation*, 1(1/2), p. 101. URL https://www.inderscienceonline.com/doi/abs/10.1504/IJMMNO.2009.03009.

Chan, Joshua C. C., Gary Koop, and Simon M. Potter (2013). "A new model of trend inflation." *Journal of Business & Economic Statistics*, 31(1), pp. 94–106. doi:10.1080/07350015.2012.741549.

Chan, Joshua C. C., Gary Koop, and Simon M. Potter (2016). "A bounded model of time variation in trend inflation, Nairu and the Phillips curve." *Journal of Applied Econometrics*, 31(3), pp. 551–565. doi:10.1002/jae.2442.

Chan, Joshua C. C. and Rodney Strachan (2012). "Estimation in non-linear non-Gaussian statespace models with precision-based methods." *CAMA Working Paper series 2012-13*.

Clark, Todd E. and Taeyoung Doh (2014). "Evaluating alternative models of trend inflation." *International Journal of Forecasting*, 30(3), pp. 426–448. doi:10.1016/j.ijforecast.2013.11.005.

Clark, Todd E., Matthew V. Gordon, and Saeed Zaman (2025). "Forecasting core inflation and its goods, housing, and supercore components." *International Journal of Central Banking*, 21(4), p. 351–403. URL https://www.ijcb.org/sites/default/files/journal/v21n4/ijcb-v21n4-forecasting-core-inflation-and-its-goods-housing-and-supercore-components.pdf.

Coibion, Olivier and Yuriy Gorodnichenko (2015). "Information rigidity and the expectations formation process: A simple framework and new facts." *American Economic Review*, 105, p. 2644–2678. doi:DOI: 10.1257/aer.20110306.

Cross, Jamie L., Chenghan Hou, Gary Koop, and Aubrey Poon (2023). "Large stochastic volatility in mean VARs." *Journal of Econometrics*, 236(1), p. 105,469. doi:https://doi.org/10.1016/j.jeconom.2023.05.006.

Eisenstat, Eric and Rodney W. Strachan (2016). "Modelling inflation volatility." *Journal of Monetary Economics*, 31(5), pp. 805–820. doi:https://doi.org/10.1002/jae.2469.

Eo, Yunjong, Luis Uzeda, and Benjamin Wong (2023). "Understanding trend inflation through the lens of the goods and services sectors." *Journal of Applied Econometrics*, 38(5), pp. 751–766. doi:10.1002/jae.2975.

Garcia, Juan Angel and Aubrey Poon (2022). "Inflation trends in Asia: implications for central banks." *Oxford Economic Papers*, 74(3), pp. 671—-700. doi:https://doi.org/10.1093/oep/gpab050.

Garciga, Christian, Randal J. Verbrugge, and Saeed Zaman (2025). "Improving the median CPI: Maximal disaggregation isn't necessarily optimal." *Federal Reserve Bank of Cleveland, Working Paper No 24-02R*. doi:10.26509/frbc-wp-202402r.

Harvey, David, Stephen Leybourne, and Paul Newbold (1997). "Testing the equality of prediction mean squared errors." *International Journal of Forecasting*, 13(2), pp. 281–291. doi:10.1016/S0169-2070(96)00719-4.

Koop, Gary (2003). *Bayesian Econometrics*. Wiley-Interscience.

Koop, Gary, Roberto Leon-Gonzalez, and Rodney Strachan (2010). "Dynamic probabilities of restrictions in state space models: An application to the Phillips curve." *Journal of Business and Economic Statistics*, 28(3), pp. 370–379. doi:10.1198/jbes.2009.07335.

Mertens, Elmar (2016). "Measuring the Level and Uncertainty of Trend Inflation." *The Review of Economics and Statistics*, 98(5), pp. 950–967. doi:10.1162/REST_a_00549.

Stella, Andrea and James H. Stock (2015). "A state-dependent model for inflation forecasting." In Siem Jan Koopman and Neil Shephard, editors, *Unobserved Components and Time Series Econometrics*, pp. 14–29. Oxford University Press. doi:10.1093/acprof:oso/9780199683666.003.0003.

Stock, James H. and Mark W. Watson (2007). "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking*, 39, pp. 3–33. doi:10.1111/j.1538-4616.2007.00014.x.

Stock, James H. and Mark W. Watson (2016). "Core Inflation and Trend Inflation." *Review of Economics and Statistics*, 98(4), pp. 770–784. doi:10.1162/REST_a_00608.

Tallman, Ellis W. and Saeed Zaman (2017). "Forecasting inflation: Phillips curve effects on services price measures." *International Journal of Forecasting*, 33(2), pp. 442–457. doi:10.1016/j.ijforecast.2016.10.004.

Verdinelli, Isabella and Larry Wasserman (1995). "Computing Bayes factors using a generalization of the Savage-Dickey density ratio." *Journal of the American Statistical Association*, 90(430), pp. 614–618. doi:https://doi.org/10.2307/2291073.

Watson, Mark W. (2026). "Forecasting the covid surge in inflation." *National Bureau OF Economic Research*, Working Paper 35435.

Zaman, Saeed (2024). "A unified framework to estimate macroeconomic stars." *The Review of Economics and Statistics*, pp. 1–45. doi:10.1162/rest_a_01570.

Zaman, Saeed (2025). "Disciplining trend inflation: Bounded random walk vs. survey expectations." *Manuscript*.