

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. 26-08

March 2026

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

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‡

February 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 state-of-the-art Bayesian methods. We document the competitive properties of the new model compared to variants that include only a subset of the above features. 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.

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

JEL Codes: C11, C32, E31

*We thank Todd Clark, James Mitchell, and participants at the Advances in Econometrics Bayesian Macroeconometrics Workshop 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. The addition of 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 for the stochastic volatility process in the context of the Stock and Watson model and found that all the formulations

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)'s 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

of SV yield very 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 trend in aggregate inflation 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 correlation between the innovations in the equations describing trends in the respective services and goods inflation. By doing so, they were able to explain the source of changes in the volatility of the trend in aggregate inflation. 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 the changing influence of the associated sector in determining trend inflation. One of the insightful findings of Stock and Watson (2016) was that food components have gotten a non-trivial weight in the determination of trend inflation since the 1990s, suggesting approaches that rely on inflation excluding food and energy (i.e., core inflation) for trend estimation may not be optimal. Compared to the univariate Stock and Watson model, Stock and Watson (2016) document competitive but not better forecast performance. Similarly, Eo et al. (2023) document forecast accuracy comparable to that of both the univariate Stock and Watson and the multivariate Stock and Watson (2016) models.

Although the models discussed above yield comparable average forecast accuracy, the resulting estimates of trend inflation differ substantially at times, in turn implying contrast-

ing 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 cyclical variations in goods and 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 data 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, the inclusion of survey expectations data has the most impact, whereas allowing for the Phillips curve relationship has the least. The modeling of the aggregate inflation trend using sectoral data has implications for trend inflation estimates, but mainly after 1990, when the dynamics of goods inflation began to diverge from those of services inflation; see, for example, Clark (2004). The role of uncertainty in explaining inflation fluctuations is found to be very important, and the evidence strongly suggests its time-varying effects on the level of 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.

We summarize our findings as follows. First, the results indicate economically and statistically significant evidence of time variation in the effects of uncertainty on the levels of goods and services inflation. 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 the survey data, the trend component explains a much greater share of the variation in inflation, reducing the share of variation 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, modeling the aggregate inflation trend with sectoral data yields inflation trend estimates that are generally lower than those from a model based solely on aggregate inflation data (i.e., the single-sector model). Since 1990, the differences in trend estimates have ranged from 0.1 to 0.6 percentage points (ppt), arguably economically meaningful differences for a central bank targeting 2 percent inflation.

Third, 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.[4] However, in contrast to previous research, including survey data worsens the model's fit to the inflation data, as measured by the one-step-ahead predictive likelihood. Furthermore, the dynamic posterior probabilities provide limited support for the inclusion of survey expectations over most of the sample period considered.

That said, pseudo-out-of-sample forecasting exercises evaluating forecast performance at horizons of 1 to 3 years suggest comparable performance for the model with and without survey data. Moreover, survey expectations data played a crucial role in reliably inferring the inflation trend during the post-COVID inflation surge. Excluding survey data from

[4]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.

the model yielded markedly different inferences, attributing most of the pandemic surge to trend inflation, similar to the model of Stock and Watson (2007), which, with the benefit of hindsight, seems unreasonable. From mid-2021 through 2023, inflation had moderated from a high of 8% to a rate closer to 3%, indicating that most of the surge was transitory.

We illustrate various factors, including differences in the sample period, that help explain the contrasting findings regarding the usefulness of survey expectations for trend inflation estimation relative to prior research.

Fourth, 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 cyclical effects in explaining inflation fluctuations.

Lastly, the predictive accuracy of the baseline model that combines all features yields point and density forecasts that generally rival those of the univariate Stock and Watson benchmark model. 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 has a supplementary online appendix that discusses Bayesian estimation steps, prior and posterior comparison, and additional results.

2 Empirical Framework

2.1 Two-Sector Empirical Model: Baseline

Our empirical model substantially builds on the two-sector UCSV model of Eo et al. (2023). Eo et al. (2023) 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 between the innovations in the equations describing trends in the respective services and goods inflation. We extend their bivariate UCSV model 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 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.[5]

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

Services inflation is modeled as

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

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, which are assumed stationary components, converge to zero, services inflation gravitates to τ^S.

Goods inflation is modeled as (measurement equation)

[5]To keep the estimation tractable, we do not allow for the "direct" link between the two inflation measures via the second moments (error-covariance structure), but they are linked "indirectly" because of the joint modeling. Furthermore, similar to Eo et al. (2023), we do allow the "direct" first-moment link between the two inflation measures via the trend inflation accounting equation (see equation 7).

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

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 the fact that trend inflation in the model 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,

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

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 of various empirical features incorporated into the Baseline model. These additional specifications are listed below.

M2: Baseline but no TVP in α^S. To assess the empirical support of time variation in the parameter α^S, which captures the effect of uncertainty on the level of services inflation, we estimate a variant of the baseline model with a time-invariant parameter α^S.

M3: Baseline but no TVP in α^S and α^G . To assess the empirical support for time variation in the parameter α^G, which captures the effect of uncertainty on the level of goods inflation, we take the model M2 and shut down time variation in parameter α^G. Comparing

the model fit of M2 with M3 provides an assessment of the usefulness of time variation in α^G, and comparing the model fit of M3 with the baseline M1 provides a joint assessment of the usefulness of allowing for time variation in both parameters α^S and α^G.

M4: Baseline but no TVP in λ. To assess the empirical support for time variation in the parameter λ, which captures the business cycle effects on the level of services inflation (Phillips curve slope), we estimate a variant of the baseline model with a time-invariant parameter λ.

M5: Baseline but no TVP in α^S, α^G, and λ . To jointly assess the empirical support for time variation in the parameters α^S, α^G, and λ, we estimate a variant of the baseline model that shuts off time variation for all three parameters.

M6: Baseline but No α^S, α^G, and λ. To assess the usefulness of allowing for uncertainty effects on the levels of services and goods inflation, and business cycle effects on the level of services inflation, we estimate a variant of the baseline model that zeroes out α^S, α^G, and λ for all time periods.

M7: Baseline but No Survey data. To assess the usefulness of survey expectations, we estimate a variant of the baseline model that excludes the equation linking survey data to trend inflation.

M8: Single-sector. This variant models inflation dynamics using aggregate PCE inflation rather than the two-sector goods-and-services inflation breakdown used in the baseline model. It keeps all other features of the baseline model: uses survey data, allows time-varying uncertainty effects on the level of aggregate inflation, allows a time-varying Phillips curve relationship, and includes SV in the shocks to trend inflation and in the measurement equation defining inflation. (See Section 2.3 for model details.) Comparing the model fit to aggregate inflation data for this model (M8) with that of the baseline model (M1) assesses the usefulness of modeling aggregate PCE inflation using sectoral inflation data (goods and services). It is worth pointing out that, even though this model does not use sector-level data, the fact that it incorporates three other features makes it a novel empirical model.

M9: Single-sector but No α^S, α^G, and λ. To assess the usefulness of allowing for uncertainty and business cycle effects on the level of aggregate PCE inflation, we estimate a variant of the single-sector model that zeroes out α^S, α^G, and λ for all time periods.

M10: Single-sector but No Survey data. To assess the usefulness of survey expectations, we estimate a variant of the single-sector model that excludes the equation linking survey data to trend inflation.

2.3 Single-Sector Empirical Model

The specification of our single-sector model closely follows that of the two-sector model, except that, instead of using sector-level data, it directly models the dynamics of aggregate inflation.

Aggregate PCE inflation is modeled as

$$\pi_t = \underbrace{\tau_t}_{\text{Trend}} + \underbrace{\lambda_t \ (U_t - U_t^*)}_{\text{Cyclical effect}} + \underbrace{\alpha_t \ (e^{h_t} - e^{\overline{h}_t})}_{\text{Uncertainty effect}} + \varepsilon_t^{\pi}, \ \ \varepsilon_t^{\pi} \sim N(0, e^{h_t}) \tag{11}$$

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

Like before, the time-varying parameter, λ, the slope of the price Phillips curve is constrained between (-1,0), and the parameter α, which captures the volatility feedback, evolves according to an RW process.

Trend in aggregate PCE inflation, τ_t is modeled as a driftless random walk, with a shock term, whose variance is allowed to vary over time,

$$\tau_t = \tau_{t-1} \ + \ \varepsilon_t^{\tau}, \ \ \varepsilon_t^{\tau} \sim N(0, e^{h_t^{\tau}}) \tag{12}$$

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

$$Z_t^{\pi} = C_t^{\pi} \;+\; \beta^{\pi}\tau_t \;+\; \varepsilon_t^{z\pi},\;\; \varepsilon_t^{z\pi} \sim N(0, \sigma_{z\pi}^2) \tag{13}$$

Like before, C_t^{π} is a time-varying intercept assumed to evolve as an RW process.

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

$$\begin{aligned} h_t &= h_{t-1} \;+\; \varepsilon_t^{h},\;\; \varepsilon_t^{h} \sim N(0, \sigma_h^2) \\ h_t^{\tau} &= h_{t-1}^{\tau} \;+\; \varepsilon_t^{h\tau},\;\; \varepsilon_t^{h\tau} \sim N(0, \sigma_{h\tau}^2) \end{aligned} \tag{14}$$

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 spans the period 1959Q2 to 2025Q2 and are downloaded from Haver Analytics, except for the inflation expectations data, which are downloaded from the Fed Board's website (PTR) and the Philadelphia Fed's website (SPF).

As evident in the model description, 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 renders estimation using Kalman Filter-based methods unfeasible. But there are other methods for solving such nonlinear state-space models, and 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). And 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. Discard the first 500,000 draws, and of the remaining, keep every 100th draw. The reported results are based on 5000 retained draws.

3.2 Marginal Likelihood and the Posterior Odds Ratio

Bayesian model comparison based on the marginal likelihood criterion is used to compare the fit of the baseline model and its variants, providing an assessment of 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 $y_{1:t}^j = (y_1^j, ...y_t^j)'$ 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,

$$p(y^j|X_i^j, M_i) = \prod_{t=3}^{T} p(y_t^j|y_{1:t-1}^j, X_{1:t,i}^j, M_i) \tag{15}$$

where $j = \{\text{Services inflation, goods inflation}\}$, $p(y_t^j|y_{1:t-1}^j, X_{1:t,i}^j, M_i)$ is the predictive likelihood 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 the 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 provide the proof of equality between the BF and 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 two main variants: Baseline-NoSurvey (model M6) and Single-sector (model M8), and and compare the forecasting performance of each of these three models to the univariate Stock and Watson benchmark model, denoted as UCSV in the reported results.

We assess both point and density forecasts for each model. Point forecasts, which are the

posterior mean of the density forecasts, are evaluated using the root mean squared forecast error (RMSE) 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. The density forecasts are evaluated using the logarithmic predictive score, and the statistical significance of gains is assessed 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 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.

4 Empirical results

4.1 Trend Inflation

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 from 2012 to 2019, when inflation remained stubbornly below 2 percent, trend inflation is estimated to be stable at a touch 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.[6] 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% to 2.4%). 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 4, 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).

4.1.1 Influence of survey expectations data on trend

The parameter estimates suggest a strong influence of the survey data on the 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 5, 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}.

The strong influence of survey data is evident in Figure 2, which builds on Figure 1 by plotting the trend inflation from the model excluding survey data. Comparing estimates from the baseline model and the model without survey data, the level of 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) that are very volatile.

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

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 2.5 percentage points, from 1.5 percent to a level closer to 4.0 percent, primarily driven by the trend in services inflation, which rises from 2.0 percent to 5.5 percent.

4.1.2 Sensitivity of trend to other modeling assumptions

The figures in Appendix A6 illustrate the sensitivity of trend inflation to other modeling assumptions, specifically the assumption of time variation in the parameters that capture uncertainty and cyclical effects. Shutting down time variation in the parameters capturing the uncertainty effects has a non-trivial impact on trend estimates for goods and services, especially since 1990 (Figure 10 in A6). Prior to 1990, the model with time-invariant uncertainty parameters estimated a higher trend in services and a lower trend in goods inflation than the baseline model. These offsetting changes in trends for services and goods yield estimates of aggregate trend inflation similar to those from the baseline model. Post-1990, trends in both goods and services are estimated to be lower than the baseline model's estimates, resulting in an aggregate inflation trend that is 0.3 to 0.4 ppts lower than the baseline model's estimate.

Shutting down time variation in the Phillips curve relationship has a small impact on estimated aggregate trend inflation, reflecting the net effect of the mostly marginal impact on trend goods inflation and a non-trivial impact on trend services inflation (Figure 11 in A6). The aggregate trend is estimated to be 0.2 ppts lower than the baseline model's estimate since 1995, and prior to 1995, the estimates are similar to the baseline model's. This similarity in the estimates earlier in the sample reflects offsetting differences in the trend estimates of goods and services inflation; for example, between the late 1960s and the early 1970s, the estimate of the services trend from the baseline was 0.3 to 0.5 ppts lower than that inferred from the model variant without TVP PC (Model M4), but the trend in goods inflation during

this period was 0.3 to 0.5 ppts higher, leaving the estimated trend in aggregate inflation similar to that in the baseline. The periods in which the trend services inflation estimates differed notably across the two models are also the periods when the difference between the time-invariant and the TVP PC parameters was sizable (Appendix A9, Figure 16, panel c).

Figure 12 in Appendix A6 compares the trend estimates from the baseline model with those from the model variant that excludes uncertainty effects and the Phillips curve (Model 6). As is evident, shutting down uncertainty and cyclical channels substantially alters estimates of the trend in services and goods inflation, but the implied aggregate trend in inflation remains remarkably similar across the two models, particularly from 1990. This similarity in the aggregate trend is due to offsetting movements in the inflation trends of goods and services. From 1980 to the present, the model that excludes uncertainty yields estimates of the trend in services inflation that are consistently lower than those of the baseline model, and of the trend in goods inflation that are consistently higher. The story is the opposite for the period from 1960 through the end of 1970, where the movements in the trends of goods and services inflation are not fully offsetting, resulting in non-trivial differences in the aggregate trend inflation estimates between the two models.

4.1.3 Comparison of trend to external models

Figure 1 compares the trend estimates from the baseline model with those from several established models in the literature. Panel (a) compares the estimates over the full sample spanning 1960 through 2025Q2, and panel (b) compares the estimates over the recent sample period spanning the onset of the COVID pandemic through the present, i.e., 2020Q1 through 2025Q2. Shown are the posterior mean (or median) smoothed estimates, i.e., estimates based on the full sample, from the models of Stock and Watson (2007), Chan et al. (2016), Tallman and Zaman (2017), and Chan et al. (2018). As shown, there are notable differences across the model estimates. Even if estimates are similar across two or more models within a given period, they can differ substantially across the same models in a different period. For example,

over the recent period from 2020 through 2025, both the baseline model and the model of Chan et al. (2016), which is a bivariate PC model with bounds on the process defining trend inflation, yield similar estimates of trend inflation, but outside of this period, there are significant differences in the inference about the trend inflation. As another example, for the most part, both the univariate Stock and Watson model and the Tallman and Zaman (2017) model, which builds on the Stock and Watson model by bringing in sectoral data and PC relationship, yield similar estimates, but during some periods, e.g., when the labor market is very tight (unemployment gap is very negative) or very loose (unemployment gap is positive), there are non-trivial differences in the trend estimates between the two models.

It is also worth noting that models that include survey expectations data—the baseline model and the Chan, Clark, and Koop model—yield stable estimates of trend inflation over the past 25 years, a period when inflation expectations have been relatively stable at around 2 percent. That said, the difference in the trend estimates across the two models may appear small (with the Chan, Clark, and Koop model inferring trend a few tenths below 2 percent, whereas the baseline model is sometimes at or below or slightly above 2 percent), but from the monetary policy perspective, it is sizable enough to have monetary policy implications.

Overall, the comparison with external models provides further evidence that modeling choices notably influence inferences about the inflation trend. Furthermore, these external models have established competitive (not better) forecast accuracy (in predicting inflation) relative to the univariate Stock and Watson model, making it difficult to select a model based solely on forecast evaluation. Our forecast evaluation examination, as reported in Table 3, demonstrates the generally competitive forecast accuracy of our baseline model and its (selected) variants to the Stock and Watson model, which, on the one hand, is encouraging, but on the other hand, highlights the difficulty of picking a model based on out-of-sample evaluation from a pool of competitive model candidates, something emphasized by Clark and Doh (2014). 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 in inflation and should be preferred.

4.2 Role of Uncertainty Effects

Figure 4 (top row) presents estimates of uncertainty, as measured by stochastic volatility. Panel (a) represents 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.[7]

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, the estimated uncertainty (the estimated stochastic volatility) for both services and goods inflation has moderated, but there remains substantial uncertainty around the uncertainty estimates, as evidenced by very wide credible bands around the posterior median estimates of the stochastic volatility.

The broad contours of the uncertainty estimates shown in Figure 4 are consistent with those reported in the literature on inflation. However, on closer inspection, the comparison of uncertainty estimates across variants of the baseline model reveals that modeling choices affect the uncertainty estimates (i.e., 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

[7] 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. (2019), Cross et al. (2023), and Canova and Forero (2024)).

between the baseline model and the model without the survey data.

Figure 5, panels (a) and (b), report the estimates of the parameters α^S and α^G. These estimates indicate substantial time variation in the effects of uncertainty on the levels of services and 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 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)).

According to the Bayesian model comparison reported in Table 2, the data strongly favor the role of uncertainty (or volatility feedback) for both goods and services inflation, as evidenced by the model fit, with the baseline model's MDD for services inflation -310.4 and goods inflation -592.6 compared to the model variant excluding uncertainty and PC, yielding MDD for services inflation -325.7 and for goods inflation -677.4. Furthermore, the data strongly favor allowing time-varying effects of uncertainty on the levels of goods and services inflation, as the MDD of the model variant without TVP in the parameters α^S and α^G (model M3) is inferior to that of the baseline model (model M1).

The dynamic posterior probabilities shown in Figure 6, 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, it is highly likely that $\alpha^G \neq 0$ from 1960-1990 and from 2020Q4-2023Q1, whereas for the period 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 5. For example, panel (a) of Figure 5 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 the 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 cyclical component.[8]

4.3 Importance of Sectoral Data

A growing body of research has illustrated the usefulness of measuring the trend in aggregate inflation using sector-level inflation data; see, for example, Stock and Watson (2016), Tallman and Zaman (2017), and Eo et al. (2023), among others. 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.

[8]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, as evidenced by the size of the blue bars during these specific periods, hence, it is reasonable to assume the estimated effects are less biased.

Our results 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, the estimate of trend inflation from a single-sector model 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.

Given the above findings, not surprisingly, the model comparison indicates that the baseline model fits the historical data better than the baseline model variant that has all the same features as the baseline model but excludes sectoral data (single-sector model, M8): Baseline model's MDD of aggregate inflation of -398.0 vs. Single-sector model's MDD of -415.8.

An out-of-sample forecast comparison between the baseline and single-sector models suggests comparable point-forecast accuracy, as both models perform competitively with the UCSV benchmark. However, the baseline model produces slightly inferior density forecasts compared to the single-sector model, as shown in Table 3 by comparing the rows labeled "Baseline/UCSV" and "Single-sector/UCSV." These forecasting results are 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). Although our baseline model with sectoral data does not dominate the single-sector model in predictive accuracy, it provides a more interpretable decomposition of historical inflation data.

4.4 Role of Cyclical Effects

Figure 5, panel (c) presents the posterior estimates of parameter λ, which is 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. As is evident, the PC slope parameter is estimated with very little precision. In contrast, a model variant that removes time variation in the PC slope is estimated more precisely, with a posterior mean of -0.3; see Appendix A8, Figure 15, Panel (c). Bayesian model comparison provides support for the time-varying PC slope parameter. As reported in Table 2, the model variant that removes time variation in the PC slope (model M4) has a worse fit to both services and goods inflation data. Furthermore, dynamic posterior probabilities provide strong support for the time-varying PC slope parameter; see Panel (c) in Figure 5.

The model-based decomposition of inflation shown in Figure 7 indicates that the cyclical 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 cyclical effects played a non-trivial role in explaining inflation fluctuations have generally coincided with recessions and the first year or two of the expansions that followed them. In the past couple of years, according to the model, the cyclical effects have exerted a positive influence on inflation, but this influence pales in comparison to the uncertainty factor.

4.5 Usefulness of Survey Expectations

A growing body of research has shown the benefits of incorporating survey-based inflation expectations into unobserved components models. These benefits include yielding more reasonable estimates of model-based trends (e.g., Mertens (2016); Chan et al. (2018); Zaman (2024)), improved accuracy of density forecasts of inflation that adapt more quickly to perceived structural changes in the economy (e.g., Chan et al. (2018)), and tractable estimation of high-dimensional models during periods of high uncertainty (e.g., Zaman (2024)).

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. Both the Baseline model and the Single-sector model, which includes survey data, generate trend estimates that are much less volatile and less responsive to arguably large transitory spikes compared to variants without survey data, as can be seen in Figure 3 for the Baseline model and in Appendix A10 (Figure 17) for the Single-sector model.

However, somewhat discouragingly, adding survey data worsens the fit of both the Baseline model and the Single-sector model to the inflation data, as evidenced in Table 2. 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 yields predictive densities that are significantly 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.

Focusing beyond near-term predictions, the pseudo-out-of-sample forecasting evaluation exercises suggest that the baseline model yields point and density forecasts of inflation competitive with but not better than those of the model without survey data, when evaluated on the sample ending prior to the COVID-19 pandemic (see Table 3). However, for the period that includes the post-COVID inflation surge, the model without survey data yields more

accurate forecasts. This is because during the inflation surge, survey expectations remained low and stable, and including them in the model yielded lower inflation trajectories than without them. That said, since early 2024, as inflation fell to a level just below 3 percent and has remained steady since, the accuracy of the Baseline model, which uses survey data, has been comparable to that of the model without survey data.

Lastly, the dynamic posterior probability metric indicates substantial time variation in support of including survey data, as shown in Figure 6, panel (d). For example, for the periods 1990 through 2000, and 2010 through 2017, when inflation was relatively stable, the dynamic posterior probability for the inclusion of survey data (i.e., β and C_t^{π} are $\neq 0$) is above 0.5, providing support for the inclusion of survey data in the model. A closer examination of the time series of the marginal data density (one-step-ahead predictive likelihoods) and out-of-sample forecasting exercises suggests that these are also the periods when the baseline model outperforms the model variant without the survey data.

It is worth mentioning that our finding that the model fit worsens with the inclusion of survey data is in contrast to prior research by Chan et al. (2018) and Zaman (2024), who, using model formulations similar to those developed in this paper, found that adding survey expectations data neither hurts nor improves the fit of the model to the aggregate PCE inflation data when evaluated over a long sample.[9] A detailed supplementary examination reveals that the contrasting findings on model fit from prior research are due to the combination of different estimation samples, different evaluation samples for model fit comparison, and differences in the model formulations, including prior specification. For example, not shown, but the MDD computed using a specific sample period spanning 1983-2016 (similar

[9]See Table 5, page 31 in Chan et al. (2018); the model M1 (with survey data) has an MDD of -367.28 compared to -366.35 for the model M4 (without survey data). Similarly, see Table 1, panel (b) in Zaman (2024), the Base model (with survey data) has an MDD of -365.2 compared to -366.0 for the Base-NoSurvey (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.

to the period used in the baseline setup of the Chan, Clark, and Koop model) implies a much smaller gap between the fit of the baseline model and the variant without the survey data: Baseline model's MDD (1983-2016) for services inflation is -142.2 compared to the Baseline-NoSurvey model's (M7) -139.5; similarly, for goods inflation, the corresponding MDD values are -324.5 and -318.1, respectively.

We end this 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 available prior to 1990). To do so, we re-estimate our baseline model (M1), the baseline model variant without the survey data (M7), the single-sector model (M8), and a variant of the single-sector model without survey data (M10). The plots of the resulting trend inflation estimates are shown in Appendix A11. As can be seen, the 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 they appear sensible. Furthermore, as shown in Table 3 of Appendix A11, the differences in model fit between models with and without survey data, when estimated using a shorter sample, are smaller than those shown in Table 2, highlighting the role of the estimation period in influencing 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 without survey data has estimation issues, highlighting the important role of survey data not only helping provide sensible estimates of trend but also helping feasibly estimate the richer specifications of unobserved components models when faced with data limitations.

5 Conclusion

This paper proposes a new empirical model to estimate the (long-run) trend in inflation. The elements of the new model are inspired by prior research, which 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), cyclical 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.

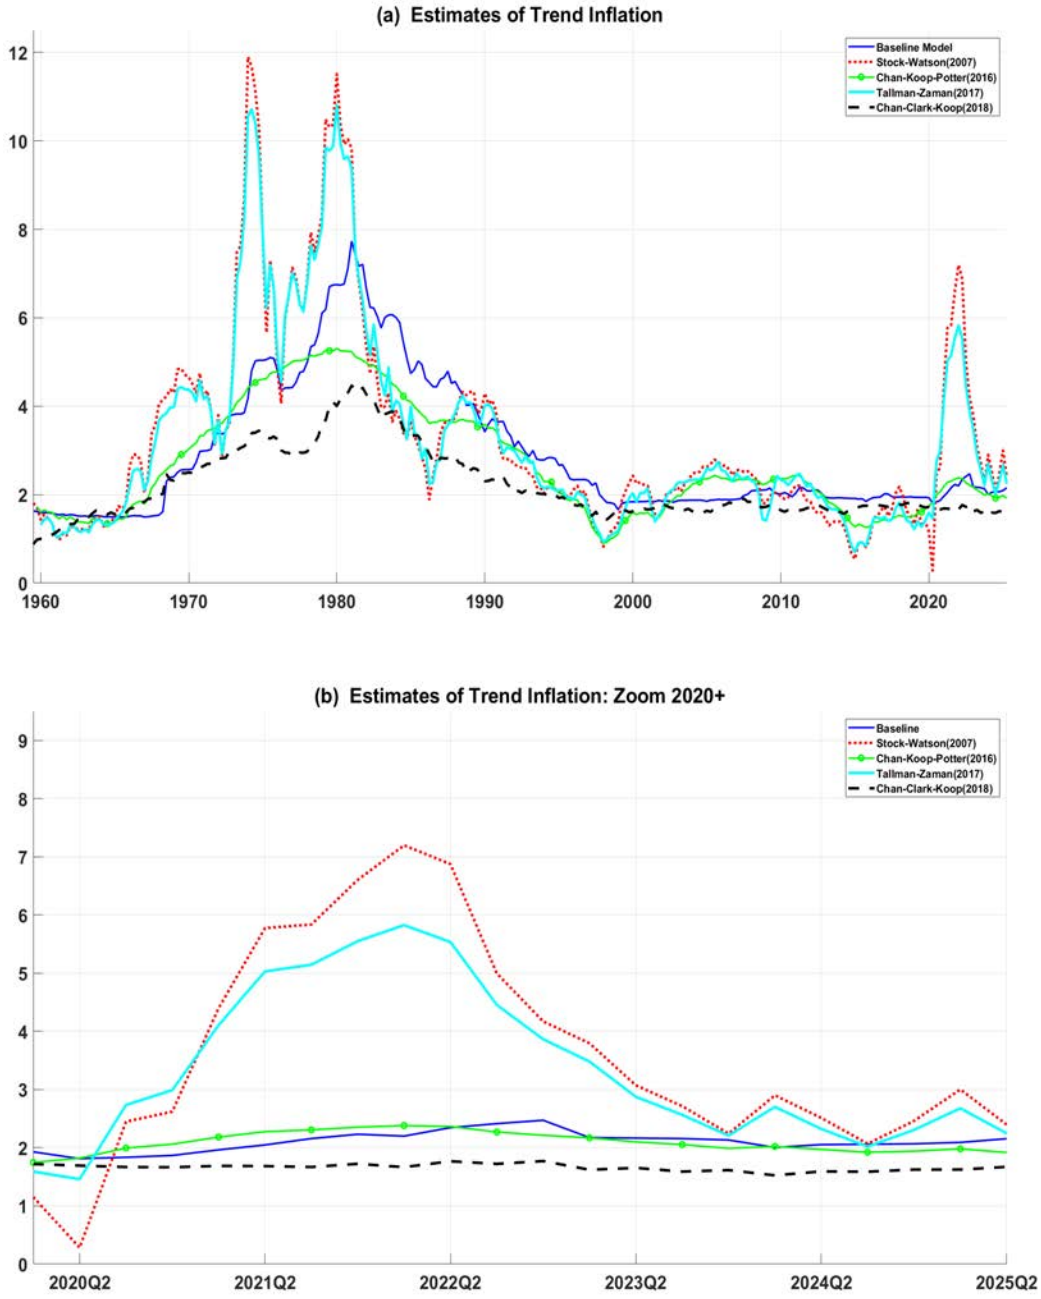
Our results indicate the importance of jointly modeling all four elements, with some elements being more informative for estimating trend inflation than others. Survey expectations have the greatest impact on trend estimates, followed by uncertainty and the use of sector-level data, and the Phillips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

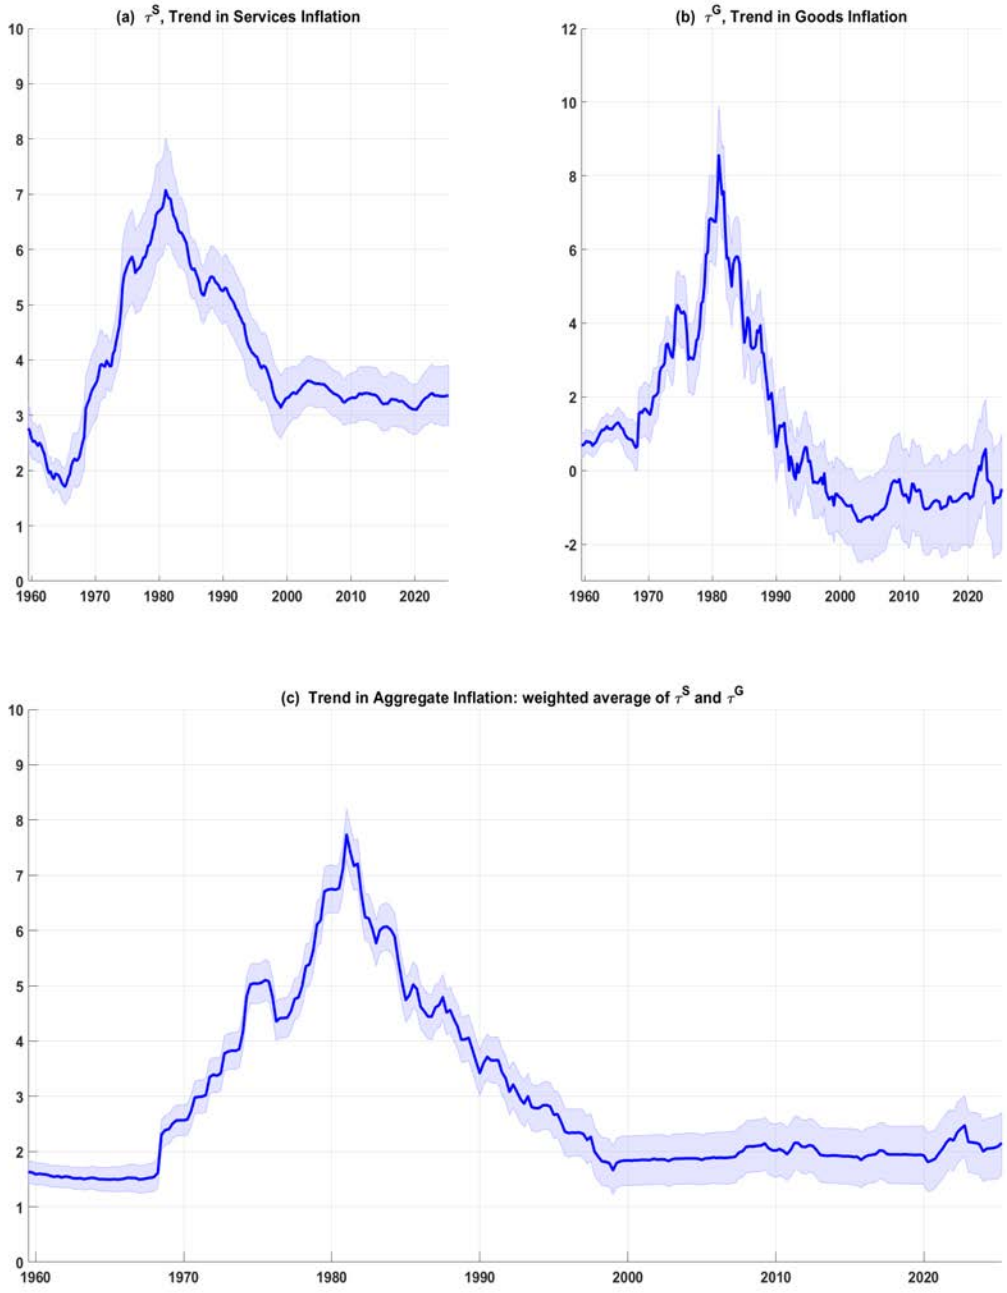
of inflation data than the models it extends. The decomposition suggests that uncertainty effects play a greater role than cyclical effects in explaining inflation fluctuations.

Figure 1: Trend Inflation: Comparison with External Models



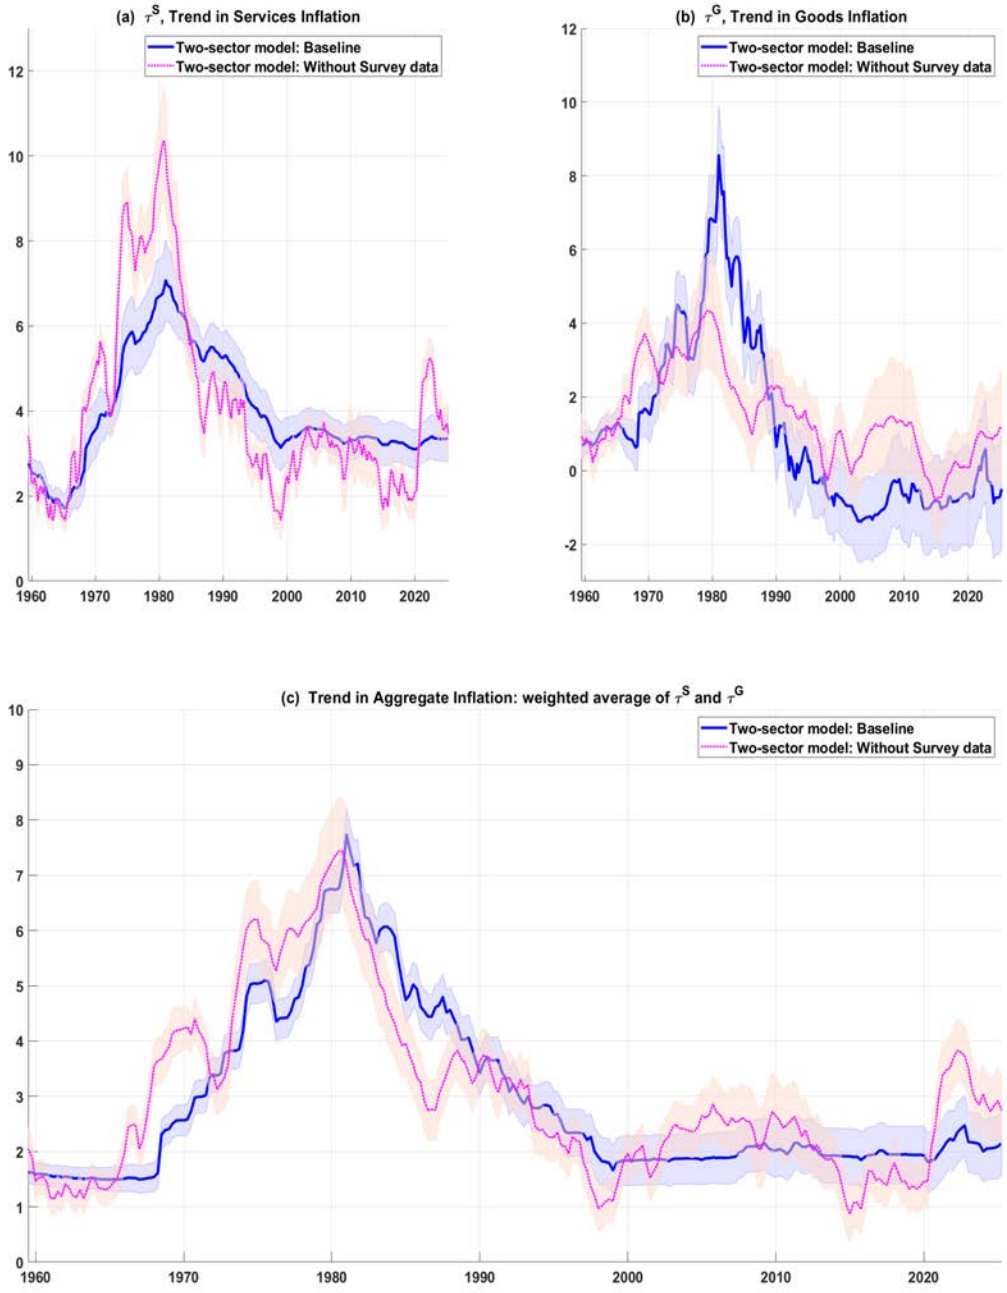
Note: Plotted are posterior mean estimates based on the full sample.

Figure 2: Trend Inflation, Baseline Model



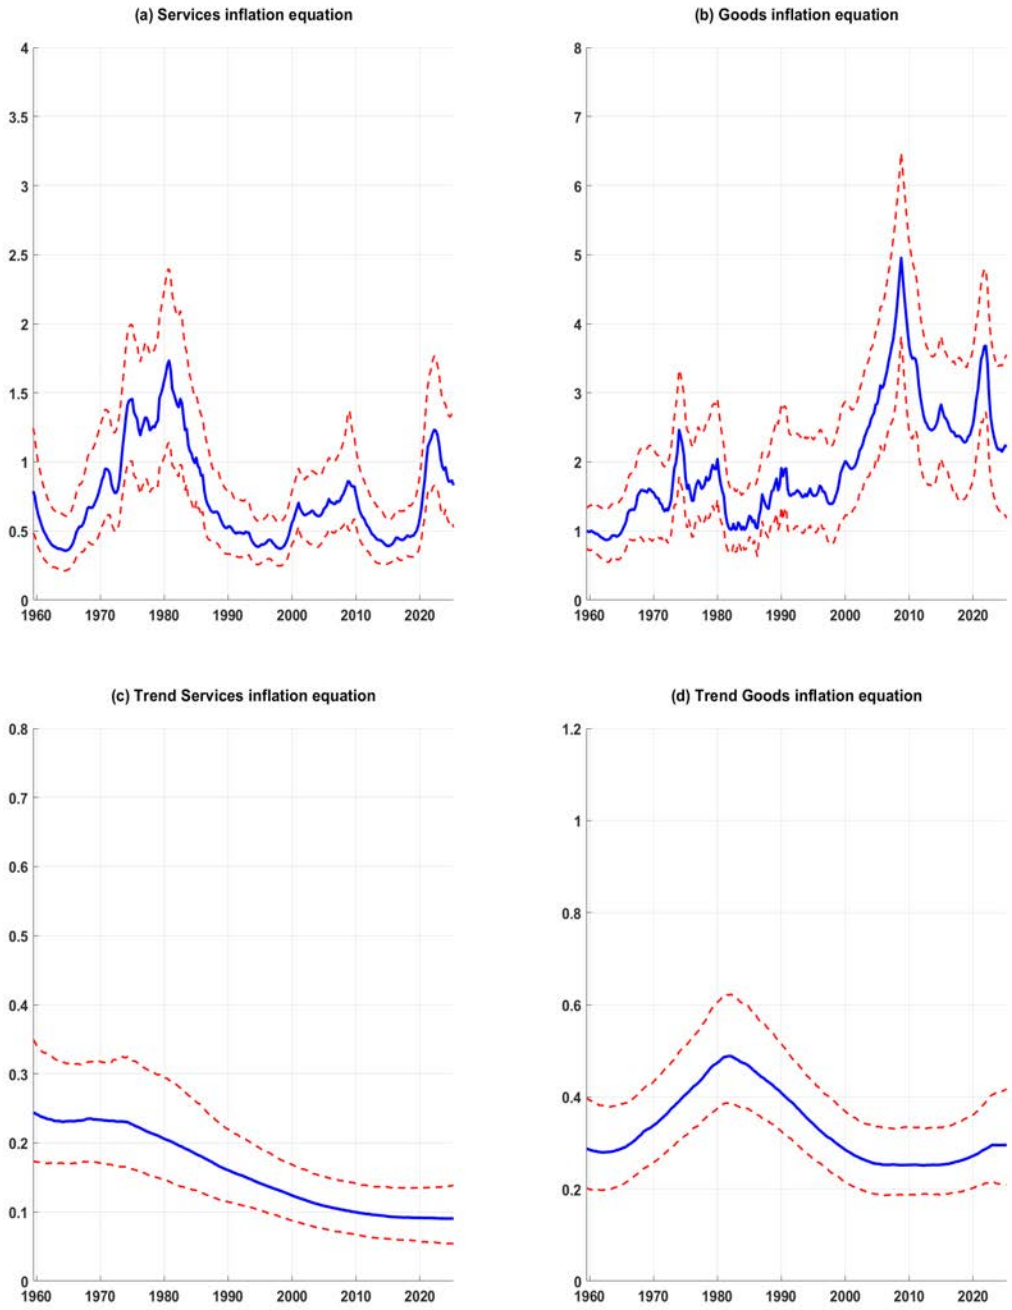
Note: The posterior estimates are based on the full sample (from 1959Q3 through 2025Q2). The shaded area represents the 68% posterior coverage interval. Panel (a), posterior estimates of the trend in services PCE inflation. Panel (b), posterior estimates of the trend in goods PCE inflation. Panel (c), 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: Trend Inflation: Baseline vs. Variant without Survey data



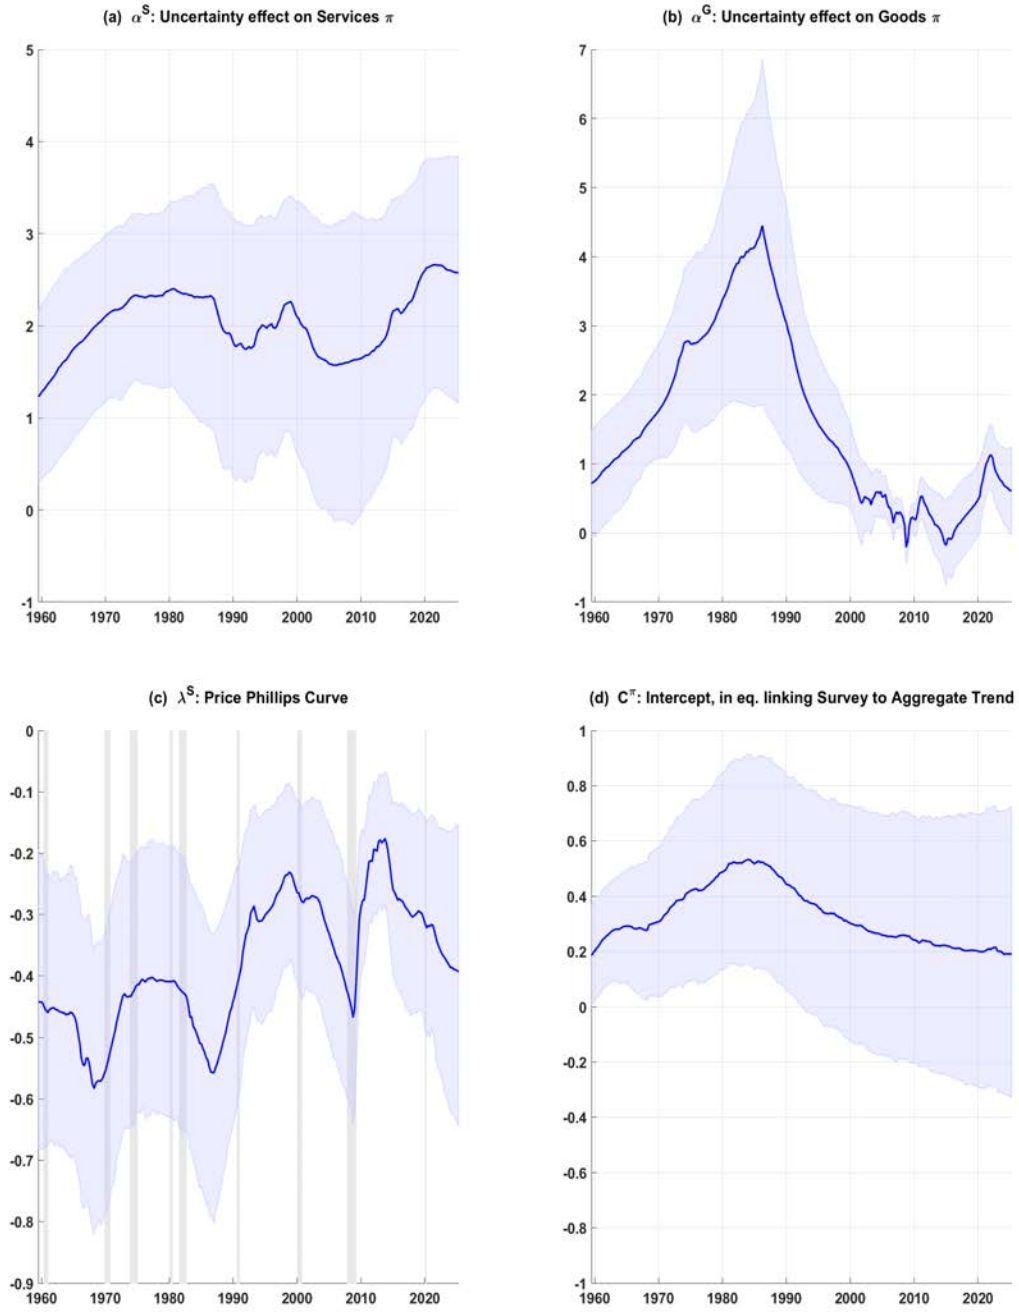
Note: The posterior estimates are based on the full sample (from 1959Q3 through 2025Q2). The shaded area represents the 68% posterior coverage interval. Panel (a), posterior estimates of the trend in services PCE inflation. Panel (b), posterior estimates of the trend in goods PCE inflation. Panel (c), 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: Estimates of Stochastic Volatility



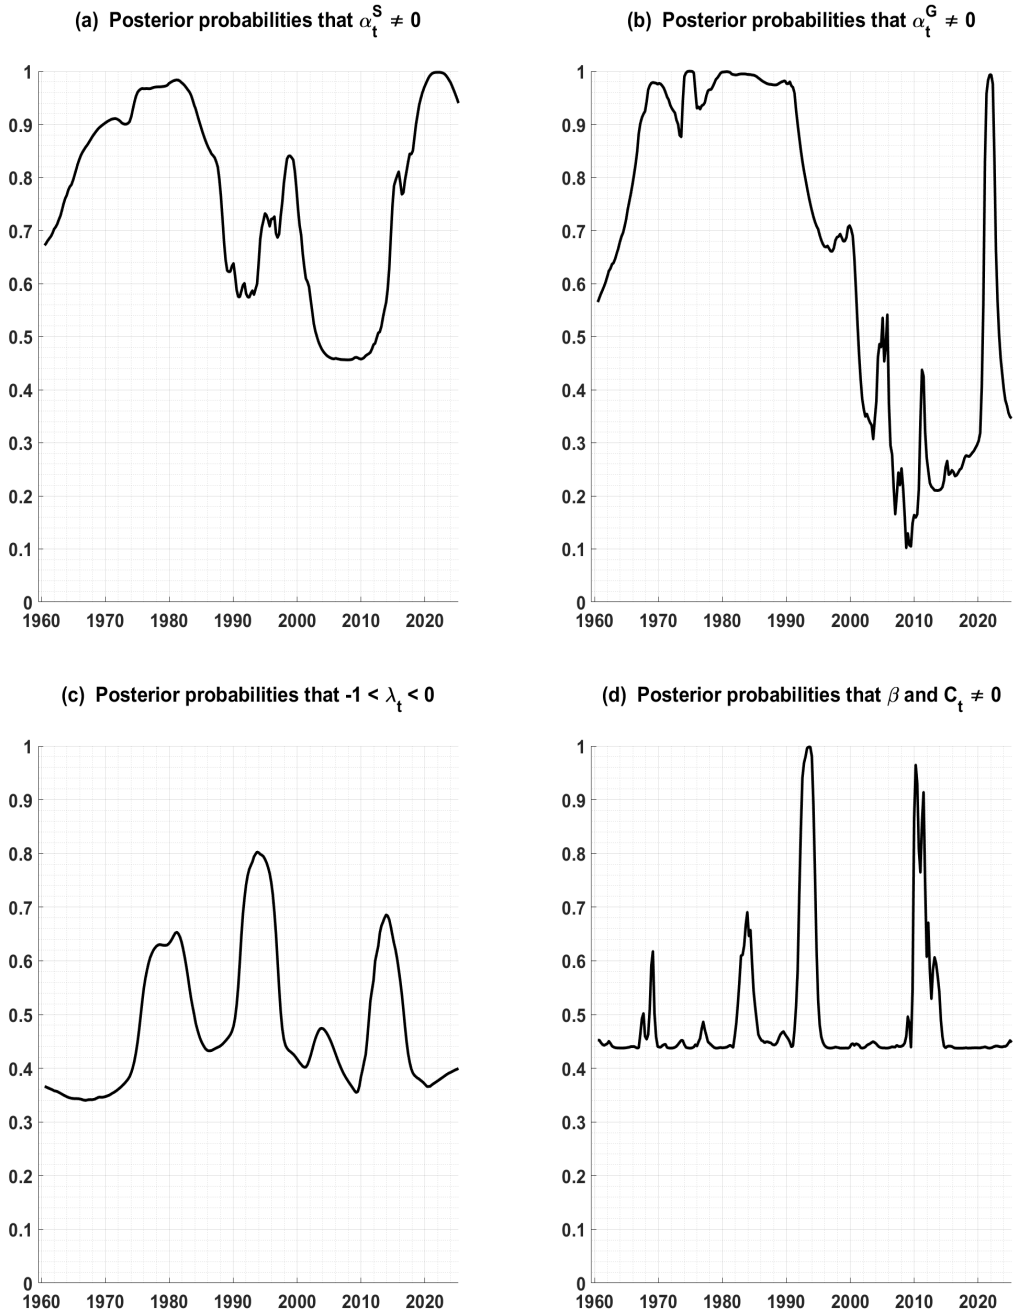
Note: Panel (a) and Panel (b) present the posterior estimates of the standard deviation (sd) of the shocks defining the services inflation and goods inflation measurement equations. Panel (c) and Panel (d) present the estimates of the sd of the shocks to the processes governing the evolution of the trend 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 5: Estimates of Time-varying Parameters



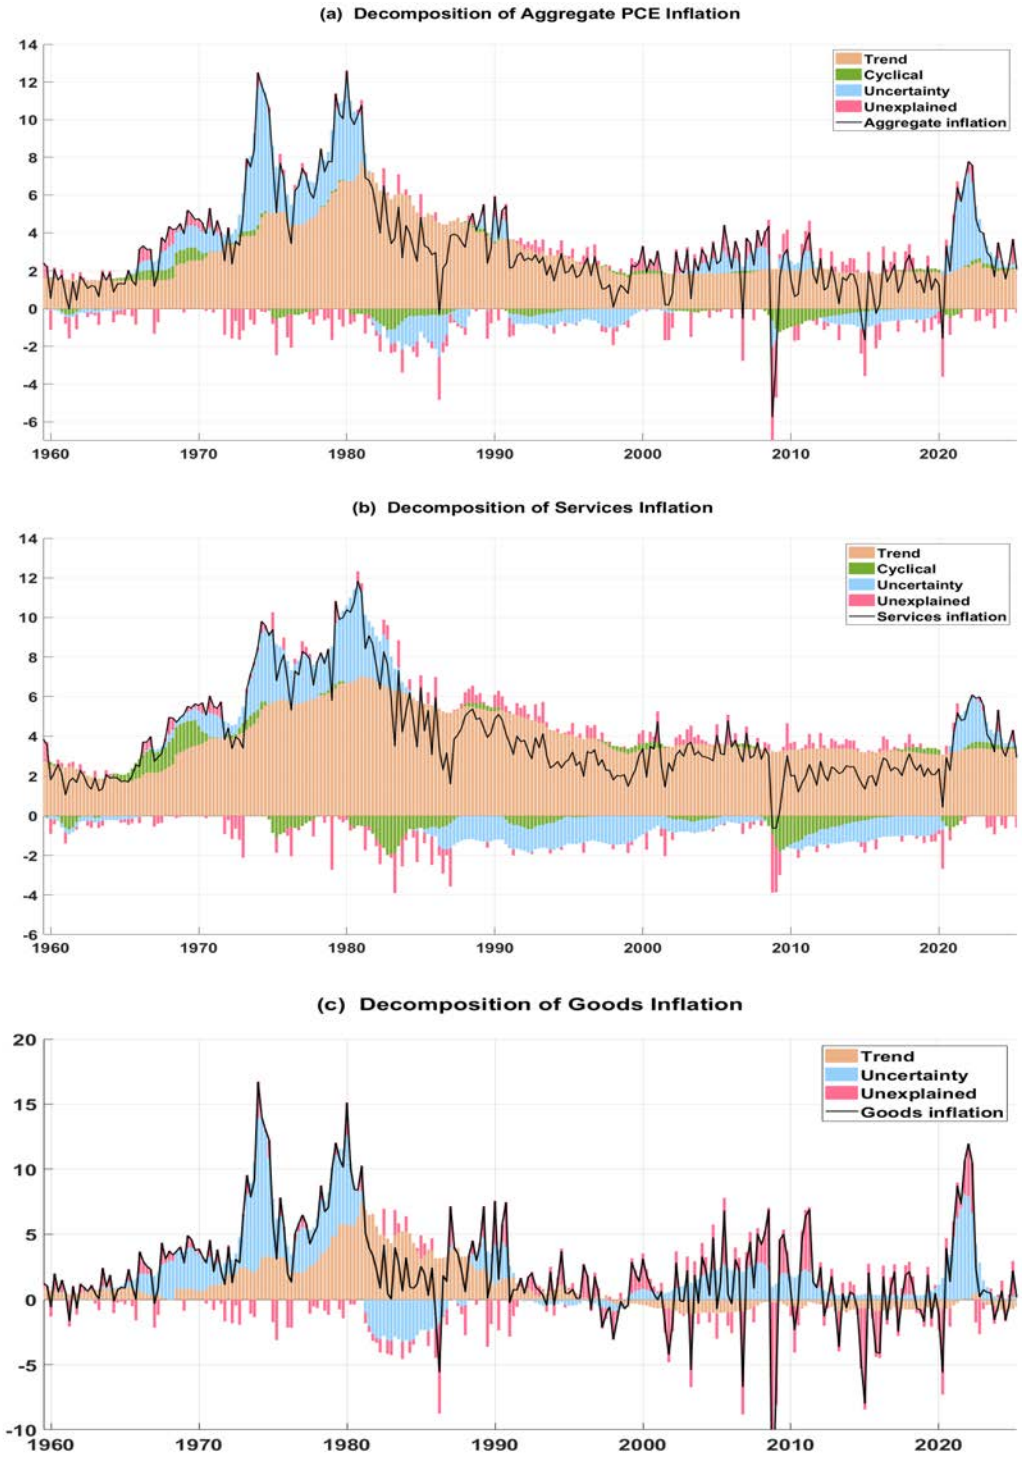
Note: Plotted are the posterior estimates of the time-varying parameters computed using the full sample (from 1959Q3 through 2025Q2). The solid lines represent the posterior mean, and the shaded area represents the 68% credible intervals.

Figure 6: 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. Panel (c), formal assessment over time of the existence of the Phillips curve relationship. Panel (d), formal assessment over time of inclusion of survey data.

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).

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: Baseline (Two-sector model) vs. Variants of Two-sector model

	1960-2025Q2, MDD		1960-2019, MDD	
Model	Services π	Goods π	Services π	Goods π
M1: Baseline	-310.4	-592.6	-276.1	-536.1
M2: Baseline but no TVP in α^S	-322.1	-600.1	-286.1	-545.6
M3: Baseline but no TVP in α^S and α^G	-319.7	-626.3	-283.5	-567.6
M4: Baseline but no TVP in λ	-312.4	-596.5	-282.5	-551.6
M5: Baseline but no TVP in α^S, α^G, and λ	-324.5	-628.4	-288.1	-569.5
M6: Baseline but No α^S, α^G, and λ	-325.7	-677.4	-286.4	-613.4
M7: Baseline but No Survey data	-282.8	-583.0	-254.7	-527.3

Panel B: Baseline (Two-sector model) vs. Single-sector model and its Variants

	MDD Aggregate PCE π	
Model	1960-2025Q2 sample	1960-2019 sample
M1: Baseline	-398.0	-358.4
M7: Baseline but No Survey data	-386.9	-349.2
M8: Single-sector	-415.8	-372.4
M9: Single-sector but No α, and λ	-455.4	-407.1
M10: Single-sector but No Survey data	-393.7	-355.3

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 single-sector model (which models directly the aggregate PCE inflation) and its variants, to facilitate comparison, the implied model fit for the Baseline model and Baseline model without survey data is also reported. The term "no TVP" means no time variation is permitted in the parameter of interest.

Table 3: Out-of-Sample Forecasting Performance

Panel A: Point forecast recursive evaluation

	1999-2025Q2, sample				1999-2019Q4, sample			
Relative RMSFE	h=2Q	4Q	8Q	12Q	2Q	4Q	8Q	12Q
Baseline/UCSV	0.92	0.99	1.07*	1.13	0.97	1.05	1.05	1.10
Baseline-NoSurvey/UCSV	0.90	0.96	0.99	0.99	0.96	1.07	1.04	1.00
Single-sector/UCSV	0.93	1.00	1.09*	1.13	0.96	1.04	1.06	1.04

Panel B: Density forecast recursive evaluation

	1999-2025Q2, sample				1999-2019Q4, sample			
Relative Log Score	h=2Q	4Q	8Q	12Q	2Q	4Q	8Q	12Q
Baseline - UCSV	0.00	0.00	0.00	-0.01***	0.00	0.00	-0.01	-0.02**
Baseline-NoSurvey - UCSV	0.02***	0.02***	0.01	0.00	0.02**	0.01**	0.01	0.00
Single-sector - UCSV	0.01***	0.01**	0.01	0.01	0.01**	0.01**	0.01**	0.01

Note: The forecasts and associated accuracy correspond to the quarterly annualized rate of aggregate PCE inflation. UCSV forecast corresponds to the forecast from the univariate unobserved-component stochastic volatility model, similar to Stock and Watson (2007). Baseline-NoSurvey forecast corresponds to the forecast from the model defined as M6, which is the Baseline model excluding survey data. Single-sector forecast corresponds to the forecast from the model defined as M8, which does not feature the disaggregates (services and goods) but includes all features of the baseline model. The top panel reports results for the point forecast accuracy (relative root mean squared errors) and the bottom panel reports the corresponding density forecast accuracy (the difference between the mean log predictive score of the model of interest and that of the benchmark model). 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.

References

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, CEPR Discussion Papers*. URL `https://cepr.org/publications/dp18993`.

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino (2019). "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 and Rodney Strachan (2012). "Estimation in non-linear non-gaussian state-space models with precision-based methods." *CAMA working paper series 2012-13*.

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.

Clark, Todd E. (2004). "An evaluation of the decline in goods inflation." *Economic Review, Federal Reserve Bank of Kansas City*, 89(Q II), pp. 19—-51. URL https://ideas.repec.org/a/fip/fedker/y2004iqiip19-51nv.89no.2.html.

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.

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*.

Federal Reserve Bank of Cleveland Working Paper Series

Online Appendix:

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

Ellis Tallman and Saeed Zaman

Working Paper No. 26-08

March 2026

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

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/

Online Appendix

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

Ellis Tallman† Saeed Zaman‡

Feb 27, 2026

*We thank Todd Clark, James Mitchell, and participants at the Advances in Econometrics Bayesian Macroeconometrics Workshop 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, Cleveland Fed, 1455 E. 6th St., Cleveland, OH 44114. Email: saeed.zaman@clev.frb.org

Contents

A1. Prior Elicitation **3**

A2. MCMC Algorithm **5**

A3. MCMC Convergence Diagnostics **20**

A4. Prior and Posterior Distributions **27**

A5.CBO Unemployment Rate Gap **29**

A6.Trend Estimates: Sensitivity to Modeling Assumptions **30**

A7.Uncertainty Estimates: Sensitivity to Modeling Assumptions **34**

A8.Comparing Parameter Estimates: Baseline vs. Model with time-invariant parameters **35**

A9.Comparing Parameter Estimates: Baseline vs. Model without Survey Expectations Data **36**

A10. Trend Estimates: Single-Sector Model vs. Single-Sector Model without Survey data **37**

A11. Trend estimates based on estimating models with data starting from 1980 onward **38**

A1. Prior Elicitation

In the table below, the notation $N(a, b)$ denotes Normal distribution with mean a, and variance b; and $IG(\nu, S)$ denotes Inverse Gamma distribution with degrees of freedom parameter ν, and scale parameter S.

Table 1: Prior settings for the Baseline Model

Parameter	**Parameter Description**	**Prior**
σ^2_{hS}	Var. of the Volatility – Services π measurement eq.	$IG(10, 0.25^2 \times 9)$
σ^2_{hG}	Var. of the Volatility – Goods π measurement eq.	$IG(10, 0.25^2 \times 9)$
$\sigma^2_{h\tau S}$	Var. of the Volatility – Trend Services π eq.	$IG(10, 0.05^2 \times 9)$
$\sigma^2_{h\tau G}$	Var. of the Volatility – Trend Goods π eq.	$IG(10, 0.05^2 \times 9)$
σ^2_{λ}	Var. of the shocks to TVP λ (Phillips curve slope)	$IG(10, 0.05^2 \times 9)$
$\sigma^2_{\alpha S}$	Var. of the shocks to TVP α^S	$IG(10, 0.1^2 \times 9)$
$\sigma^2_{\alpha G}$	Var. of the shocks to TVP α^G	$IG(10, 0.1^2 \times 9)$
$\sigma^2_{c\pi}$	Var. of the shocks to TVP C^{π}	$IG(10, 0.1^2 \times 9)$
$\sigma^2_{z\pi}$	Var. of the shocks to the eq. linking survey to π^*	$IG(10, 0.447^2 \times 9)$
β^{π}	Link between π^* and survey	$N(1, 0.05^2)$
λ_0	Initial value of λ, $t = 0$	$TN(-1, 0; 0, 1)$
α^S_0	Initial value of α^S, $t = 0$	$N(0, 1)$
α^G_0	Initial value of α^G, $t = 0$	$N(0, 1)$
C^{π}_0	Initial value of C^{π}, $t = 0$	$N(0, 0.1)$
h^{Ser}_0	Initial value of h^{Ser}, $t = 0$	$N(0, 1)$
h^{Goods}_0	Initial value of h^{Goods}, $t = 0$	$N(0, 1/5)$
$h^{\tau S}_0$	Initial value of $h^{\tau S}$, $t = 0$	$N(0, 1)$
$h^{\tau G}_0$	Initial value of $h^{\tau G}$, $t = 0$	$N(0, 1)$
τ^G_0	Initial value of τ^G, $t = 0$	$N(1, 100)$
τ^S_0	Initial value of τ^S, $t = 0$	$N(3, 100)$

Table 2: Prior settings for the Single-sector Model

Parameter	**Parameter Description**	**Prior**
σ^2_{hS}	Var. of the Volatility – aggregate π measurement eq.	$IG(10, 0.25^2 \times 9)$
$\sigma^2_{h\tau S}$	Var. of the Volatility – aggregate Trend π eq.	$IG(10, 0.05^2 \times 9)$
σ^2_{λ}	Var. of the shocks to TVP λ (Phillips curve slope)	$IG(10, 0.05^2 \times 9)$
$\sigma^2_{\alpha S}$	Var. of the shocks to TVP α	$IG(10, 0.1^2 \times 9)$
$\sigma^2_{c\pi}$	Var. of the shocks to TVP C^{π}	$IG(10, 0.1^2 \times 9)$
$\sigma^2_{z\pi}$	Var. of the shocks to the eq. linking survey to π^*	$IG(10, 0.447^2 \times 9)$
β^{π}	Link between π^* and survey	$N(1, 0.05^2)$
λ_0	Initial value of λ, $t = 0$	$TN(-1, 0; 0, 1)$
α_0	Initial value of α, $t = 0$	$N(0, 1)$
C^{π}_0	Initial value of C^{π}, $t = 0$	$N(0, 0.1)$
h_0	Initial value of h, $t = 0$	$N(0, 1)$
h^{τ}_0	Initial value of h^{τ}, $t = 0$	$N(0, 1)$
τ_0	Initial value of τ, $t = 0$	$N(2, 100)$

A2. MCMC Algorithm

The Baseline model is estimated by sampling from its joint posterior distribution, which entails sequentially drawing from a set of conditional posterior densities, some standard and some non-standard.

Collect all the time-invariant model parameters into θ:
$\theta = (\sigma^2_{hS}, \sigma^2_{hG}, \sigma^2_{h\tau S}, \sigma^2_{h\tau G}, \sigma^2_{\lambda}, \sigma^2_{\alpha S}, \sigma^2_{\alpha G}, \sigma^2_{c\pi}, \sigma^2_{z\pi}, \beta^{\pi})$

Let's denote $\bullet$ as representing all other model parameters. Then, the sequential steps to draw from the conditional posterior densities are:
1. $p(\alpha^S|Y, \bullet)$
2. $p(\alpha^G|Y, \bullet)$
3. $p(\lambda|Y, \bullet)$
4. $p(\tau^S|Y, \bullet)$
5. $p(\tau^G|Y, \bullet)$
6. $p(h^{Ser}|Y, \bullet)$
7. $p(h^{Goods}|Y, \bullet)$
8. $p(h^{\tau S}|Y, \bullet)$
9. $p(h^{\tau G}|Y, \bullet)$
10. $p(C^{\pi}|Y, \bullet)$
11. $p(\theta|Y, \bullet)$

In the steps below, $Y = \pi^S$ when sampling parameters defining dynamics of services inflation, and $Y = \pi^G$ when sampling parameters defining the dynamics of goods inflation.

Step 1. Derive the conditional distribution $p(\alpha^S|Y, \bullet)$

The information about α^S comes from two sources. Below, we derive an expression for each of these two sources.

First, some notation,

$\tilde{\pi}^S_t = \pi^S_t - \tau^S_t$
$\tilde{u}_t = U_t - U^*_t$
$\tilde{v}_t = e^{h^{Ser}_t} - e^{\bar{h}^{Ser}_t}$

The first source is the services inflation measurement equation. Rewrite it in a matrix notation,

$$N^{\alpha s} = X_v \alpha^S + \varepsilon^S \qquad \varepsilon^S \sim N(0, \Omega_s), \quad where \ \Omega_s = diag(e^{h_1^{Ser}}, e^{h_2^{Ser}}, ..., e^{h_T^{Ser}}) \tag{1}$$

where,
$X_v = diag(\tilde{v}_1, ..., \tilde{v}_T)$
$N^{\alpha s} = (\tilde{\pi}_1^S - \lambda_1 \tilde{u}_1, ..., \tilde{\pi}_T^S - \lambda_T \tilde{u}_T)'$

Ignoring any terms not involving α^S, the likelihood is

$$log \ \ p(\pi^S | \alpha^S, \bullet) \propto -\tfrac{1}{2}(N^{\alpha s} - X_v \alpha^S)' \Omega_s^{-1} (N^{\alpha s} - X_v \alpha^S)$$

The second source comes from the state equation for α^S. We rewrite it in a matrix notation as follows,

$$H\alpha^S = \varepsilon^{\alpha s} \qquad \varepsilon^{\alpha s} \sim N(0, \Omega_{\alpha s}), \quad where \ \Omega_{\alpha s} = diag(\omega_{\alpha s}^2, \sigma_{\alpha s}^2, ..., \sigma_{\alpha s}^2) \tag{2}$$

Ignoring any terms not involving α^S, the prior density for α^S is given by

$$log \ \ p(\alpha^S | \sigma_{\alpha s}^2, \Omega_{\alpha s}) \propto -\tfrac{1}{2}(\alpha^S)' H' \Omega_{\alpha s}^{-1} H(\alpha^S)$$

Combining the above two conditional densities we obtain,

$$log \ \ p(\alpha^S | Y, \bullet) \propto -\tfrac{1}{2}(\alpha^S - \hat{\alpha}^S)' D_{\alpha s}^{-1} (\alpha^S - \hat{\alpha}^S)$$

where,

$$D_{\alpha s} = (H' \Omega_{\alpha s}^{-1} H + X_v' \Omega_s^{-1} X_v)^{-1}$$

$$\hat{\alpha}^S = D_{\alpha s}(X_v' \Omega_s^{-1} N^{\alpha s})$$

The candidate draws are sampled from $N(\hat{\alpha}^S, D_{\alpha s})$ using the precision-based algorithm.

Step 2. Derive the conditional distribution $p(\alpha^G | Y, \bullet)$

The information about α^G comes from two sources. Below, we derive an expression for each of these two sources.

First, some notation,

$\tilde{\pi}_t^G = \pi_t^G - \tau_t^G$
$\tilde{u}_t = U_t - U_t^*$
$\tilde{vg}_t = e^{h_t^{Goods}} - e^{\bar{h}_t^{Goods}}$

The first source is the goods inflation measurement equation. Rewrite it in a matrix notation,

$$N^{\alpha g} = X_{vg}\alpha^G + \varepsilon^G \qquad \varepsilon^G \sim N(0, \Omega_g), \quad where\ \ \Omega_g = diag(e^{h_1^{Goods}}, e^{h_2^{Goods}}, ..., e^{h_T^{Goods}}) \quad (3)$$

where,
$X_{vg} = diag(\tilde{vg}_1, ..., \tilde{vg}_T)$
$N^{\alpha g} = (\tilde{\pi}_1^G, ..., \tilde{\pi}_T^G)'$

Ignoring any terms not involving α^S, the likelihood is

$$log\ \ p(\pi^G|\alpha^G, \bullet) \propto -\tfrac{1}{2}(N^{\alpha g} - X_{vg}\alpha^G)'\Omega_g^{-1}(N^{\alpha g} - X_{vg}\alpha^G)$$

The second source comes from the state equation for α^G. We rewrite it in a matrix notation as follows,

$$H\alpha^G = \varepsilon^{\alpha g} \qquad \varepsilon^{\alpha g} \sim N(0, \Omega_{\alpha g}), \quad where\ \ \Omega_{\alpha g} = diag(\omega_{\alpha g}^2, \sigma_{\alpha g}^2, ..., \sigma_{\alpha g}^2) \qquad (4)$$

Ignoring any terms not involving α^G, the prior density for α^G is given by

$$log\ \ p(\alpha^G|\sigma_{\alpha g}^2, \Omega_{\alpha g}) \propto -\tfrac{1}{2}(\alpha^G)'H'\Omega_{\alpha g}^{-1}H(\alpha^G)$$

Combining the above two conditional densities we obtain,

$$log\ \ p(\alpha^G|Y, \bullet) \propto -\tfrac{1}{2}(\alpha^G - \hat{\alpha}^G)'D_{\alpha g}^{-1}(\alpha^G - \hat{\alpha}^G)$$

where,
$D_{\alpha g} = (H'\Omega_{\alpha g}^{-1}H + X_{vg}'\Omega_g^{-1}X_{vg})^{-1}$

$$\hat{\alpha}^G = D_{\alpha g}(X'_{vg}\Omega_g^{-1}N^{\alpha g})$$

The candidate draws are sampled from $N(\hat{\alpha}^G, D_{\alpha g})$ using the precision-based algorithm.

Step 3. Derive the conditional distribution $p(\lambda|Y, \bullet)$

The information about λ comes from two sources. Below, we derive an expression for each of these two sources.

First, some notation,

$\tilde{\pi}_t^S = \pi_t^S - \tau_t^S$
$\tilde{u}_t = U_t - U_t^*$
$\tilde{v}_t = e^{h_t^{Ser}} - e^{\overline{h}_t^{Ser}}$
$NW = (\tilde{\pi}_1^S - \alpha_1^S\tilde{v}_1, ..., \tilde{\pi}_T^S - \alpha_T^S\tilde{v}_T)'$

The first source is the services inflation measurement equation. Rewrite it in a matrix notation,

$$NW = X_u\lambda + \varepsilon^S \qquad \varepsilon^S \sim N(0, \Omega_s) \tag{5}$$

where,
$X_u = diag(\tilde{u}_1, ..., \tilde{u}_T)$

Ignoring any terms not involving λ, the likelihood is

$$log\ \ p(\pi^S|\lambda, \bullet) \propto -\tfrac{1}{2}(NW - X_u\lambda)'\Omega_s^{-1}(NW - X_u\lambda)$$

The second source comes from the state equation for λ. We rewrite it in a matrix notation as follows,

$$H\lambda = \varepsilon^\lambda \qquad \varepsilon^\lambda \sim N(0, \Omega_\lambda), \quad where\ \ \Omega_\lambda = diag(\omega_\lambda^2, \sigma_\lambda^2, ..., \sigma_\lambda^2) \tag{6}$$

$-1 < \lambda_t < 0$ for t=1,....,T

Ignoring any terms not involving λ, the prior density for λ is given by

$$log\ \ p(\lambda|\sigma_\lambda^2, \Omega_\lambda) \propto -\tfrac{1}{2}(\lambda)'H'\Omega_\lambda^{-1}H(\lambda) + g_\lambda(\lambda, \sigma_\lambda^2)$$

where,

$$g_\lambda(\lambda, \sigma_\lambda^2) = -\sum_{t=2}^{T} log\left(\Phi\left(\frac{0-\lambda_{t-1}}{\sigma_\lambda}\right) - \Phi\left(\frac{-1-\lambda_{t-1}}{\sigma_\lambda}\right)\right)$$

Combining the above two conditional densities we obtain,

$$log\ \ p(\lambda|Y, \bullet) \propto -\tfrac{1}{2}(\lambda - \hat{\lambda})' D_\lambda^{-1}(\lambda - \hat{\lambda}) + g_\lambda(\lambda, \sigma_\lambda^2)$$

where,

$D_\lambda = (H'\Omega_\lambda^{-1}H + X_u'\Omega_s^{-1}X_u)^{-1}$

$\hat{\lambda} = D_\lambda(X_u'\Omega_s^{-1}NW)$

The addition of the term $g_\lambda(\lambda, \sigma_\lambda^2)$ leads to a non-standard density. Accordingly, we sample λ using an independence-chain Metropolis-Hastings (MH) procedure. This involves first generating candidate draws from $N(\hat{\lambda}, D_\lambda)$ using the precision-based algorithm that are then accepted or rejected based on the accept-reject Metropolis-Hastings (ARMH) algorithm (discussed in Chan and Strachan, 2012).

Step 4. Derive the conditional distribution $p(\tau^S|Y, \bullet)$

The information about τ^S comes from three sources. Below, we derive an expression for each of these sources.

The first source is the services inflation measurement equation. Rewrite it in a matrix notation as,

$$\pi^S = \Phi_{\tau s} + \tau^S + \varepsilon^S \qquad \varepsilon^S \sim N(0, \Omega_S), \quad where\ \ \Omega_S = diag(e^{h_1^{Ser}}, e^{h_2^{Ser}}, ..., e^{h_T^{Ser}}) \qquad (7)$$

where, $\Phi_{\tau s} = \begin{pmatrix} \lambda_1(U_1 - U_1^*) + \alpha_1^S(e^{h_1^{Ser}} - e^{\overline{h}_1^{Ser}}) \\ \lambda_1(U_2 - U_2^*) + \alpha_2^S(e^{h_2^{Ser}} - e^{\overline{h}_2^{Ser}}) \\ \vdots \\ \lambda_1(U_T - U_T^*) + \alpha_T^S(e^{h_T^{Ser}} - e^{\overline{h}_T^{Ser}}) \end{pmatrix}$,

The likelihood is

$log\ \ p(\pi^S|U, U^*, \bullet) \propto -\frac{1}{2}\iota_T h^{Ser} - \frac{1}{2}(\pi^S - \Phi_{\tau s})'\Omega_S^{-1}(\pi^S - \Phi_{\tau s})$

where ι_T is a $T \times 1$ column vector of ones.

The second source of information is from the state equation of τ^S. Rewrite it in a matrix notation,

$$H\tau^S = \mu_s + \varepsilon^{\tau S} \qquad \varepsilon^{\tau S} \sim N(0, \Omega_{\tau s}), \quad where\ \ \Omega_{\tau s} = diag(\omega_{\tau s}^2, e^{h_1^{\tau S}}, ..., e^{h_{T-1}^{\tau S}}) \tag{8}$$

where, $\mu_s = \begin{pmatrix} \tau_0^S \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$

That is, the prior density for τ^S is given by

$$p(\tau^S|h^{\tau S}) \propto -\tfrac{1}{2}(\tau^S - H^{-1}\mu_s)' H' \Omega_{\tau s}^{-1} H(\tau^S - H^{-1}\mu_s)$$

The third source of information comes from the measurement equation that links surveys to τ^*. Rewrite the equation in a matrix notation,

$$F^{\pi s} = \beta^\pi \tau^S + \varepsilon^{z\pi} \qquad \varepsilon^{z\pi} \sim N(0, \Omega_{z\pi}) \tag{9}$$

where,

$f_t^{\pi s} = \frac{Z_t^\pi - C^\pi - \beta^\pi(w_t^G \tau^G)}{w_t^S}$
$F^{\pi s} = (f_1^{\pi s}, f_2^{\pi s}, ..., f_T^{\pi s})'$
$\Omega_{z\pi} = diag(\sigma_{z\pi}^2, ..., \sigma_{z\pi}^2)$
Ignoring any terms not involving τ^S, we have

$$log\ \ p(Z^\pi|\tau^S, \bullet) \propto -\tfrac{1}{2}(F^{\pi s} - \beta^\pi \tau^S)' \Omega_{z\pi}^{-1}(F^{\pi s} - \beta^\pi \tau^S)$$

Combining the above three conditional densities we obtain,

$$log\ \ p(\tau^S|Y,\bullet) \propto -\tfrac{1}{2}(\tau^S-\hat{\tau}^S)'D_{\tau s}^{-1}(\tau^S-\hat{\tau}^S)$$

where,

$$D_{\tau s} = (\Omega_s^{-1} + H'\Omega_{\tau s}^{-1}H + (\beta^\pi)^2\Omega_{z\pi}^{-1})^{-1}$$

$$\hat{\tau}^S = D_{\tau s}(\Omega_s^{-1}(\pi^S - \Phi_{\tau s}) + H'\Omega_{\tau s}^{-1}\mu_s + \beta^\pi\Omega_{z\pi}^{-1}F^{\pi s})$$

The candidate draws are sampled from $N(\hat{\tau}^S, D_{\tau s})$ using the precision-based algorithm.

Step 5. Derive the conditional distribution $p(\tau^G|Y,\bullet)$

The information about τ^G comes from three sources. Below, we derive an expression for each of these sources.

The first source is the services inflation measurement equation. Rewrite it in a matrix notation as,

$$\pi^S = \Phi_{\tau g} + \tau^G + \varepsilon^G \qquad \varepsilon^G \sim N(0,\Omega_G), \quad where\ \ \Omega_G = diag(e^{h_1^{Goods}}, e^{h_2^{Goods}}, ..., e^{h_T^{Goods}}) \tag{10}$$

where, $\Phi_{\tau g} = \begin{pmatrix} \alpha_1^G(e^{h_1^{Goods}} - e^{\bar{h}_1^{Goods}}) \\ \alpha_2^G(e^{h_2^{Goods}} - e^{\bar{h}_2^{Goods}}) \\ \vdots \\ \alpha_T^G(e^{h_T^{Goods}} - e^{\bar{h}_T^{Goods}}) \end{pmatrix}$,

The likelihood is

$$log\ \ p(\pi^G|U,U^*,\bullet) \propto -\tfrac{1}{2}\iota_T h^{Goods} - \tfrac{1}{2}(\pi^G-\Phi_{\tau g})'\Omega_G^{-1}(\pi^G-\Phi_{\tau g})$$

where ι_T is a $T\times 1$ column vector of ones.

The second source of information is from the state equation of τ^G. Rewrite it in a matrix notation,

$$H\tau^G = \mu_g + \varepsilon^{\tau G} \qquad \varepsilon^{\tau G} \sim N(0,\Omega_{\tau g}), \quad where\ \ \Omega_{\tau g} = diag(\omega_{\tau g}^2, e^{h_1^{\tau G}}, ..., e^{h_{T-1}^{\tau G}}) \tag{11}$$

where, $\mu_g = \begin{pmatrix} \tau_0^G \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$

That is, the prior density for τ^G is given by

$$p(\tau^G|h^{\tau G}) \propto -\tfrac{1}{2}(\tau^G - H^{-1}\mu_g)' H' \Omega_{\tau g}^{-1} H(\tau^G - H^{-1}\mu_g)$$

The third source of information comes from the measurement equation that links surveys to τ^*. Rewrite the equation in a matrix notation,

$$F^{\pi g} = \beta^\pi \tau^G + \varepsilon^{z\pi} \qquad \varepsilon^{z\pi} \sim N(0, \Omega_{z\pi}) \tag{12}$$

where,

$f_t^{\pi g} = \frac{Z_t^\pi - C^\pi - \beta^\pi (w_t^S \tau^S)}{w_t^G}$
$F^{\pi g} = (f_1^{\pi g}, f_2^{\pi g}, ..., f_T^{\pi g})'$
$\Omega_{z\pi} = diag(\sigma_{z\pi}^2, ..., \sigma_{z\pi}^2)$
Ignoring any terms not involving τ^G, we have

$$log\ \ p(Z^\pi|\tau^G, \bullet) \propto -\tfrac{1}{2}(F^{\pi g} - \beta^\pi \tau^G)' \Omega_{z\pi}^{-1} (F^{\pi g} - \beta^\pi \tau^G)$$

Combining the above three conditional densities we obtain,

$$log\ \ p(\tau^G|Y, \bullet) \propto -\tfrac{1}{2}(\tau^G - \hat{\tau}^G)' D_{\tau g}^{-1} (\tau^G - \hat{\tau}^G)$$

where,

$$D_{\tau g} = (\Omega_G^{-1} + H' \Omega_{\tau g}^{-1} H + (\beta^\pi)^2 \Omega_{z\pi}^{-1})^{-1}$$

$$\hat{\tau}^G = D_{\tau g}(\Omega_G^{-1}(\pi^G - \Phi_{\tau g}) + H' \Omega_{\tau g}^{-1} \mu_s + \beta^\pi \Omega_{z\pi}^{-1} F^{\pi g})$$

The candidate draws are sampled from $N(\hat{\tau}^G, D_{\tau g})$ using the precision-based algorithm.

Step 6. Derive the conditional distribution $p(h^{Ser}|Y, \bullet)$

This is a high-dimensional and non-standard density. To derive the conditional density $p(h^{Ser}|Y, \bullet)$, we follow the steps detailed in Chan (2017).

a. Obtain the gradient and Hessian of the log-density $log\ \ p(h^{Ser}|Y, \alpha^S, \sigma^2_{hs}, \bullet)$

According to the Bayes's theorem

$$log\ \ p(h^{Ser}|Y, \alpha^S, \sigma^2_{hs}, h^{Ser}_0) \propto p(Y|\alpha^S, h^{Ser}) \quad p(h^{Ser}|\sigma^2_{hs}, h^{Ser}_0)$$

The prior density of h^{Ser}, $p(h^{Ser}|\sigma^2_{hs}, h^{Ser}_0)$ is Gaussian with log-density

$$log\ \ p(h^{Ser}|\sigma^2_{hs}, h^{Ser}_0) = \tfrac{-1}{2\sigma^2_{hs}}\ \ (h^{Ser'}H'Hh^{Ser} - 2h^{Ser'}H'Hh^{Ser}_0\iota_T) + c_1$$

where c_1 is a constant independent of h^{Ser}.

The gradient and Hessian of the above log conditional density is

$$\tfrac{d}{dh^{Ser}}\ \ log\ \ p(h^{Ser}|\sigma^2_{hs}, h^{Ser}_0) = \tfrac{-1}{\sigma^2_{hs}}H'H(h^{Ser} - h^{Ser}_0\iota_T)$$

$$\tfrac{d^2}{dh^{Ser}dh^{Ser'}}\ \ log\ \ p(h^{Ser}|\sigma^2_{hs}, h^{Ser}_0) = \tfrac{-1}{\sigma^2_{hs}}H'H$$

Next, let f and G denote, respectively, the gradient and the negative Hessian of $log\ \ p(Y|\alpha^S, h^{Ser})$

$f = \tfrac{d}{dh^{Ser}}\ \ log\ \ p(Y|\alpha^S, h^{Ser})$
$G = \tfrac{d^2}{dh^{Ser}dh^{Ser'}}\ \ log\ \ p(Y|\alpha^S, h^{Ser})$

Note that the log conditional likelihood $log\ \ p(Y|\alpha^S, h^{Ser})$ has the form
$log\ \ p(Y|\alpha^S, h^{Ser}) = \sum_{t=1}^{T} log\ \ p(\pi^S_t|\alpha^S_t, h^{Ser}_t)$.

The t-th element of f is simply $f_t = \frac{d\ \ log\ \ p(\pi^S_t|\alpha^S_t, h^{Ser}_t)}{dh^{Ser}_t}$

And the negative Hessian G is diagonal, where the t-th diagonal element is
$G_t = \frac{-d^2\ \ log\ \ p(\pi^S_t|\alpha^S_t, h^{Ser}_t)}{(dh^{Ser}_t)^2}$

Putting all these together, the gradient and Hessian of $log\ \ p(h^{Ser}|Y, \alpha^S, \sigma^2_{hs}, h^{Ser}_0)$ are

$$S(h^{Ser}) = \frac{d}{dh^{Ser}} log\ \ p(h^{Ser}|Y, \alpha^S, \sigma^2_{hs}, h_0^{Ser}) = \frac{-1}{\sigma^2_{hs}} H'H(h^{Ser} - h_0^{Ser}\iota_T) + f$$

$$F(h^{Ser}) = \frac{d^2}{dh^{Ser} dh^{Ser\prime}} log\ \ p(h^{Ser}|Y, \alpha^S, \sigma^2_{hs}, h_0^{Ser}) = -(\frac{1}{\sigma^2_{hs}} H'H + G)$$

Solve the above using the Newton-Raphson method:

1. Initialize the algorithm with $h^{Ser} = h^{Ser(1)}$ for some constant vector $h^{Ser(1)}$

2. For t=1,2,..., use $h^{Ser} = h^{Ser(t)}$ in the evaluation of $S(h^{Ser})$ and $F(h^{Ser})$, and compute the Newton-Raphson recursion $h^{Ser(t+1)} = h^{Ser(t)} - F(h^{Ser(t)})S(h^{Ser(t)})$

3. Repeat this procedure until some convergence criterion is reached, e.g., when $||h^{Ser(t+1)} - h^{Ser(t)||} < \epsilon$ for some prefixed tolerance level ϵ.

The convergence is very fast because the Hessian is globally negative definite.

Let $\hat{h}^{Ser}$ be the mode obtained from the Newton-Raphson method and let K_{hs} denote the negative Hessian evaluated at the mode, i.e., $K_{hs} = \frac{H'H}{\sigma^2_{hs}} + G$

Then use $N(\hat{h}^{Ser}, K_{hs}^{-1})$ as the proposal density in an independence-chain MH step. Since the precision matrix is again a band matrix, we rely on the precision sampler to obtain candidate draws from $N(\hat{h}^{Ser}, K_{hs}^{-1})$ efficiently.

The candidate draw is then accepted with a probability, which is computed in the MH step.

Step 7. Derive the conditional distribution $p(h^{Goods}|Y, \bullet)$

Same procedure as discussed in the previous step, but applied to goods inflation.

Step 8. Derive the conditional distribution $p(h^{\tau S}|Y, \bullet)$

Following, Chan, Koop, and Potter (2013; 2016), we draw $h^{\tau S}$ using the accept-reject independence-chain Metropolis Hastings (ARMH) algorithm of Chan and Strachan (2012; page 32-34).

Step 9. Derive the conditional distribution $p(h^{\tau G}|Y, \bullet)$

As done in the previous step, we follow Chan, Koop, and Potter (2013; 2016) to draw $h^{\tau G}$ using the accept-reject independence-chain Metropolis Hastings (ARMH) algorithm of Chan and Strachan (2012; page 32-34).

Step 10. Derive the conditional distribution $p(C^{\pi}|Y, \bullet)$

The information about C^{π} comes from two sources. Below, we derive an expression for each of these two sources.

The first source is the measurement equation linking survey to τ^*. Rewrite it in a matrix notation,

$$N^{z\pi} = C^{\pi} + \varepsilon^{z\pi} \qquad \varepsilon^{z\pi} \sim N(0, \Omega_{z\pi}) \tag{13}$$

where,

$n_t^{z\pi} = Z_t^{\pi} - \beta^{\pi}\tau^*$
$N^{z\pi} = (n_1^{z\pi}, n_2^{z\pi}, ..., n_T^{z\pi})'$
$\Omega_{z\pi} = diag(\sigma_{z\pi}^2, ..., \sigma_{z\pi}^2)$

Ignoring any terms not involving C^{π}, the likelihood is

$$log\ \ p(Z^{\pi}|C^{\pi}, \bullet) \propto -\tfrac{1}{2}(N^{z\pi} - C^{\pi})'\Omega_{z\pi}^{-1}(N^{z\pi} - C^{\pi})$$

The second source comes from the state equation for C^{π}. We rewrite it in a matrix notation as follows,

$$HC^{\pi} = \mu_{c\pi} + \varepsilon^{c\pi} \qquad \varepsilon^{c\pi} \sim N(0, \Omega_{c\pi}), \quad where\ \ \Omega_{c\pi} = diag(\omega_{c\pi}^2, \sigma_{c\pi}^2, ..., \sigma_{c\pi}^2) \tag{14}$$

where,

$$\mu_{c\pi} = \begin{pmatrix} C_0^{\pi} \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

Ignoring any terms not involving C^{π}, the prior density for C^{π} is given by

$$log\ \ p(C^{\pi}|\sigma^2_{c\pi}, \Omega_{c\pi}) \propto -\tfrac{1}{2}(C^{\pi} - H^{-1}\mu_{c\pi})' H' \Omega^{-1}_{c\pi} H (C^{\pi} - H^{-1}\mu_{c\pi})$$

Combining the above two conditional densities we obtain,

$$log\ \ p(C^{\pi}|Y, \bullet) \propto -\tfrac{1}{2}(C^{\pi} - \hat{C}^{\pi})' D^{-1}_{C^{\pi}} (C^{\pi} - \hat{C}^{\pi})$$

where,

$D_{C^{\pi}} = (H' \Omega^{-1}_{c\pi} H + \Omega^{-1}_{z\pi})^{-1}$

$\hat{C}^{\pi} = D_{C^{\pi}} (H' \Omega^{-1}_{c\pi} \mu_{c\pi} + \Omega^{-1}_{z\pi} N^{z\pi})$

The candidate draws are sampled from $N(\hat{C}^{\pi}, D_{C^{\pi}})$ using the precision-based algorithm.

Step 11. Derive the conditional distribution $p(\theta|Y, \bullet)$

There are 10 parameters in the vector θ. These parameters are drawn in 10 separate blocks using standard regression procedures.

Substep 11.1 Derive the conditional distribution $p(\beta^{\pi}|Y, \bullet)$

Candidate draws are sampled from $N(\hat{\beta}^{\pi}, D_{\beta\pi})$ using the precision-based algorithm.

where,

$D_{\beta\pi} = (V^{-1}_{\beta\pi} + X^{z\pi\prime} \Omega^{-1}_{z\pi} X^{z\pi})^{-1}$
$\hat{\beta}^{\pi} = D_{\beta\pi} (V^{-1}_{\beta\pi} \beta^{\pi}_0 + X^{z\pi\prime} \Omega^{-1}_{z\pi} J^{z\pi})$

$j_t^{z\pi} = Z_t^{\pi} - C_t^{\pi}$
$J^{z\pi} = (j_1^{z\pi}, ..., j_T^{z\pi})'$
$x_t^{z\pi} = w_t^S \tau_t^S + w_t^G \tau_t^G$
$X^{z\pi} = (x_1^{z\pi}, ..., x_T^{z\pi})'$

$V_{\beta\pi}^{-1}$ is the prior variance and β_0^{π} is the prior mean for β^{π}

Substep 11.2 Derive the conditional distribution $p(\sigma_{c\pi}^2|Y, \bullet)$

$p(\sigma_{c\pi}^2|Y, \bullet)$ is a standard inverse-Gamma density,

Candidate draws are sampled from

$$p(\sigma_{c\pi}^2|Y, \bullet) \sim IG(\nu_{c\pi 0} + \tfrac{T-1}{2}, S_{c\pi 0} + \tfrac{1}{2}\sum_{t=2}^{T}(C_t^{\pi} - C_{t-1}^{\pi})^2)$$

Substep 11.3 Derive the conditional distribution $p(\sigma_{z\pi}^2|Y, \bullet)$

$p(\sigma_{z\pi}^2|Y, \bullet)$ is a standard inverse-Gamma density,

Candidate draws are sampled from

$$p(\sigma_{z\pi}^2|Y, \bullet) \sim IG(\nu_{z\pi 0} + \tfrac{T}{2}, S_{z\pi 0} + \tfrac{1}{2}\sum_{t=1}^{T}(Z_t^{\pi} - C_t^{\pi} - \beta^{\pi}\tau^*)^2)$$

Substep 11.4 Derive the conditional distribution $p(\sigma_{\alpha s}^2|Y, \bullet)$

$p(\sigma_{\alpha s}^2|Y, \bullet)$ is a standard inverse-Gamma density,

Candidate draws are sampled from

$$p(\sigma_{\alpha s}^2|Y, \bullet) \sim IG(\nu_{\alpha s0} + \tfrac{T-1}{2}, S_{\alpha s0} + \tfrac{1}{2}\sum_{t=2}^{T}(\alpha_t^S - \alpha_{t-1}^S)^2)$$

Substep 11.5 Derive the conditional distribution $p(\sigma_{\alpha G}^2|Y, \bullet)$

$p(\sigma^2_{\alpha G}|Y, \bullet)$ is a standard inverse-Gamma density,

Candidate draws are sampled from

$$p(\sigma^2_{\alpha G}|Y, \bullet) \sim IG(\nu_{\alpha g0} + \tfrac{T-1}{2}, S_{\alpha g0} + \tfrac{1}{2}\sum_{t=2}^{T}(\alpha^G_t - \alpha^G_{t-1})^2)$$

Substep 11.6 Derive the conditional distribution $p(\sigma^2_{hs}|Y, \bullet)$

$p(\sigma^2_{hs}|Y, \bullet)$ is a standard inverse-Gamma density,

Candidate draws are sampled from

$$p(\sigma^2_{hs}|Y, \bullet) \sim IG(\nu_{hs0} + \tfrac{T-1}{2}, S_{hs0} + \tfrac{1}{2}\sum_{t=2}^{T}(h^{Ser}_t - h^{Ser}_{t-1})^2)$$

Substep 11.7 Derive the conditional distribution $p(\sigma^2_{hG}|Y, \bullet)$

$p(\sigma^2_{hG}|Y, \bullet)$ is a standard inverse-Gamma density,

Candidate draws are sampled from

$$p(\sigma^2_{hG}|Y, \bullet) \sim IG(\nu_{hg0} + \tfrac{T-1}{2}, S_{hg0} + \tfrac{1}{2}\sum_{t=2}^{T}(h^{Goods}_t - h^{Goods}_{t-1})^2)$$

Substep 11.8 Derive the conditional distribution $p(\sigma^2_{h\tau s}|Y, \bullet)$

$p(\sigma^2_{h\tau s}|Y, \bullet)$ is a standard inverse-Gamma density,

Candidate draws are sampled from

$$p(\sigma^2_{h\tau s}|Y, \bullet) \sim IG(\nu_{h\tau s0} + \tfrac{T-1}{2}, S_{h\tau s0} + \tfrac{1}{2}\sum_{t=2}^{T}(h^{\tau S}_t - h^{\tau S}_{t-1})^2)$$

Substep 11.9 Derive the conditional distribution $p(\sigma^2_{h\tau G}|Y, \bullet)$

$p(\sigma^2_{h\tau G}|Y, \bullet)$ is a standard inverse-Gamma density,

Candidate draws are sampled from

$$p(\sigma^2_{h\tau G}|Y, \bullet) \sim IG(\nu_{h\tau g0} + \tfrac{T-1}{2}, S_{h\tau g0} + \tfrac{1}{2}\sum_{t=2}^{T}(h_t^{\tau G} - h_{t-1}^{\tau G})^2)$$

Substep 11.10 Derive the conditional distribution $p(\sigma^2_\lambda|Y, \bullet)$

$p(\sigma^2_\lambda|Y, \bullet)$ is a non-standard density because of the constraints on λ,

$$log\ \ p(\sigma^2_\lambda|Y, \bullet) \propto -(\nu_{\lambda 0}+1)\ \ log\ \ \sigma^2_\lambda - \tfrac{S_{\lambda 0}}{\sigma^2_\lambda} - \tfrac{T-1}{2}\ \ log\ \ \sigma^2_\lambda - \tfrac{1}{2\sigma^2_\lambda}\textstyle\sum_{t=2}^{T}(\lambda_t - \lambda_{t-1})^2 + g_\lambda(\lambda, \sigma^2_\lambda)$$

The candidate draws from $p(\sigma^2_\lambda|Y, \bullet)$ are obtained via the MH step with the proposal density $IG(\nu_{\lambda 0} + \frac{T-1}{2}, S_{\lambda 0} + \frac{1}{2}\sum_{t=2}^{T}(\lambda_t - \lambda_{t-1})^2)$

A3. MCMC Convergence Diagnostics

This section documents the diagnostic properties of the Base model's MCMC algorithm. Following Koop, Leon-Gonzalez, and Strachan (2010), among others, we report the autocorrelation functions of the posterior draws (10th and 50th order sample autocorrelation), inefficiency factors (IFs), and convergence diagnostic (CD) of Geweke (1992) of the model parameters.[1]

One of the most common metrics examined to assess the efficiency of the MCMC sampler is to look at the autocorrelation function of the draws, which indicates how well the chain is mixing. Low autocorrelations are preferred to higher because the lower the autocorrelation, the closer the draws are to being independent and the higher the efficiency of the algorithm. The plots shown in the left panels of Figures 1,...,5 correspond to the 10th and 50th order autocorrelations in the draws, and as can be seen, they indicate very low autocorrelation. In the case of the 50th order autocorrelation, most indicate correlation less than 0.1, and in the case of 10th order except for a few, most indicate correlation below 0.3.

To assess whether the MCMC sampler has converged, a rough rule of thumb is to look at the CDs and see whether 95% of them are less than 2. If they are, then convergence is likely achieved. Based on the plots in the right panels of Figures 1,...,5, majority of CDs are within +/- 2.

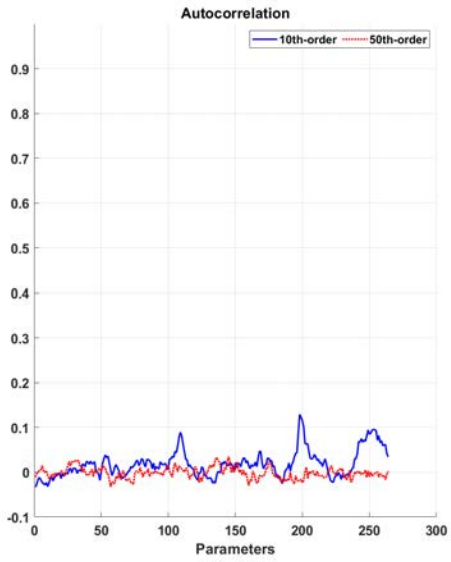
The inefficiency factor related to the autocorrelation functions is the inverse of Geweke's (1992) relative numerical efficiency measure (RNE). It is computed using the following formula, $(1 + 2\sum_{i=1}^{\infty} \rho_i)$, where ρ_i refers to the $k - th$ order autocorrelation of the chain. Figure 6 plots the IF for all the model parameters of the Baseline model. The values lower than or close to 20 are considered desirable. As shown, most of the IFs are below 20. (Note: IFs are computed using the default setting in LeSage's toolbox: estimation of spectral density at frequency zero uses a tapered window of 4%.)

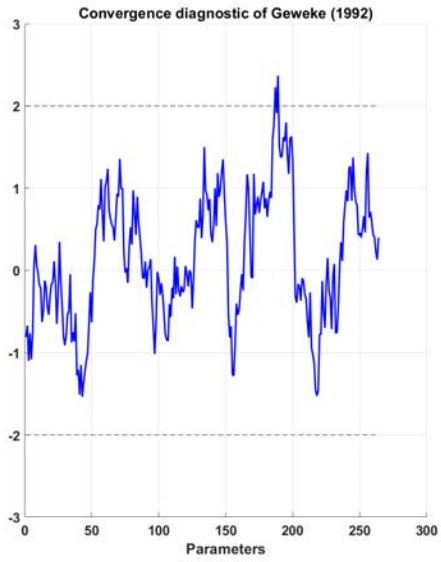
Overall, these diagnostic measures suggest good convergence properties of the MCMC algorithm developed for the Baseline model.

[1] In computing some of these metrics, we have benefitted from the Matlab toolbox developed by James P. LeSage. A detailed explanation, including intuition for these convergence diagnostics, is provided in Koop (2003; page 67-68) and Chan et al. (2019; page 209).

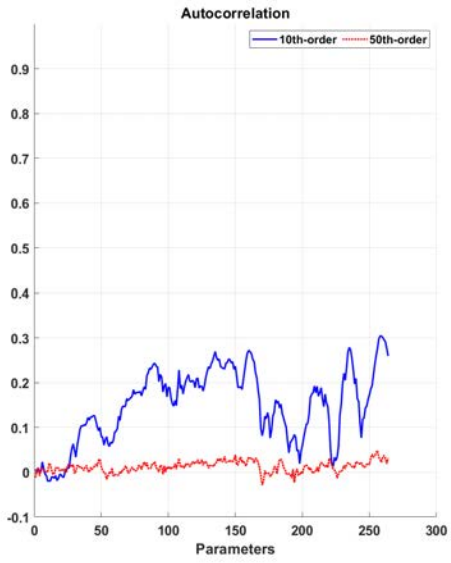
Figure 1: MCMC Diagnostics: Autocorrelation and CDs

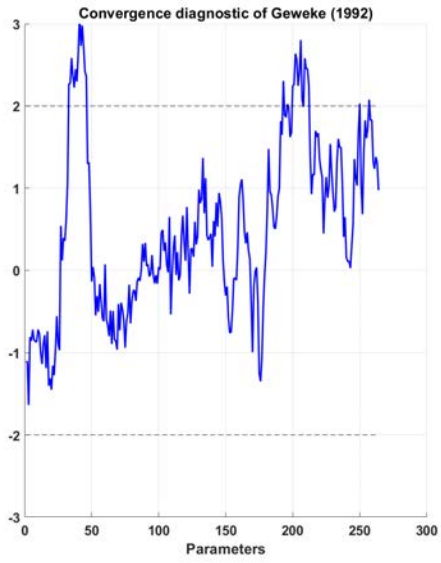
SV parameter: h^{Ser}





SV parameter: h^{Goods}

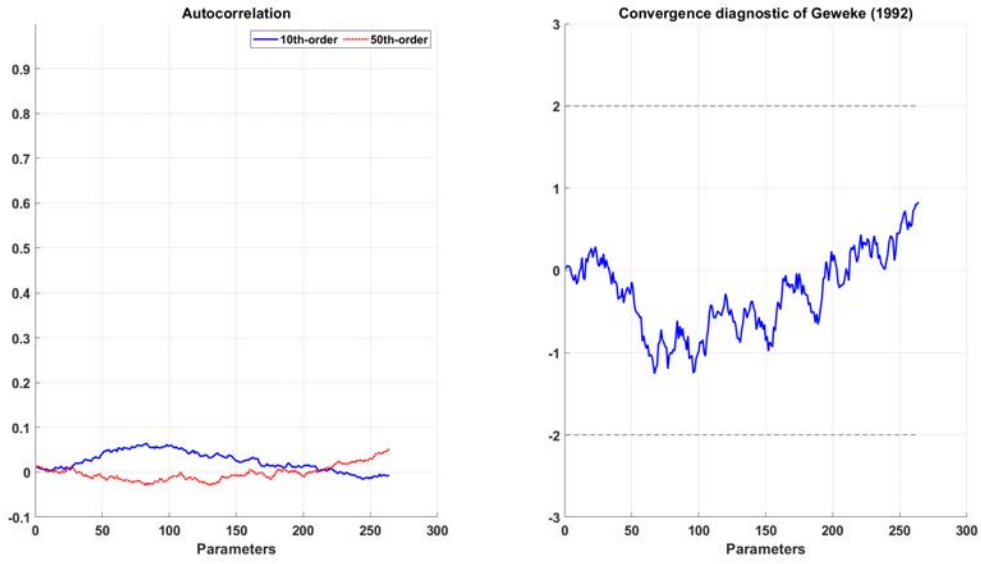




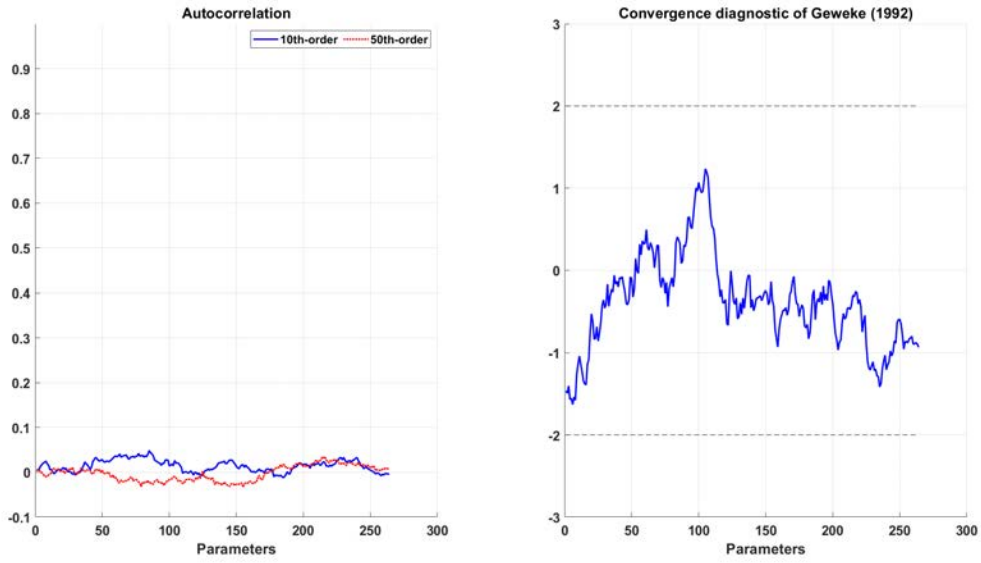
Note: Diagnostics based on estimating data from 1959Q3 through 2025Q2.

Figure 2: MCMC Diagnostics: Autocorrelation and CDs

SV parameter: $h^{\tau S}$



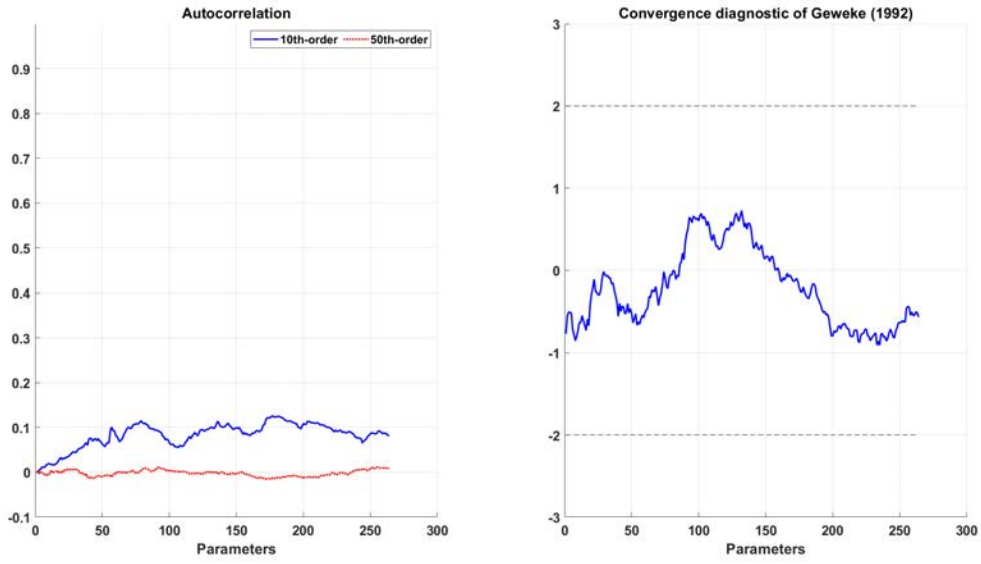
SV parameter: $h^{\tau G}$



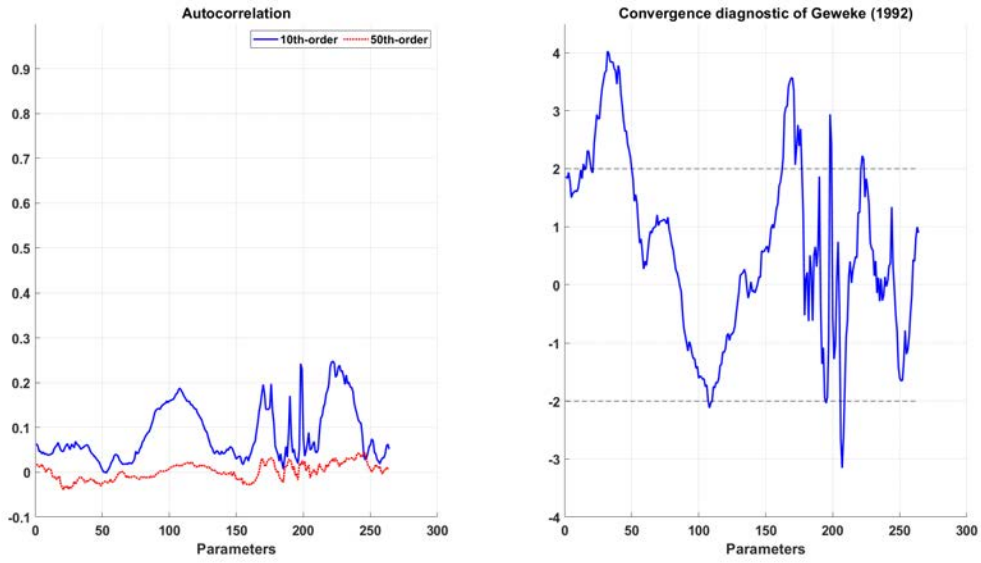
Note: Diagnostics based on estimating data from 1959Q3 through 2025Q2.

Figure 3: MCMC Diagnostics: Autocorrelation and CDs

TVP parameter: α^S



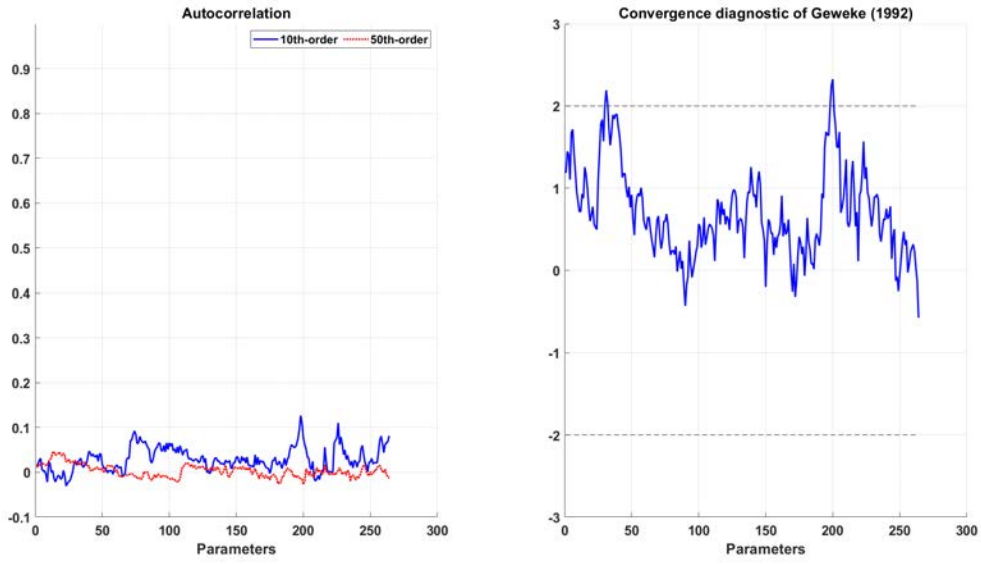
TVP parameter: α^G



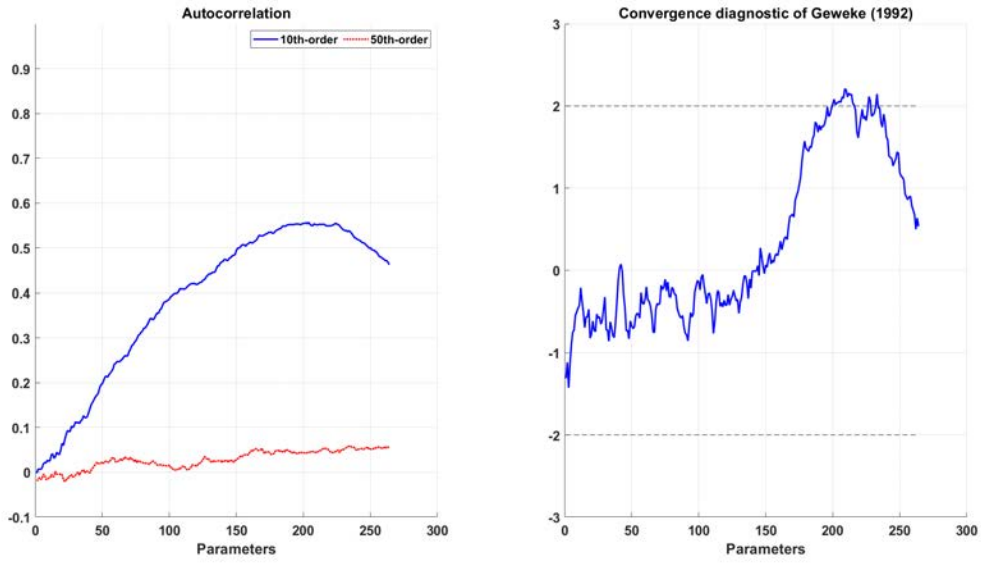
Note: Diagnostics based on estimating data from 1959Q3 through 2025Q2.

Figure 4: MCMC Diagnostics: Autocorrelation and CDs

TVP parameter: λ



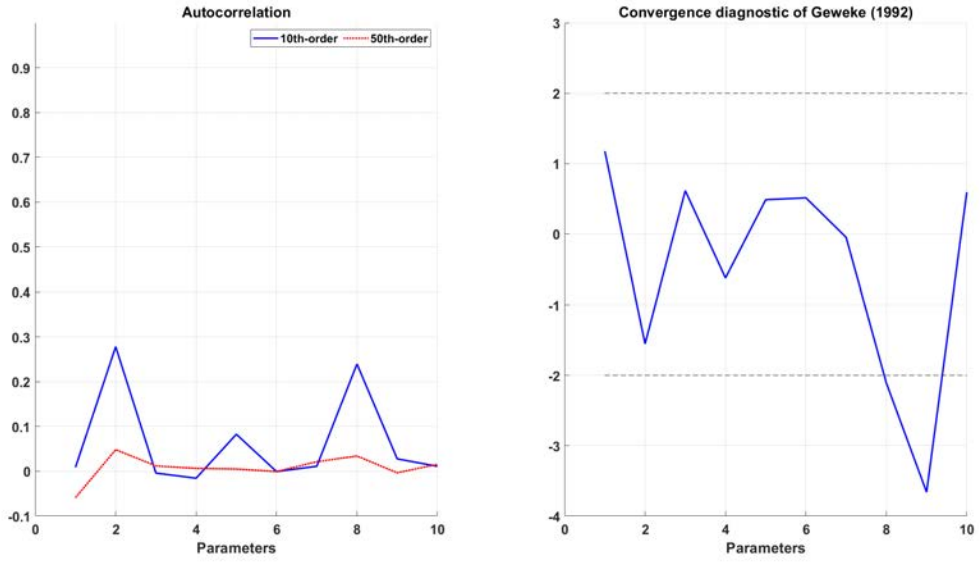
TVP parameter: C^{π}



Note: Diagnostics based on estimating data from 1959Q3 through 2025Q2.

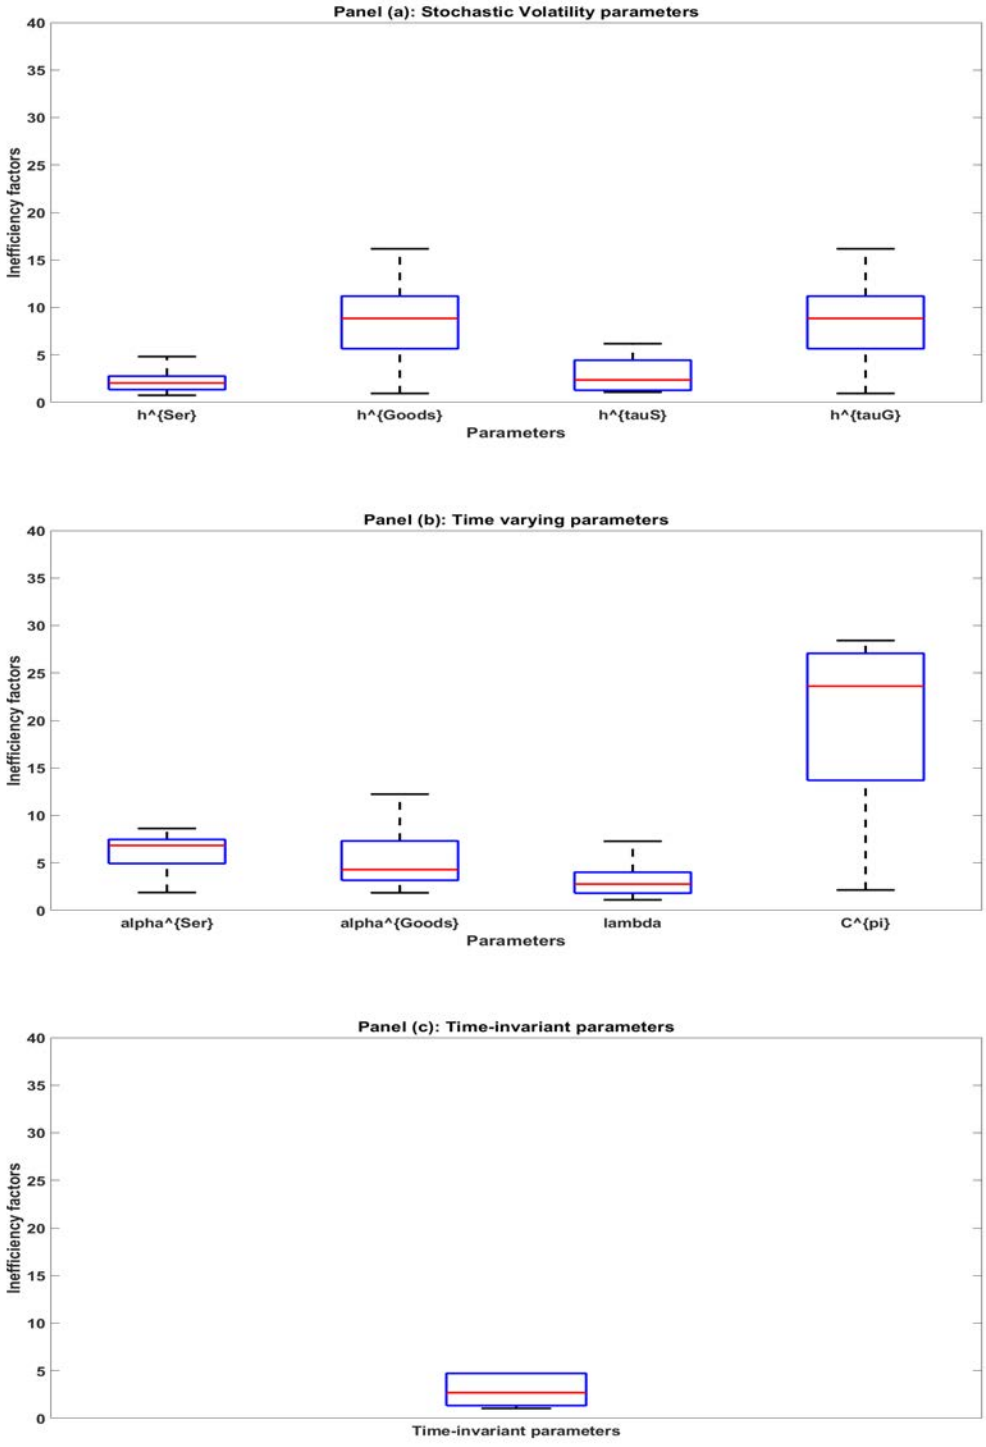
Figure 5: MCMC Diagnostics: Autocorrelation and CDs

Time-invariant parameters



Note: Diagnostics based on estimating data from 1959Q3 through 2025Q2.

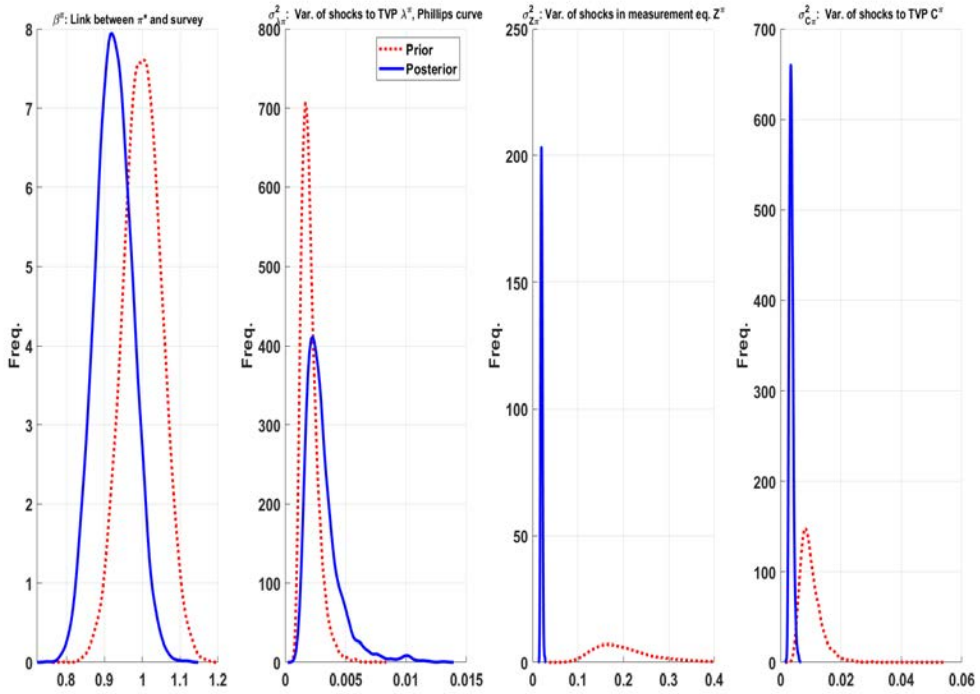
Figure 6: MCMC Diagnostics: Inefficiency Factors



Note: Diagnostics based on estimating data from 1959Q3 through 2025Q2. The red line represents the median, the lower and upper lines represent the 25th and 75th percentiles, respectively. The whiskers extend to the maximum and minimum.

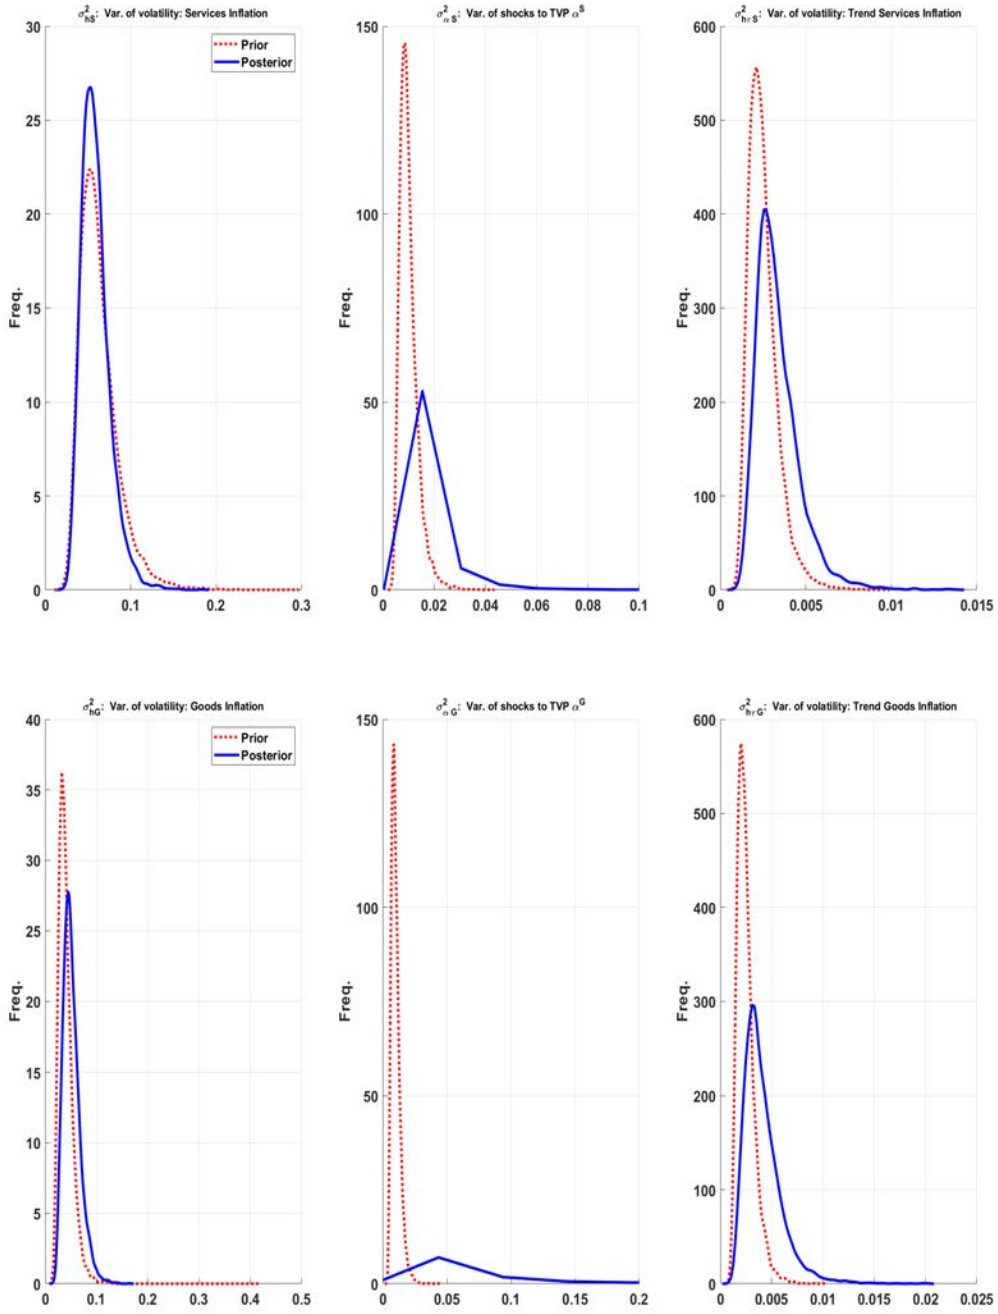
A4. Prior and Posterior Distributions of the Parameters: Baseline Model

Figure 7: Prior and Posterior Distributions of the Parameters: Baseline Model



Note: Plotted in dotted red lines are the prior distributions, in blue are the posterior distributions based on estimating data from 1959Q3 through 2025Q2.

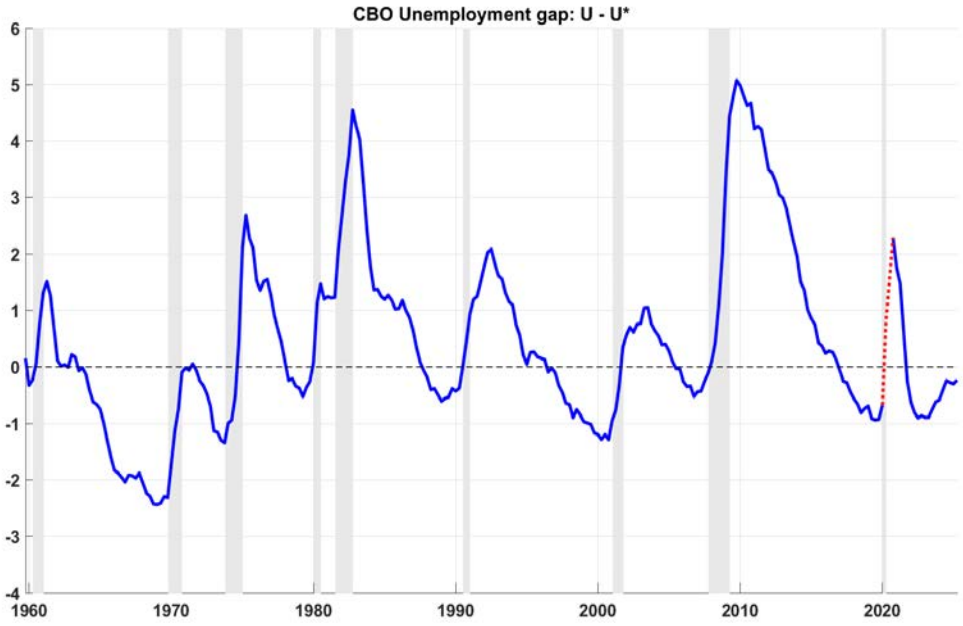
Figure 8: Prior and Posterior Distributions of the Parameters: Baseline Model



Note: Plotted in dotted red lines are the prior distributions, in blue are the posterior distributions based on estimating data from 1959Q3 through 2025Q2.

A5. CBO Unemployment Rate Gap

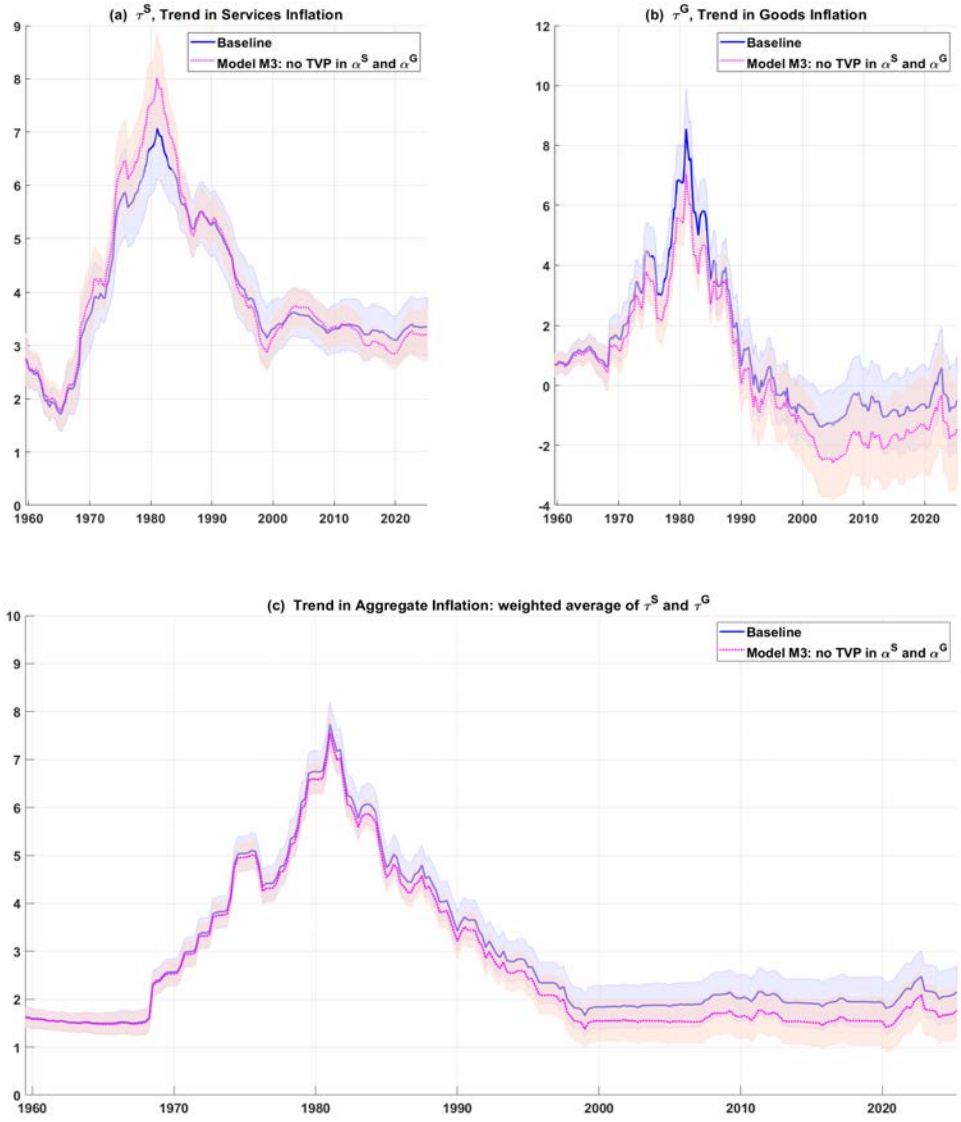
Figure 9: Congressional Budget Office estimate of the Unemployment Gap



Note: The plot in blue represents the Congressional Budget Office (CBO) estimate of the Unemployment gap, defined as the difference between the unemployment rate and the CBO estimate of the natural rate of unemployment. The plot in red represents the inferred value of the unemployment gap for 2020Q2 and 2020Q3, the COVID-pandemic quarters; the inferred value for 2020Q2 is based on averaging the actual unemployment gap values for 2020Q1 and 2020Q4, and for 2020Q3 is based on averaging the inferred 2020Q2 value and the actual 2020Q4 value. The shaded bars represent the NBER-defined recession bars.

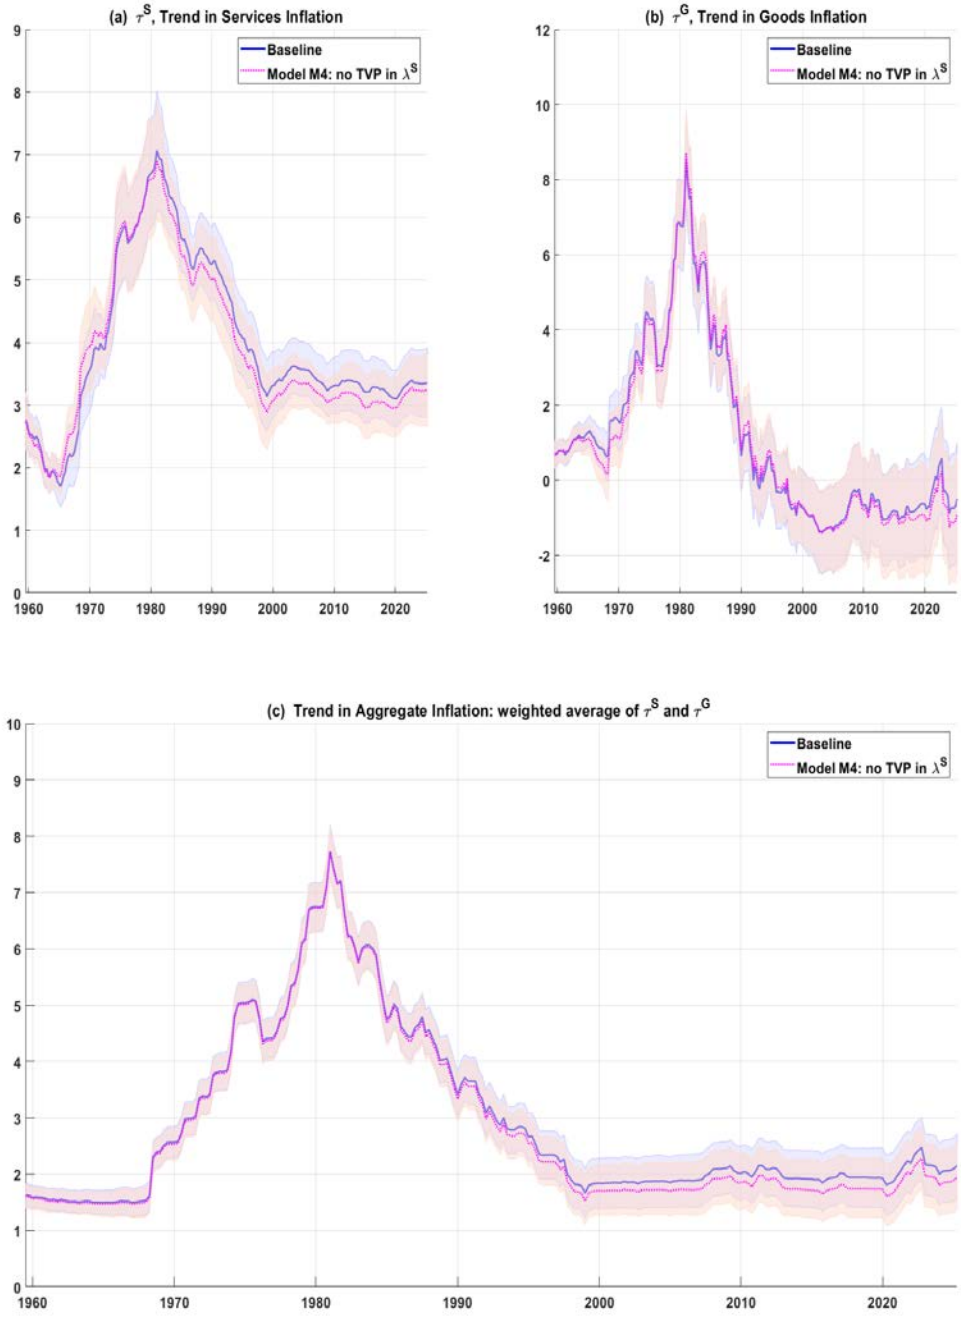
A6. Trend Estimates: Sensitivity to Modeling Assumptions

Figure 10: Trend Inflation: Baseline vs. Variant with time-invariant Uncertainty effects



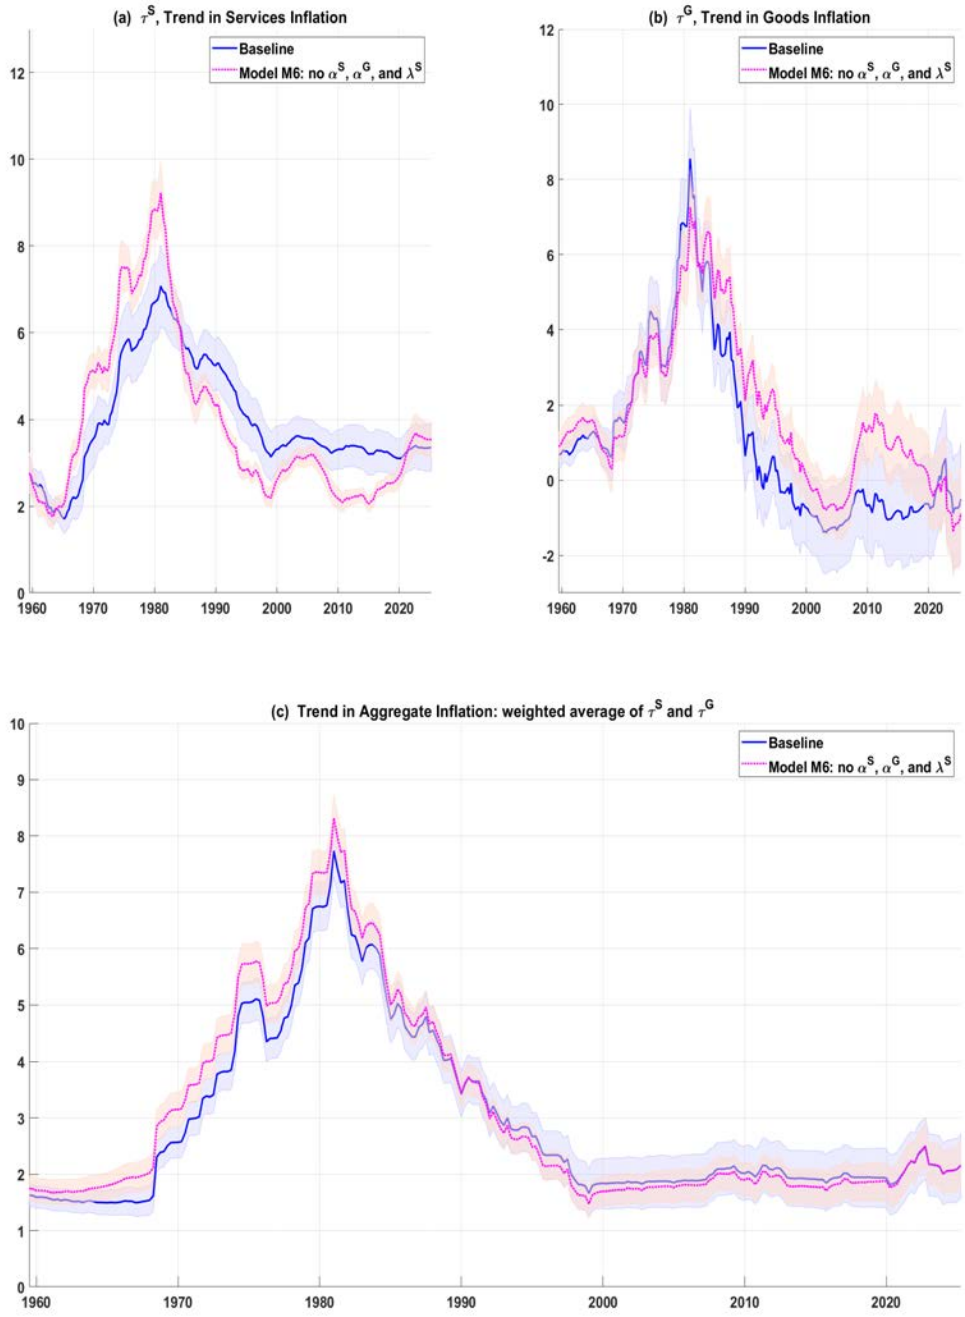
Note: The posterior estimates are based on the full sample (from 1959Q3 through 2025Q2). The shaded area represents the 68% posterior coverage interval. Panel (a), posterior estimates of the trend in services PCE inflation. Panel (b), posterior estimates of the trend in goods PCE inflation. Panel (c),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

Figure 11: Trend Inflation: Baseline vs. Variant with time-invariant Phillips curve



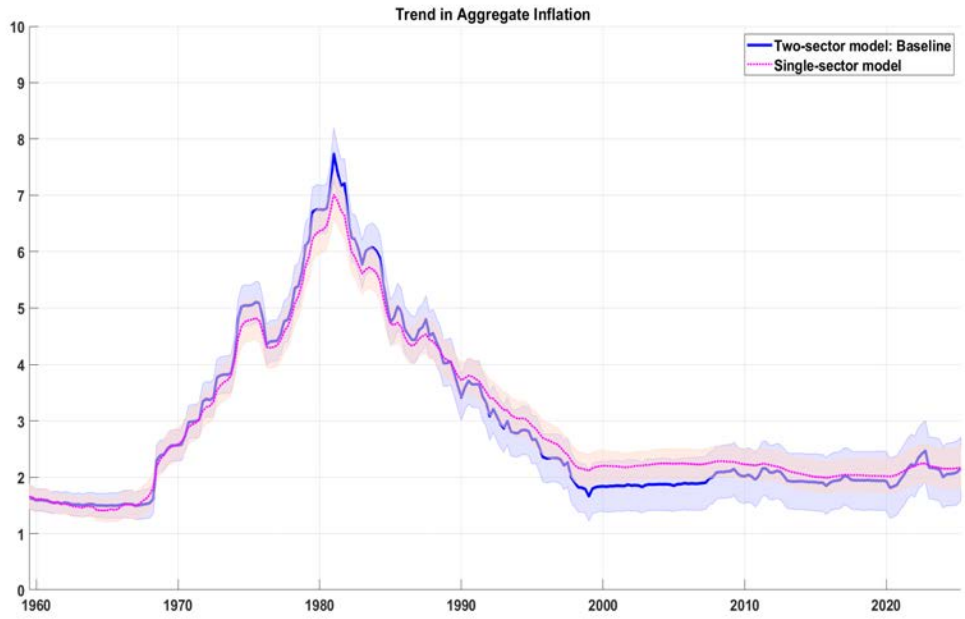
Note: The posterior estimates are based on the full sample (from 1959Q3 through 2025Q2). The shaded area represents the 68% posterior coverage interval. Panel (a), posterior estimates of the trend in services PCE inflation. Panel (b), posterior estimates of the trend in goods PCE inflation. Panel (c),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

Figure 12: Trend Inflation: Baseline vs. Variant with No Uncertainty and Phillips Curve



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.

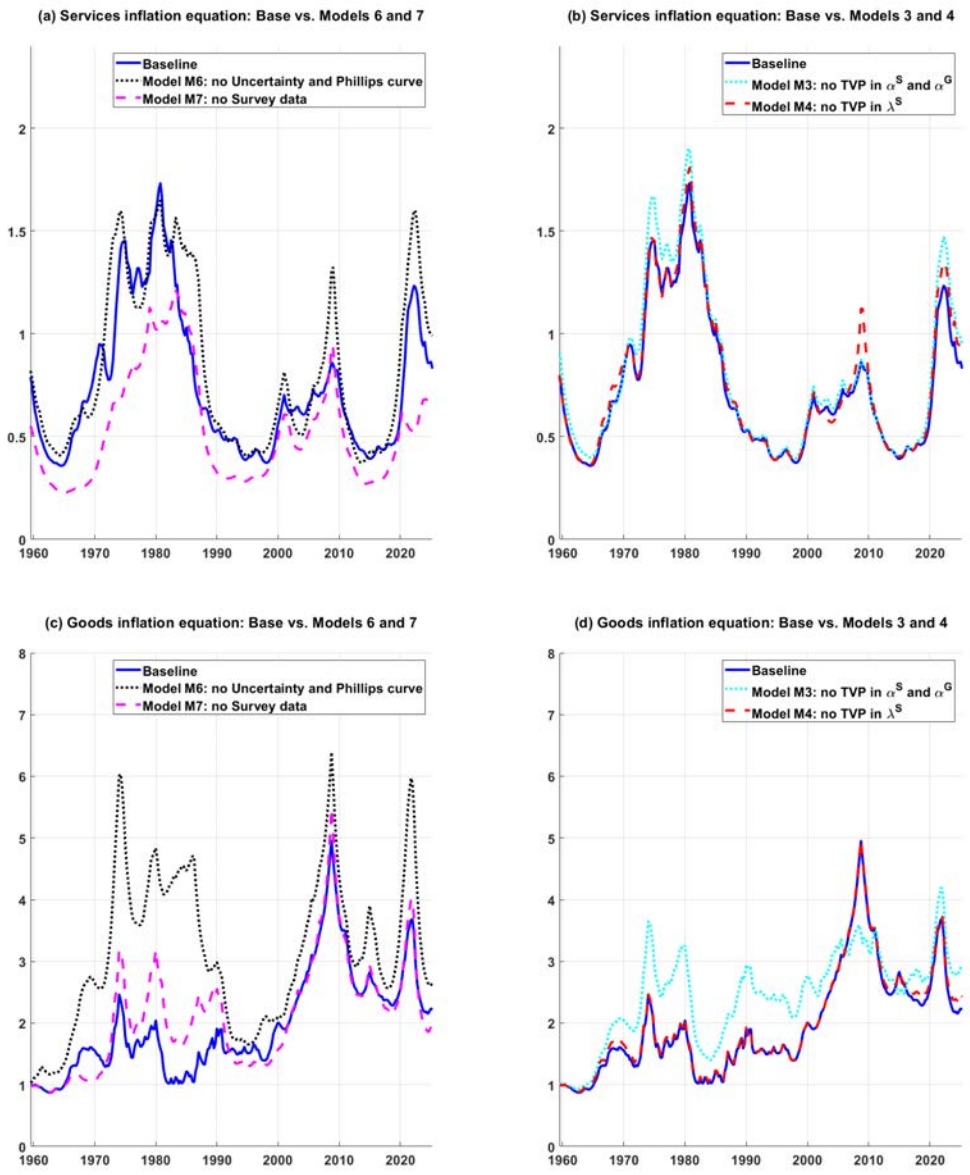
Figure 13: Trend Inflation: Baseline vs. Single-Sector Model



Note: The posterior estimates are based on the full sample (from 1959Q3 through 2025Q2). The shaded area represents the 68% posterior coverage interval. The solid and dashed lines represent the posterior mean estimates from the (two-sector) baseline model and single-sector models, respectively (posterior estimates of the trend in services PCE inflation.

A7. Uncertainty Estimates: Sensitivity to Modeling Assumptions

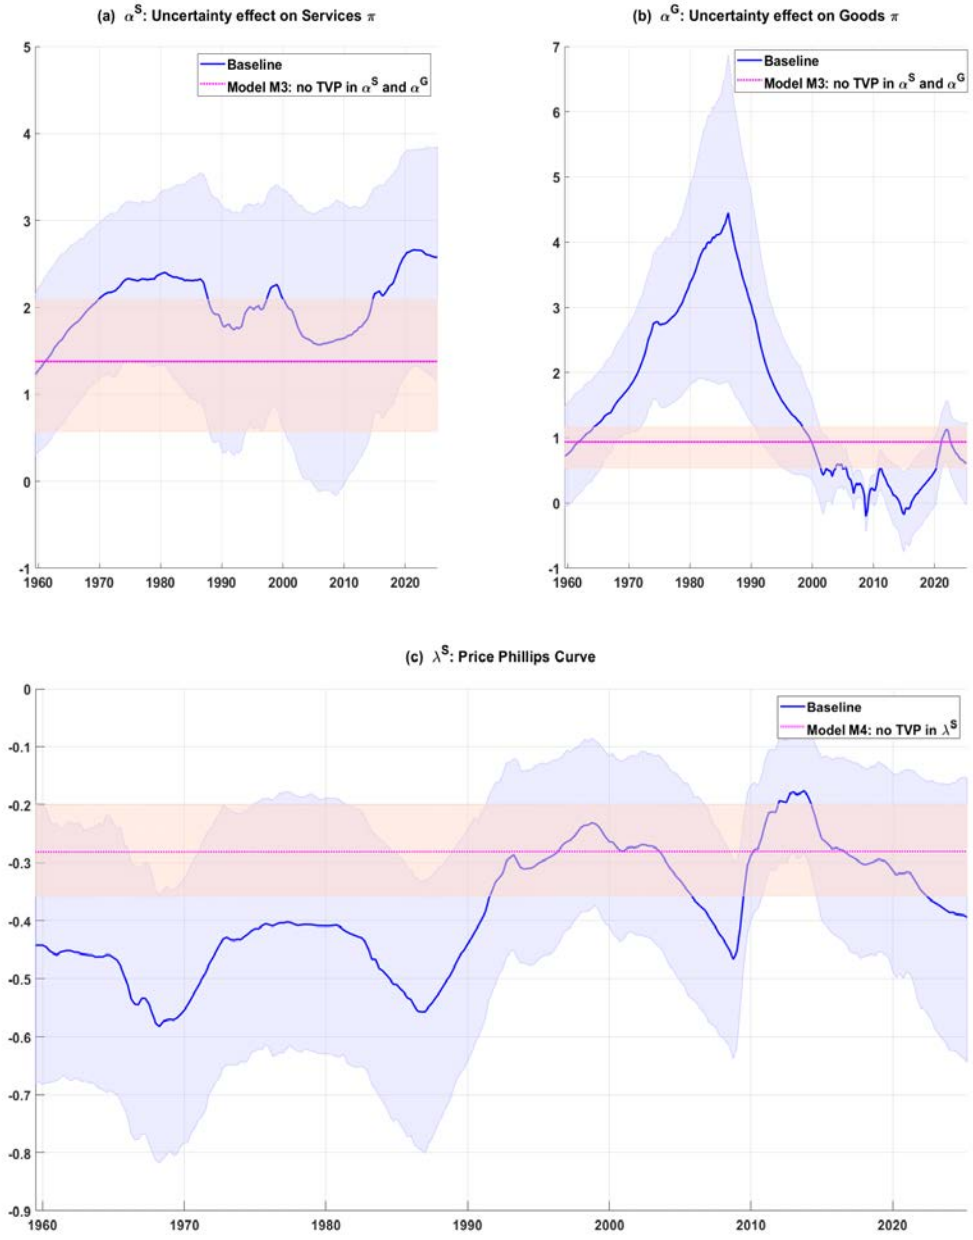
Figure 14: Estimates of Stochastic Volatility



Note: Panel (a) and Panel (b) present the uncertainty estimates associated with services inflation, i.e., the posterior median estimates of the standard deviation (sd) of the shocks defining the services inflation measurement equation. Panel (c) and Panel (d) present the uncertainty estimates associated with goods inflation. The estimates are computed using the full sample (from 1959Q3 through 2025Q2).

A8. Comparing Parameter Estimates: Baseline vs. Model with time-invariant parameters

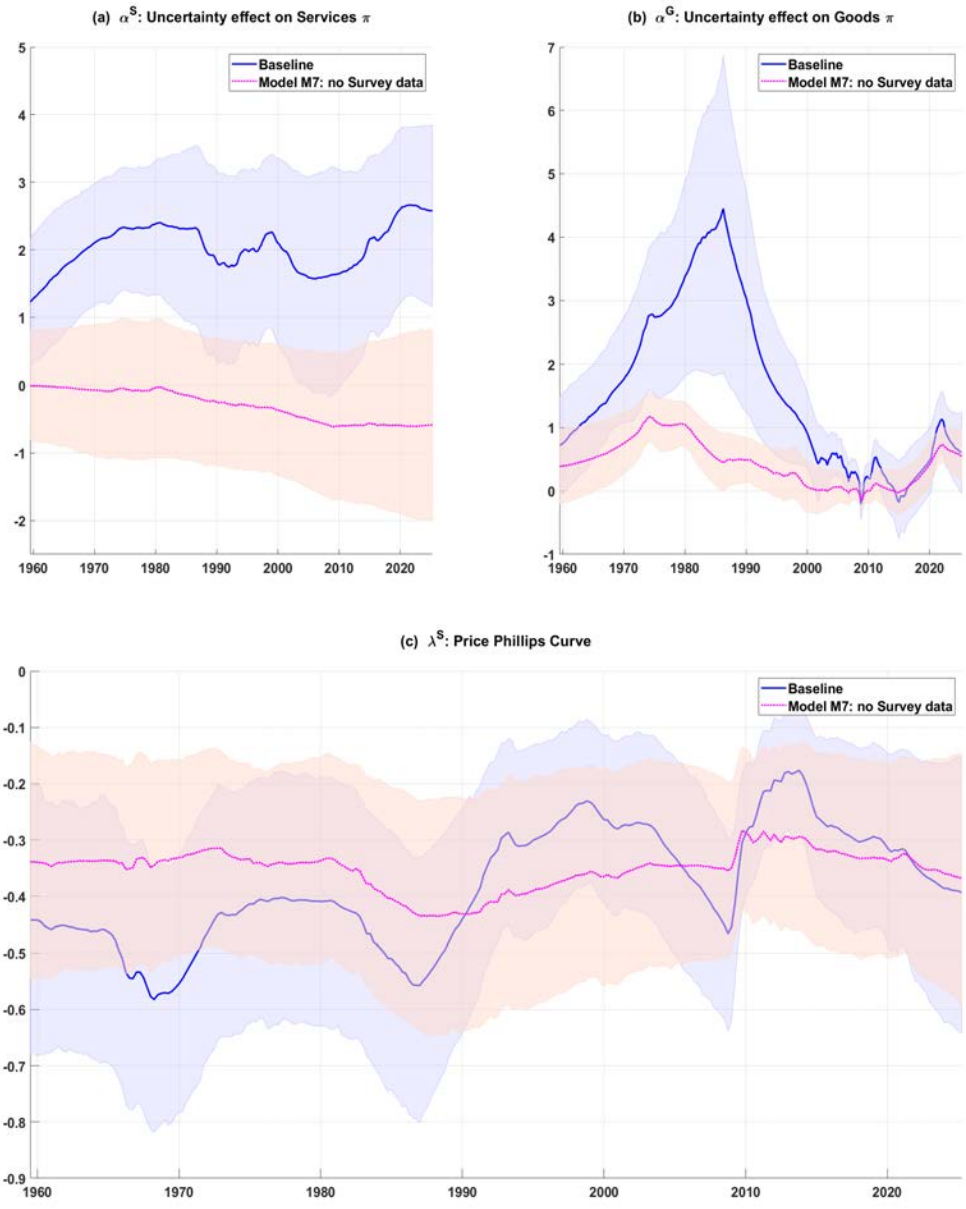
Figure 15: Estimates of TVPs: Base vs. Variant with time-invariant parameters



Note: Plotted are the posterior estimates of the 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.

A9. Comparing Parameter Estimates: Baseline vs. Model without Survey Expectations Data

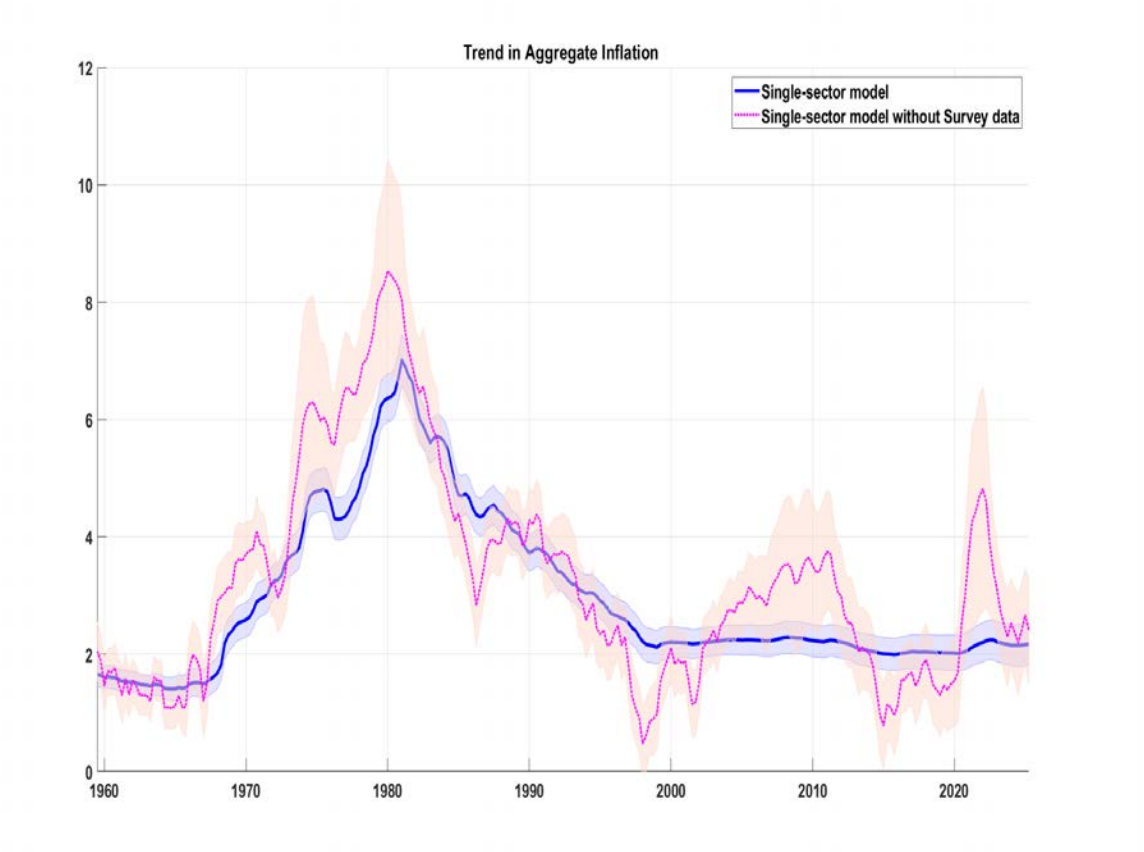
Figure 16: Estimates of TVPs: Base vs. Variant with no Survey data



Note: Plotted are the posterior estimates of the 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.

A10. Trend Estimates: Single-Sector Model vs. Single-Sector Model without Survey data

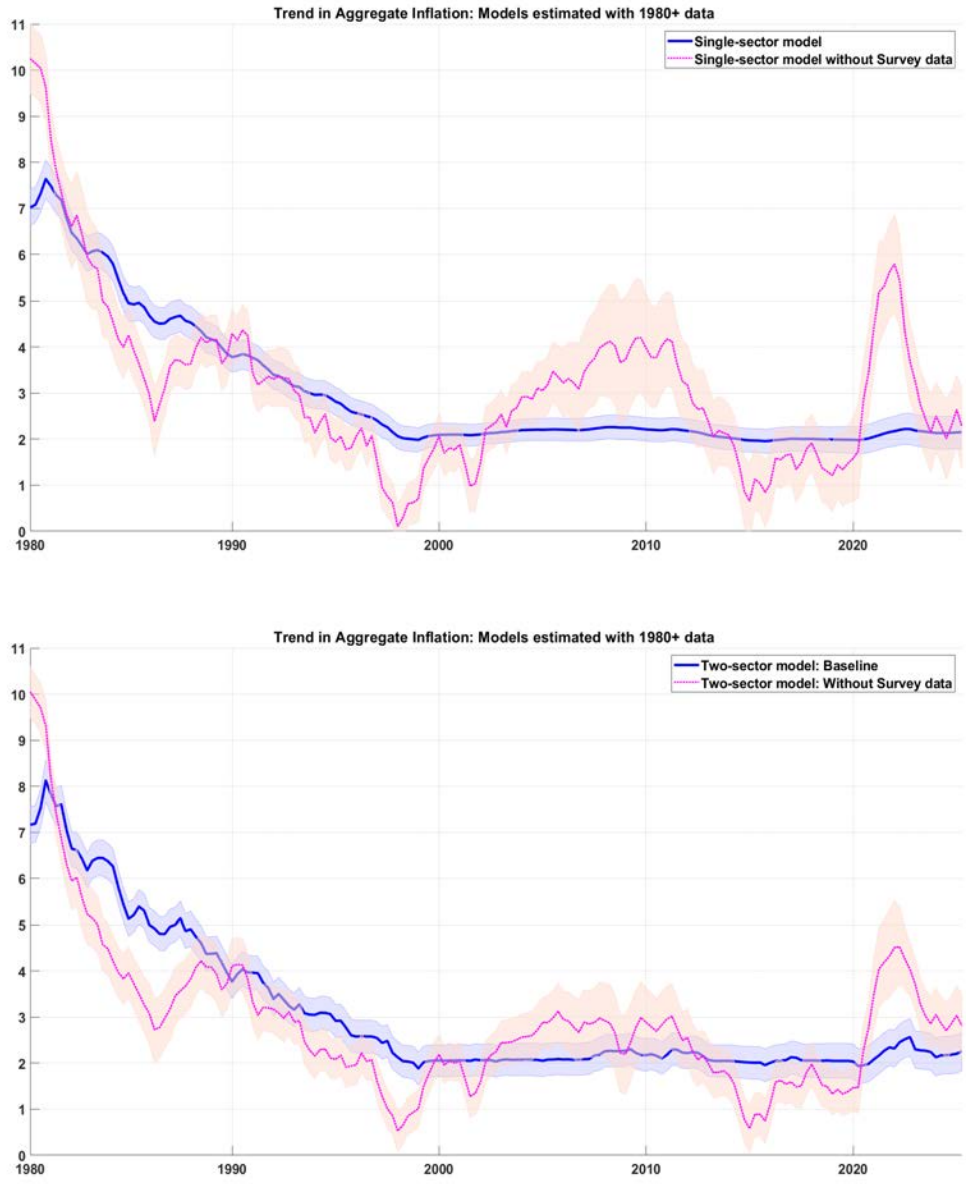
Figure 17: Trend Inflation: Single-Sector Model



Note: The posterior estimates are based on the full sample (from 1959Q3 through 2025Q2). The shaded area represents the 68% posterior coverage interval.

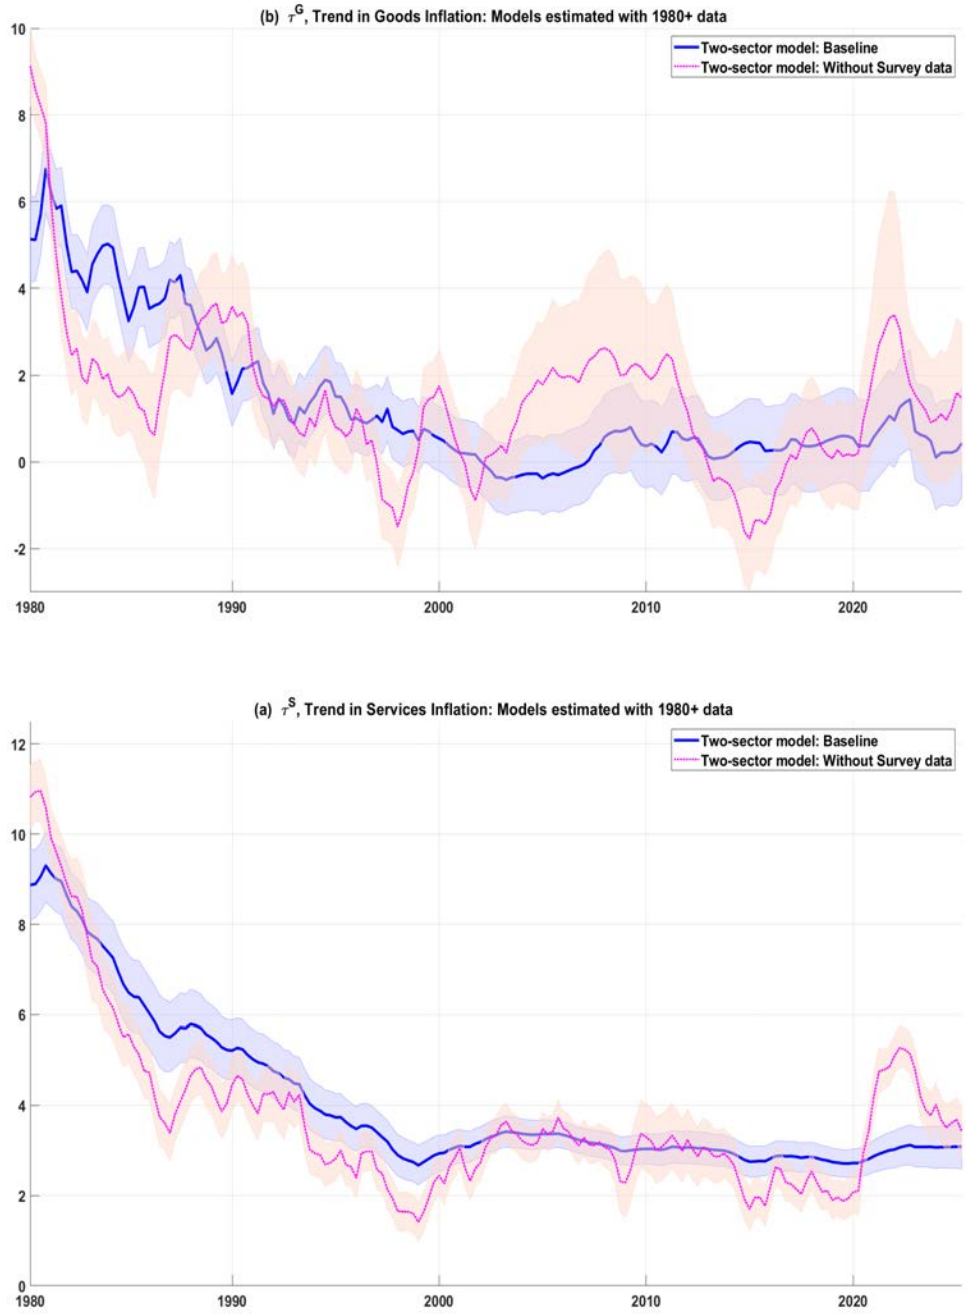
A11. Trend estimates based on estimating models with data starting from 1980 onward

Figure 18: Trend Estimates: Baseline and Single-Sector Models



Note: Plotted are the posterior estimates of the parameters computed using the full sample (from 1980Q2 through 2025Q2). The solid lines represent the posterior mean, and the shaded area represents the 68% credible intervals.

Figure 19: Goods and Services Trend Estimates: Base vs. Variant with no Survey data



Note: Plotted are the posterior estimates of the parameters computed using the full sample (from 1980Q2 through 2025Q2). The solid lines represent the posterior mean, and the shaded area represents the 68% credible intervals.

Table 3: Bayesian Model Comparison: Models estimated with data 1980+

*Panel A: Baseline (Two-sector model) vs. Variant without Survey data

	1980-2025Q2, MDD		1980-2019, MDD	
Model	Services π	Goods π	Services π	Goods π
M1: Baseline	-204.3	-437.8	-168.5	-379.8
M7: Baseline but No Survey data	-192.0	-418.0	-163.7	-362.6

*Panel B: Baseline, Single-sector, and Variants without Survey data

	MDD Aggregate PCE π	
Model	1980-2025Q2 sample	1980-2019 sample
M1: Baseline	-287.1	-246.1
M7: Baseline but No Survey data	-267.9	-230.3
M8: Single-sector	-290.2	-246.1
M10: Single-sector but No Survey data	-273.2	-236.0

Note: MDD refers to the log of the marginal data density, which measures the model's fit to the services inflation data and to goods inflation data. The top panel reports the model's fit to the data over the entire sample (1980Q4-2025Q2) and the pre-COVID sample (1980Q4-2019Q4). The bottom panel reports the model fit to the aggregate PCE inflation for the single-sector model (which models directly the aggregate PCE inflation) and its variants, to facilitate comparison, the implied model fit for the Baseline model and Baseline model without survey data are also reported.