

Forecasting Core Inflation and Its Goods, Housing, and Supercore Components*

Todd E. Clark,^{a,b} Matthew V. Gordon,^c and Saeed Zaman^b

^aJohns Hopkins University

^bFederal Reserve Bank of Cleveland

^cUniversity of Pennsylvania

This paper examines the forecasting efficacy and implications of the recently popular breakdown of core inflation into three components: goods excluding food and energy, services excluding energy and housing, and housing. A comprehensive historical evaluation of the accuracy of point and density forecasts from a range of models and approaches shows that a BVAR with stochastic volatility in aggregate core inflation, its three components, and wage growth is an effective tool for forecasting inflation's components as well as aggregate core inflation. Looking ahead, the model's baseline projection puts core inflation at 2.7 percent in 2026, well below its 2023 level but still elevated relative to the Federal Reserve's 2 percent objective. The probability that core inflation will return to 2 percent or less is much higher when conditioning on goods or nonhousing services inflation slowing to prepandemic levels than when conditioning on these components remaining above the same thresholds. Scenario analysis indicates that slower wage growth will likely be associated with reduced inflation in all three components, especially goods and nonhousing

*We gratefully acknowledge helpful comments from the editor Òscar Jordà, two anonymous referees, Elena Bobeica, and James Mitchell. We also thank Joshua Chan, Niko Hauzenberger (who also kindly provided guidance on usage), and Elmar Mertens for their publicly available replication code, which we use within our paper. Finally, we acknowledge the use of computing resources provided by the Center for the Advancement of Data and Research in Economics (CADRE) at the Federal Reserve Bank of Kansas City. Code and data for the results in this paper are available on our websites. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Cleveland or the Federal Reserve System. Author emails: tclarkresearch@gmail.com (corresponding author), mvgordon@sas.upenn.edu, saeed.zaman@clev.frb.org.

services, helping to return core inflation to near the 2 percent target by 2026.

JEL Codes: C32, C53, E17, E31, E37.

1. Introduction

Since inflation surged following the outbreak of the COVID-19 pandemic and Russia's war against Ukraine, central banks around the world have focused on the challenge of reducing inflation. In these efforts, forecasts of inflation have played an important role in policymakers' assessments. Recently, in commenting on inflation developments, policymakers in the United States have often drawn attention to a three-component breakdown of core inflation into goods excluding food and energy, services excluding energy and housing (also known as supercore inflation), and housing. As summarized in sources such as Powell (2022), a common narrative has maintained that the normalization of pandemic-driven changes in supply and demand would contribute to significant declines in the goods and housing components of core inflation, with some progress already evident. That same narrative has indicated that a slowing of inflation in nonhousing services, which has a weight of more than 50 percent in core inflation, would be crucial to returning inflation to the Federal Reserve's 2 percent target and perhaps more challenging to achieve than reductions in the other components. Because labor is even more important as an input to services than goods, a common suggestion has been that a slowing of aggregate wage growth to prepandemic levels will be needed to achieve a reduction in core services inflation. Federal Reserve Chair Powell has suggested that a decline of inflation to 2 percent will likely be associated with reductions in its core goods, supercore, and housing components.¹

These policy-driven interests raise a number of questions for research. First, how well can the goods, housing, and services components of core inflation be forecast as compared with aggregate core inflation? Second, is aggregate inflation best forecast directly or by aggregating forecasts of the three components? Third, how

¹See pages 5–6 of the transcript of the Chair's March 20, 2024 press conference: <https://www.federalreserve.gov/mediacenter/files/FOMCpresconf20240320.pdf>.

important is wage growth to slower inflation in services excluding housing as well as the other components of core inflation? Finally, based on historical comovement, how are the components likely to evolve as aggregate inflation slows? Is it most likely the case that all three components will need to return to prepandemic levels in order for aggregate core inflation to reach 2 percent?

Of course, to varying degrees, previous research has shed some light on these and related questions, although not necessarily directly in the current context. In particular, the forecasting literature has established the following. First, simple univariate models for forecasting aggregate core inflation are very hard to beat compared with multivariate models that include elements of the Phillips curve or other variables such as wage growth (see, e.g., Faust and Wright 2013 and Stock and Watson 2007). Second, most studies (but not all; some have obtained positive results) have found that it is generally difficult to consistently improve accuracy by combining predictions of components rather than forecasting aggregate core inflation directly (see, e.g., Eo, Uzeda, and Wong 2023; Faust and Wright 2013; Hubrich 2005; Joseph et al. 2022; McGregor and Toscani 2022; Tallman and Zaman 2017; and Verbrugge and Zaman 2024).² Third, allowing some time variation in means and other regression parameters of a model can help forecast accuracy (see, e.g., Banbura and van Vlodrop 2018 and D'Agostino, Gambetti, and Giannone 2013), although it is not assured of doing so (e.g., Chan 2023). Finally, allowing time variation in conditional variances can improve the accuracy of both point and density forecasts (see, e.g., Chan and Eisenstat 2018; Clark 2011; Clark and Ravazzolo 2015; and D'Agostino, Gambetti, and Giannone 2013).

Building on existing work, in this paper we examine the efficacy and implications of forecasts of the three-component breakdown of core inflation in the U.S. Of course, the accuracy of inflation forecasts has been widely studied; our examination is novel in its focus on the three-component breakdown of core inflation. We begin by

²While most studies noted in the text have focused on point forecasts, some work has examined whether density accuracy can be improved by aggregating component forecasts. For example, Ravazzolo and Vahey (2014) find some benefits to using linear opinion pooling to combine the forecast densities of inflation components.

evaluating the accuracy of inflation and component rate forecasts from a range of models and approaches that include some of the key features noted above, particularly stochastic volatility. While our analysis has considered additional models with larger variable sets and nonlinearities (without finding much payoff), we primarily focus on forecasts from Bayesian vector autoregressions (BVARs) that include as variables aggregate core inflation, its goods-housing-services excluding housing components, and wage growth. Our specifications vary in the features of the model—for example, we consider some models with time-varying parameters—as well as in the formulation of the coefficient priors. We compare the forecast performance of this range of models with the univariate unobserved components-stochastic volatility model that Stock and Watson (2007) established as a good model for forecasting aggregate inflation.

After establishing the historical efficacy of our preferred BVAR model specification, we examine medium-term forecasts for the period ahead, 2024 through 2026, and the likelihood (according to our model) of inflation returning to the Federal Reserve's 2 percent target. We conclude the paper with scenario analysis, that is, forecasts of inflation and its components conditional on selected scenarios or conditions. We implement conditional forecasting with entropic tilting (introduced to macroeconomic forecasting by Robertson, Tallman, and Whiteman 2005), which can be seen as a nonparametric approach to conditional forecasting. One scenario conditions on a gradual decline of wage growth to near its prepandemic level to assess the likelihood that such a development may be key to continued disinflation. The second scenario conditions on a decline of aggregate core inflation to 2 percent to provide a joint assessment of the paths of component inflation rates and wage growth most likely associated with a return of aggregate inflation to target.

Overall, we find that a BVAR with stochastic volatility estimated with an adaptive Minnesota-style prior (as developed in Chan 2021) is an effective tool for forecasting inflation in the core services excluding housing, housing, and core goods components of core inflation, as well as aggregate core inflation. Not surprisingly, compared with the range of models we consider, this BVAR model performs better for some components than others. For some inflation measures, its advantage over univariate benchmarks in forecast accuracy has

increased over the past several years, a period in which the pandemic caused variability in inflation. Regarding the specific questions noted above, our findings indicate that the inflation components can be forecast about as well as aggregate core inflation; in fact, our preferred BVAR tends to perform better, compared with a univariate benchmark, for inflation components than for aggregate inflation. However, aggregate inflation is best forecast directly; no consistent benefit can be achieved by aggregating disaggregate components. In our baseline predictive distributions, a return of core inflation to 2 percent or less over the next few years is much higher when conditioning on goods or supercore inflation slowing to near prepandemic levels than when conditioning on these components remaining above the same thresholds. Our scenario analyses indicate that slower wage growth will likely be associated with reduced inflation in all three components, not just supercore, but housing less so than goods and supercore inflation. In addition, it is most likely the case that a return of core inflation to the Federal Reserve's 2 percent target will be associated with all three components returning to near prepandemic levels.

The paper proceeds as follows. Section 2 describes the data used. Section 3 reviews the models used and details our approaches to estimation and forecasting. Section 4 presents our historical forecast evaluation, and Section 5 provides forecasts and scenario analysis for the next few years. Section 6 concludes.

2. Data

Because the Federal Reserve's inflation target is stated in terms of the price index for personal consumption expenditures (PCE), we focus on forecasting PCE inflation excluding food and energy (core PCE inflation) and its goods, housing, and services excluding housing components.³ Our primary results use quarterly data on PCE price indices—aggregate core inflation and its three components—as

³The PCE time series feature a methodologically consistent treatment of housing over time, based on owner-equivalent rent. In contrast, in published data on the consumer price index (CPI), methodological improvements are applied going forward but not to past data. Consequently, published CPI series suffer a significant break in the treatment of housing in 1983, when measurement changed to use the current approach based on owner-equivalent rent.

well as the employment cost index (ECI).⁴ We obtained all of these raw data series from Haver Analytics.⁵ (Some additional results in the online appendix mentioned in footnote 5 add to the models 19 other macroeconomic variables taken from the FRED-MD database maintained by the Federal Reserve Bank of St. Louis and from Haver Analytics.) We abstract from real-time data considerations and focus on currently available time series due to binding (from a forecast evaluation perspective) constraints on the availability of vintages of the components of PCE inflation.⁶

Inflation and wage growth rates are measured as 400 times log changes in index levels. The raw data sample (before transformations) for our forecast evaluation exercise is 1959:Q1 through 2023:Q4; Section 5's forecasts for 2024–26 use data through 2024:Q1. In the case of the ECI, the raw data do not begin until 1980. We extend the series back to 1959 by backcasting the growth rate of the ECI with fitted values from a regression of ECI growth on the growth rate of the series on compensation per hour (nonfarm business sector) from the Bureau of Labor Statistics. For our forecast evaluations we use data available as of March 28, 2024 while for our conditional forecasts we use data available as of April 30, 2024.

3. Models, Estimation, and Forecasting

This section details the models used for inflation forecasting, their estimation, and our approach to forecasting. For convenience, Table 1 lists the ensemble of models that we consider in our forecasting exercises.

⁴The ECI abstracts from cyclical changes in the sectoral composition of employment, avoiding some of the composition bias that affects other wage measures.

⁵An online appendix (available at https://github.com/zamansaeed/ijcb_cgzi_inflation/blob/main/CGZinflation_onlineappendix.pdf) provides the Haver mnemonics. Our results in Section 4 (Section 5) use data downloaded from Haver on March 28, 2024 (April 30, 2024).

⁶For example, in the Federal Reserve Bank of St. Louis' ALFRED database, a primary source of real-time data, vintages of PCE price index components only begin in 2013, and these vintages do not appear to have enough detail to permit us to construct real-time data on inflation in core goods, housing, and services excluding housing.

Table 1. List of Models

Identifier	Model Type	Prior	Time-Varying Parameters	Outlier Adjustment
UC-SV	UC-SV	N/A	N/A	None
NM-BVAR-SV	BVAR-SV	Minnesota	None	None
NM-BVAR-SVO	BVAR-SV	Minnesota	None	CCMM2024
NM-BVAR-SV-TVM	BVAR-SV	Minnesota	Intercept	None
NM-BVAR-SVO-TVM	BVAR-SV	Minnesota	Intercept	CCMM2024
AM-BVAR-SV	BVAR-SV	Adaptive Minnesota	None	None
AM-BVAR-SVO	BVAR-SV	Adaptive Minnesota	None	CCMM2024
AM-BVAR-SV-TVM	BVAR-SV	Adaptive Minnesota	Intercept	None
AM-BVAR-SVO-TVM	BVAR-SV	Adaptive Minnesota	Intercept	CCMM2024
NGM-BVAR-SV	BVAR-SV	Minnesota Normal-Gamma	None	None
NGM-BVAR-SVO	BVAR-SV	Minnesota Normal-Gamma	None	CCMM2024
NGM-BVAR-SV-TVM	BVAR-SV	Minnesota Normal-Gamma	Intercept	None
NGM-BVAR-SVO-TVM	BVAR-SV	Minnesota Normal-Gamma	Intercept	CCMM2024
SVD-GP-NC	SVD-WN	g-prior, No Clustering	Coefficients	None

Note: All models feature stochastic volatility in their innovations. Models with time variation in their intercepts assume a random walk state evolution. SVD models with time variation in their coefficients assume a white-noise state evolution.

3.1 Models

We consider three families of models. The first is multivariate and the other two are univariate.

Our first family consists of BVARs with stochastic volatility, referred to as the BVAR-SV specification. This specification takes the following form, for the $n \times 1$ data vector y_t , which includes aggregate core inflation, its three components, and wage growth (we exclude aggregate core inflation from y_t when we consider the “bottom-up” approach to forecasting aggregate core inflation):⁷

$$y_t = \Pi_0 + \sum_{i=1}^p \Pi_i y_{t-i} + v_t$$

$$v_t = A^{-1} \Lambda_t^{0.5} \epsilon_t, \quad \epsilon_t \sim N(0, I_n), \quad \Lambda_t \equiv \text{diag}(\lambda_{1,t}, \dots, \lambda_{n,t}) \quad (1)$$

$$\ln(\lambda_{i,t}) = \ln(\lambda_{i,t-1}) + \nu_{i,t}, \quad i = 1, \dots, n$$

$$\nu_t \equiv (\nu_{1,t}, \nu_{2,t}, \dots, \nu_{n,t})' \sim N(0, \Phi),$$

where A is a lower triangular matrix with ones on the diagonal and nonzero coefficients below the diagonal, and the diagonal matrix Λ_t contains the time-varying variances of conditionally Gaussian innovations. This model implies that the reduced-form variance-covariance matrix of innovations to the BVAR is $\text{var}(v_t) \equiv \Sigma_t = A^{-1} \Lambda_t A^{-1'}$. Note that, as in the implementation of Primiceri (2005), innovations to log volatility are allowed to be correlated across variables; Φ is not restricted to be diagonal. For notational simplicity, let Π denote the collection of the BVAR’s coefficients. In implementation, we include four lags in the BVAR.

⁷We order our variables so that wage growth is first, followed by supercore, housing, core goods, and aggregate core inflation. As noted in the seminal work of Primiceri (2005) and some subsequent studies, in BVARs with stochastic volatility specified as in this paper, variable ordering affects estimates. Work such as Bognanni (2018) has discussed the issue from a perspective of structural inference. Recent work by Arias, Rubio-Ramirez, and Shin (2023) has shown that ordering choices in VARs with time-varying parameters and stochastic volatility can affect out-of-sample forecasts. In particular, in their results, ordering has little effect on point forecasts but measurable effects on density-related measures, including the standard deviation of the predictive density and the length of prediction intervals. Empirically speaking, ordering impacts can be smaller in models with constant regression parameters.

Furthermore, we consider versions of the BVAR-SV model extended in two dimensions (separately as well as together). First, we treat the intercepts of the VAR as time-varying. In this case, the intercepts follow random walk processes; this makes the model a restricted version of the time-varying parameter VAR of Primiceri (2005), with time variation of intercepts and conditional volatilities but constant slope parameters in the VAR. The time variation of intercepts allows for time variation of the implied conditional means of the data, i.e., time shifts in trend inflation. Second, to reduce the impacts of the extreme, temporary volatility of macroeconomic data following the COVID-19 outbreak, we extend the model by including in the volatility piece the multiplicative volatility outlier specification of Carriero et al. (2024), originally developed by Stock and Watson (2016) for modeling inflation. Adding a second volatility factor that is independent over time serves to reduce the persistence of changes in volatility that affect forecasts over unusual periods such as the pandemic.

Of course, Bayesian estimation entails priors on the parameters of models. In specifying the prior variance matrix for the BVAR coefficients Π , we consider three options: a standard Minnesota-type Gaussian prior that shrinks all coefficients to zero but allows differing degrees of shrinkage for coefficients on own lagged variables as opposed to lags of other variables; a similarly formed but adaptive Minnesota-type prior developed by Chan (2021), where the shrinkage is estimated from the data; and Chan's new adaptive Minnesota-type normal-gamma prior that combines features of the Minnesota and hierarchical priors. Details are omitted in the interest of brevity, but follow the aforementioned citations. We apply these priors to the baseline BVAR-SV model as well as the versions extended to allow time-varying intercepts and volatility outliers.

Our second considered model family, which we call the SVD-WN specification, uses one of the new, novel approaches to the specification and estimation of large time-varying parameter (TVP) regressions (single regression, not a vector autoregression) developed in Hauzenberger et al. (2022). Hauzenberger et al. (2022) develop a computationally efficient estimation algorithm based on singular value decomposition (SVD) that permits application to large regressions and focuses on models that replace the usual assumption of random walk variation in parameters with a hierarchical

mixture model on the time-varying parameters. In particular, we use their specification that features a white-noise state equation and the g-prior of Zellner (1986), without any clustering of the parameters.⁸ This model is applied to one inflation measure at a time rather than a vector of measures as in the VAR specifications. Our SVD-WN model takes the same dynamic regression form as Hauzenberger et al. (2022):

$$\begin{aligned}\pi_t &= x_t' \beta_t + \sigma_t \eta_t, \quad \eta_t \sim N(0, 1) \\ \beta_t &= \gamma + \tilde{\beta}_t \\ \tilde{\beta}_t &\sim N(0, \sigma_t^2 \cdot \Psi) \\ \log(\sigma_t^2) &= \mu_\sigma + \rho_\sigma (\log(\sigma_{t-1}^2) - \mu_\sigma) + \sigma_\sigma \nu_t, \quad \nu_t \sim N(0, 1),\end{aligned}\tag{2}$$

where π_t is our scalar inflation measure, x_t is a vector of regressors consisting of lags of the data, β_t is a set of time-varying regression coefficients composed of a time-invariant component (γ) and a white-noise time-varying component ($\tilde{\beta}_t$), σ_t^2 is the time-varying error variance where its log volatility follows an AR(1) process, and Ψ is a diagonal matrix of shrinkage parameters. In our application, x_t consists of four lags of ECI, disaggregate inflation, and aggregate inflation.

Finally, we consider the univariate unobserved components-stochastic volatility (UC-SV) model that Stock and Watson (2007) established as a good model for forecasting aggregate inflation. We apply this model to aggregate core inflation as well as core inflation's three components. Letting π_t denote a given measure of inflation, this model treats inflation as the sum of a random walk and a serially uncorrelated noise term, both featuring stochastic volatility:

$$\begin{aligned}\pi_t &= \pi_t^* + v_t, \quad \pi_t^* = \pi_{t-1}^* + n_t \\ v_t &= \lambda_{v,t}^{0.5} \epsilon_{v,t}, \quad \epsilon_{v,t} \sim N(0, 1)\end{aligned}\tag{3}$$

⁸In unreported results, we also considered a specification grouping the time-varying parameters into a smaller number of regimes/clusters, which Hauzenberger et al. (2022) found to be helpful for forecasting inflation with a large number of predictors. However, we found that the clustering approach produced forecasts that performed essentially as well as forecasts without clustering, and in the interest of brevity we omit these results.

$$\begin{aligned}
n_t &= \lambda_{n,t}^{0.5} \epsilon_{n,t}, \quad \epsilon_{n,t} \sim N(0, 1) \\
\log(\lambda_{v,t}) &= \log(\lambda_{v,t-1}) + \nu_{v,t}, \quad \nu_{v,t} \sim N(0, \phi_v) \\
\log(\lambda_{n,t}) &= \log(\lambda_{n,t-1}) + \nu_{n,t}, \quad \nu_{n,t} \sim N(0, \phi_n).
\end{aligned}$$

As a note, in our VAR models the combination of inflation variables introduces a near-singularity in the variance-covariance matrix. This is because the inflation aggregate in the model is closely approximated by a time-varying linear combination of the three disaggregates of inflation that are also included in the model, leading to strong collinearity.⁹ In empirical work, this collinearity can result in a near-singular variance-covariance matrix that can, in turn, lead to explosive model estimates and numerical issues. We have found that this issue is less pronounced for the constant-coefficient models we use (as well as the versions with time-varying intercepts), but it is increasingly prevalent when a richer degree of time variation in regression parameters is allowed, such as the hybrid TVP specification and estimation approach of Chan (2023) that has TVP in some equations but constant coefficients in others, optimized for model fit. (As a result, we omit the hybrid TVP models of Chan 2023 from our paper.)

3.2 Estimation and Forecasting

We estimate all our models with Bayesian methods. For estimation of the BVAR-SV specification, we employ Gibbs sampling and use replication code from Carriero et al. (2024) to estimate these models following Carriero et al. (2022).¹⁰ For clarity, we estimate stochastic volatility following the methods in Kim, Shephard, and Chib (1998). See sources such as Clark and Ravazzolo (2015) and Karlsson (2013) for further details on estimating this family of models. Furthermore,

⁹Under the chain weighting used to construct the published core PCE price index, the core PCE inflation rate is approximated very well but not perfectly by a weighted average of the components, where the weights for inflation in quarter t are averages of spending shares in $t-1$ and t . Because of that approximation and time variation in the spending shares, the collinearity of the components with aggregate inflation is high but not perfect.

¹⁰The algorithm of Carriero et al. (2022) corrects a problem with the original algorithm of Carriero, Clark, and Marcellino (2019) pointed out by Bognanni (2022).

we follow Carriero et al. (2024) for the estimation of our outlier adjustment. We follow Chan and Eisenstat (2018) for the estimation of our time-varying intercept but continue to use the method of Carriero et al. (2022) for the rest of the model (particularly for the estimation of time-invariant coefficients). Finally, for the estimation of our different considered Minnesota-like priors we follow the estimation procedure of Chan (2021).

In our implementation of BVAR priors, in the baseline case, we follow the prior specifications of Carriero et al. (2024), which include a standard Minnesota prior on the matrix of coefficients Π_i for $i = 0, \dots, p$. However, since our application uses a quarterly frequency instead of the monthly frequency considered in Carriero et al. (2024), we make a modification and use an inverse Wishart prior for Φ , the variance matrix of innovations to log volatility, with a scale matrix of $0.05 \times (n + 3) \times I_n$ and $n + 3$ degrees of freedom. Furthermore, for simplicity, we assume prior means of 0 for all variables. The priors on our time-varying intercepts follow Chan and Eisenstat (2018). Finally, our adaptive Minnesota-like priors follow those of the baseline model in Chan (2021) with the exception of the prior means on hyperparameters κ_1 and κ_2 , which we set so as to imply prior means of .1 and .005, respectively.

For the estimation of our SVD-WN models we use replication code from Hauzenberger et al. (2022) to implement the SVD algorithm and to estimate a time-varying parameters model with a white-noise state equation, g-prior, but no clustering. Our priors also follow the specification of the original paper, and for the prior hyperparameter κ , which determines the upper bound of our innovation variance for the time-varying coefficients, we choose a value of .005.

At each forecast origin, we obtain forecasts from the posterior predictive distribution as follows (in the case of BVAR-SV models, with some adjustments for other models).¹¹ Over the forecast horizon, we draw time-series vectors of the VAR's shocks

¹¹Particularly, for the SVD-WN models, we follow Hauzenberger et al. (2022) and use the direct forecasting approach of Stock and Watson (2002). Thus, for a given forecasting horizon, h , we estimate the SVD-WN model of the form $\pi_t = x'_{t-h}\beta_t + \sigma_t\eta_t$ and produce forecasts for $\pi_{t+h} = x'_{t+h}\beta_{t+h} + \sigma_{t+h}\eta_{t+h}$. Values for β_{t+h} and σ_{t+h} are obtained by simulating paths of these variables following their respective laws of motion.

$\epsilon_{T+h}^{(m)} \sim N(0, I_N)$ and volatility shocks $\nu_{T+h}^{(m)} \sim N(0, \Phi^{(m)})$, feed in the volatility shocks and iterate forward to obtain the draw of the time series of (log-)variances over the forecast horizon, compute the draw of the VAR's reduced-form shocks, and use the autoregressive structure of the VAR to iterate forward to obtain a draw of the time series of forecasts as $y_{T+h}^{(m)} = \Pi_0^{(m)} + \sum_{i=1}^p \Pi_i^{(m)} y_{T+h-i}^{(m)} + v_{T+h}^{(m)}$. As noted by Waggoner and Zha (1999), for a given Markov chain Monte Carlo (MCMC) node m , the computational costs of simulating forward a path of stochastic volatilities $\ln \lambda_{T+h}^{(m)}$, innovations $v_{T+h}^{(m)}$, and outcomes $y_{T+h}^{(m)}$ can be substantially lower than generating an additional MCMC draw m' . To better balance computational costs against a high degree of accuracy in the Monte Carlo approximation of the predictive density, it might thus be advantageous to draw multiple trajectories $\ln \lambda_{T+h}^{(m,j)}$, innovations $v_{T+h}^{(m,j)}$, and outcomes $y_{T+h}^{(m,j)}$, indexed by $j = 1, 2, \dots, J$, for every m .¹² The posterior predictive density is then approximated by the ensemble of draws for all m and j . We follow this approximation for our BVAR-SV models, while directly simulating from the posterior predictive density for our SVD-WN and UC-SV models. Estimates and forecasts from the BVAR-SV models are based on 10,000 retained MCMC draws with a burn-in of 2,000 draws. For each MCMC draw, we approximate the posterior predictive density using 25 trajectories, for a total of 250,000 simulated posterior predictive density draws. Estimates and forecasts from the SVD-WN models are based on 30,000 retained MCMC draws with a burn-in of 10,000 draws. For the estimation of our UC-SV model we use the precision sampler of Chan and Jeliazkov (2009) and follow both the prior specification and estimation routine of Section 7.2 in Chan (2017). Estimates and forecasts from the UC-SV model are based on 50,000 retained MCMC draws with a burn-in of 5,000 draws.

We produce out-of-sample forecasts using a recursive scheme, expanding the model estimation sample as forecasting moves forward

¹²In the context of simulating conditional forecasts with only occasionally binding conditions, Waggoner and Zha (1999) provide a formal discussion of the optimal trade-off in choosing J , which also extends to the case of unconditional forecasts (i.e., density forecasts without binding conditions imposed on the forecasts).

in time. Forecasting begins with an origin, t , of 1985:Q1, estimating models with a data sample of 1959:Q2 through 1985:Q1. Forecasts are then produced for $t + h$, where h is the forecast horizon. We then move forward a quarter to the origin of 1985:Q2, reestimate models with data through 1985:Q2, and continue from there. We examine forecast accuracy for three periods: a full-sample period from 1985:Q1 through 2023:Q4, a prepandemic period of 1985:Q1 through 2017:Q4 (with 2017:Q4 as the last forecast origin, the last forecast outcome date considered is 2019:Q4 for $h = 8$), and a period of 2018:Q1 through 2023:Q4 that includes the unusual volatility of the pandemic period. Over this subsample, core inflation swung from 1.7 percent in 2020:Q1 to -0.8 percent in 2020:Q2 and up to 3.1 percent in 2020:Q3. It then jumped to 5.8 percent in 2021:Q2. We report results for forecast horizons of $h = 1, 2, 4$, and 8 quarters ahead.

We evaluate point and density forecasts (computed from posterior predictive densities of inflation, with point forecasts being the means of these densities) based on root mean squared errors (RMSE) and continuous ranked probability scores (CRPS), respectively, as described in, among others, Clark and Ravazzolo (2015) and Krüger et al. (2021). To roughly gauge the significance of differences with respect to the baseline forecast described below, we use t -tests as in Diebold and Mariano (1995) and West (1996), computed with Newey-West standard errors and $h + 1$ lags.

We consider a multitude of models, forecast evaluation measures, forecast horizons, evaluation windows, and variables of interest. As a result, it can be difficult to interpret broad patterns from our results. We use diffusion indices to help address this problem, inspired by Knotek and Zaman (2019). For each model i , each forecast evaluation measure m , each forecast horizon h , each evaluation window w , and each variable of interest v , let $\mu_{i,m,h,w,v}$ be the ratio of the realization of this evaluation measure (that is, RMSE or CRPS) to some realized evaluation measure produced by a benchmark model, where values under 1 suggest improvement in the measure over the benchmark. For a given evaluation measure m and for each i, h, w , and v , let

$$f_{i,m,h,w,v} = \begin{cases} 1, & \text{if } \mu_{i,m,h,w,v} \geq 1 \\ -1 & \text{if } \mu_{i,m,h,w,v} < 1. \end{cases} \quad (4)$$

Then, analogous to Knotek and Zaman (2019), for a given evaluation measure m , the diffusion index at forecast horizon h , with evaluation window w , and for variable of interest v is

$$(\text{Diffusion Index})_{m,h,w,v} = \sum_{i=1}^r f_{i,m,h,w,v}, \quad (5)$$

where r is the number of models we consider.

For scenario analysis, we adopt entropic tilting, which can be seen as a nonparametric approach to conditional forecasting, used in studies such as Cogley, Morozov, and Sargent (2005), Krüger, Clark, and Ravazzolo (2017), Robertson, Tallman, and Whiteman (2005), and Tallman and Zaman (2020).¹³ Tilting consists of reweighting the draws from the model's predictive distribution so as to obtain a new distribution that satisfies the conditions of interest but remains as close as possible to the original distribution, where closeness is measured by the Kullback-Leibler divergence. In our applications, the conditions of interest are either the path of wage growth or the path of aggregate core inflation over the forecast horizon (the paths are detailed in Section 5). We tilt the first moments of our predictive distribution toward these paths.

4. Historical Forecast Evaluation

This section reports a historical evaluation of the performance of our models in unconditional forecasting. In these comparisons, we use the BVAR-SV model with the adaptive Minnesota prior (AM-BVAR-SV) as the baseline, for a few reasons. First, we are interested in being able to use multivariate models (BVARs) to capture joint dynamics and predictive distribution questions. Second, as will become clear, this model and prior choice work relatively well, and using it as the baseline facilitates comparisons and punchlines in our presentation. Third, a prior of the Minnesota form has a long, successful track record in macroeconomic forecasting, and the adaptive formulation of Chan (2021) is appealing for making the prior's tightness depend on the data.

¹³See Antolín-Díaz, Petrella, and Rubio-Ramírez (2021) for formal proof of the equivalence of entropic tilting to conditional forecasting under Gaussianity.


We present our forecasting results for core PCE inflation and its components using ratios of RMSE and CRPS to the corresponding measure for the AM-BVAR-SV baseline forecast. In this case, values below (above) 1 indicate improvement (deterioration) relative to the AM-BVAR-SV model. Forecasts for inflation variables are taken directly from the posterior predictive density implied by our models. That is, since our models include core inflation and its disaggregates, in our main results we directly use the recursive forecasts (of either core inflation or one of its components) from these models for our forecast evaluation. We separately consider the efficacy of obtaining forecasts of aggregate core inflation by aggregating forecasts of the disaggregate components. Throughout, we first consider results for the relatively long subsample of 1985–2017 that ends before the pandemic-induced additional volatility, and then we consider results for the shorter 2018–23 subsample and the full 1985–2023 sample.

This section proceeds to first present the accuracy of forecasts of aggregate inflation and its components obtained directly from our models and then examine aggregate forecasts instead obtained by aggregating forecasts of inflation components.

4.1 Accuracy of Model Forecasts of Aggregate Inflation and Its Components


To first provide a broad, visual summary of how alternative forecast models compare with the AM-BVAR-SV baseline, Figure 1 presents diffusion indices for the aggregate and disaggregate point forecasts across evaluation samples and forecasting horizons, and Figure 2 presents analogous findings for density forecasts. These visual summaries abstract from the magnitudes of differences in accuracy and their statistical significance. In these displays, results for CRPS are very similar to those for RMSE. For the 1985–2017 sample, the gray bars indicate that the baseline AM-BVAR-SV model generally performs best (yielding high index readings), although with reduced advantage for supercore and core goods inflation at horizons of $h = 4$ and 8. For the 2018–23 sample, the dominance of the baseline model is reduced, more so for the inflation disaggregates than core PCE inflation. On balance, for most measures, samples, and horizons, the AM-BVAR-SV baseline beats most models, whereas it is not normally beaten by most models.

Figure 1. Diffusion Index of RMSE Ratios for PCE Inflation



Note: Ratios are calculated so that values below 1 indicate improvement over the AM-BVAR-SV model. Diffusion indices are calculated such that we take the count of ratios less than 1 and subtract it from the count of ratios greater than 1, for each horizon, application, and evaluation window.

Figure 2. Diffusion Index of CRPS Ratios for PCE



Note: Ratios are calculated so that values below 1 indicate improvement over the AM-BVAR-SV model. Diffusion indices are calculated such that we take the count of ratios less than 1 and subtract it from the count of ratios greater than 1, for each horizon, application, and evaluation window.

Starting from that broad perspective, we turn to more detailed, quantitative comparisons of accuracy across model specifications. Tables 2–5 present our forecasting results for core PCE inflation and its components, expressed in ratios of RMSE and CRPS to the corresponding measure for the AM-BVAR-SV baseline forecast. Results for the relatively long subsample of 1985–2017 that ends before the pandemic-induced additional volatility appear in the middle panels of Table 2 for core inflation and Tables 3–5 for the three disaggregate components.

Consider first the choice of prior for the BVAR-SV model. Across measures of inflation, the baseline adaptive Minnesota prior (i.e., the AM-BVAR-SV forecast) is consistently at least as good as the simple Minnesota (NM-BVAR-SV) and adaptive Minnesota-type normal-gamma priors (NGM-BVAR-SV). For aggregate inflation, the NM and NGM priors yield forecasts (both point and density) slightly less accurate than the AM baseline. For the three disaggregates, forecasts with the NM and NGM priors are sometimes slightly better than the baseline, but by small margins and without statistical significance. More frequently, these other priors are slightly to modestly less accurate than the baseline.

Now consider the additional model features allowed, which are time-varying means (TVM) and volatility outliers (SVO). Over the 1985–2017 sample, for a given choice of prior, across measures of inflation, neither of these additional model features offers any consistent improvement over the corresponding BVAR-SV specification. But there can be occasional (across the choice of measure and horizon) gains with time-varying means and volatility outliers. More specifically, TVM improves forecast accuracy gains in some cases, more typically at longer horizons (as might be expected, because as the horizon lengthens, the model-implied mean plays a larger role than recent dynamics as the forecast gradually converges to the unconditional mean), but not consistently. Under the baseline AM prior, in forecasting aggregate core inflation, adding TVM slightly reduces forecast accuracy, but not statistically significantly. Among disaggregate components, TVM offers gains for supercore and core goods inflation at long horizons, but harms accuracy for housing. These same patterns with the addition of TVM apply under other BVAR priors. Our finding on the mixed benefits of allowing time variation in means has precedent in studies such as Chan (2023),

Table 2. Forecast Ratios for RMSEs and CRPSs, Core PCE Inflation

Model/Horizon	$h = 1$		$h = 2$		$h = 4$		$h = 8$	
	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS
1985:Q1–2023:Q4								
UC-SV	1.059	1.006	1.078	1.012	1.039	1.007	1.037	0.995
NM-BVAR-SV	1.008	1.010	1.017	1.020	1.021	1.021	1.025	1.014
NM-BVAR-SVO	1.011	1.024**	1.020	1.047***	1.024	1.084***	1.027	1.118***
NM-BVAR-SV-TVM	1.004	1.012	1.030**	1.033**	1.034*	1.039**	1.038	1.029
NM-BVAR-SVO-TVM	1.009	1.022**	1.027	1.050***	1.034	1.087***	1.031	1.100***
AM-BVAR-SVO	1.001	1.012***	1.000	1.025***	1.000	1.058***	1.000	1.093***
AM-BVAR-SV-TVM	1.007	1.009	1.030	1.023	1.025	1.022	1.023	1.012
AM-BVAR-SVO-TVM	1.011	1.018**	1.024	1.038**	1.021	1.064***	1.010	1.071***
NGM-BVAR-SV	1.008*	1.009***	1.003	1.008***	1.002	1.009***	1.004	1.013***
NGM-BVAR-SVO	1.009*	1.020***	1.004	1.035***	1.001	1.071***	1.003	1.112***
NGM-BVAR-SV-TVM	1.011	1.011	1.032	1.018	1.015	1.008	1.017	1.001
NGM-BVAR-SVO-TVM	1.017	1.025*	1.027	1.034*	1.012	1.048**	1.004	1.054**
SVD-GP-NC	1.021*	1.008	1.022*	1.007	1.052	1.021	1.052	1.007
1985:Q1–2017:Q4								
UC-SV	0.984	0.972	0.965*	0.941***	0.952*	0.927***	1.021	0.939**
NM-BVAR-SV	1.011	1.012	1.038	1.032*	1.046	1.032	1.011	1.003
NM-BVAR-SVO	1.020	1.031**	1.043	1.067***	1.048	1.119***	1.014	1.167***
NM-BVAR-SV-TVM	1.005	1.015	1.047**	1.042**	1.070**	1.051**	1.071	1.038
NM-BVAR-SVO-TVM	1.015	1.029**	1.044*	1.066***	1.067	1.119***	1.065	1.154***
AM-BVAR-SVO	1.006*	1.017***	1.002	1.035***	0.997	1.082***	1.003	1.147***
AM-BVAR-SV-TVM	1.003	1.008	1.020	1.018	1.024	1.016	1.030	1.009
AM-BVAR-SVO-TVM	1.009	1.020**	1.012	1.041***	1.012	1.078***	1.008	1.107***
NGM-BVAR-SV	1.012**	1.011***	1.007*	1.011***	1.005	1.014***	1.012**	1.021***
NGM-BVAR-SVO	1.018**	1.027***	1.009	1.048***	1.001	1.101***	1.016	1.175***
NGM-BVAR-SV-TVM	1.005	1.010	1.014	1.007	1.007	0.996	1.014	0.988
NGM-BVAR-SVO-TVM	1.018	1.028**	1.008	1.030**	0.999	1.055***	0.995	1.078**
SVD-GP-NC	1.026*	1.011	1.028	0.997	1.048	0.982	1.102	0.979

(continued)

Table 2. (Continued)

Model/ Horizon	$h = 1$		$h = 2$		$h = 4$		$h = 8$	
	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS
2018:Q1–2023:Q4								
UC-SV	1.150	1.105	1.203	1.215	1.105	1.199	1.047*	1.116**
NM-BVAR-SV	1.004	1.005	0.990	0.988	1.001	0.997	1.035**	1.038**
NM-BVAR-SVO	0.998	1.003	0.990	0.987	1.004	0.999	1.036**	1.012
NM-BVAR-SV-TVM	1.002	1.003	1.010	1.009	1.004	1.011	1.014	1.010
NM-BVAR-SVO-TVM	1.002	1.003	1.006	1.002	1.006	1.008	1.007	0.983*
AM-BVAR-SVO	0.995	0.998	0.998	0.997	1.002	1.000	0.998	0.976*
AM-BVAR-SV-TVM	1.012	1.011	1.043	1.039	1.026	1.036	1.019	1.019
AM-BVAR-SVO-TVM	1.013	1.010	1.038	1.030	1.029	1.031	1.011	0.990
NGM-BVAR-SV	1.002	1.002	0.998	0.998	1.000	0.998	0.997**	0.995***
NGM-BVAR-SVO	0.998	1.001	0.997	0.996	1.001	0.998	0.995*	0.972*
NGM-BVAR-SV-TVM	1.018	1.015	1.055	1.048	1.021	1.037	1.020	1.030
NGM-BVAR-SVO-TVM	1.017	1.016	1.051	1.043	1.023	1.032	1.010	1.000
SVD-GP-NC	1.014	1.000	1.015	1.034	1.055	1.117	1.016	1.066

Note: Values below 1 indicate improvement over the AM-BVAR-SV model. Significance assessed by Diebold-Mariano-West test using Newey-West standard errors with $h + 1$ lags. *, **, and *** represent the .10, .05, and 0.1 significance levels, respectively. Entries in bold denote ratios below 1 that are the lowest for a given forecasting horizon and measure. In columns with no bold entries, the benchmark model is best.

Table 3. Forecast Ratios for RMSEs and CRPSs,
Core PCE Services Excl. Housing Inflation

Model/Horizon	$h = 1$		$h = 2$		$h = 4$		$h = 8$	
	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS
1985:Q1–2023:Q4								
UC-SV	1.050	1.040	1.094**	1.069**	1.084**	1.075**	1.077**	1.072**
NM-BVAR-SV	0.997	1.008	1.003	1.009	1.004	1.004	0.998	0.990
NM-BVAR-SVO	1.002	1.020	1.011	1.027	1.004	1.043*	1.000	1.074**
NM-BVAR-SV-TVM	0.994	1.006	1.013	1.014	1.006	1.000	0.997	0.985
NM-BVAR-SVO-TVM	1.004	1.018	1.013	1.024*	1.013	1.040**	0.991	1.046
AM-BVAR-SVO	1.002	1.009***	1.005**	1.017***	0.999	1.037***	1.001	1.074***
AM-BVAR-SV-TVM	1.000	1.004	1.011	1.008	0.996	0.988	0.997	0.981
AM-BVAR-SVO-TVM	1.007	1.014	1.013	1.021	1.001	1.023	0.982	1.029
NGM-BVAR-SV	1.010*	1.009**	1.002	1.007**	1.001	1.007***	1.006*	1.013***
NGM-BVAR-SVO	1.011	1.017***	1.007	1.024***	0.999	1.048***	1.007	1.093***
NGM-BVAR-SV-TVM	1.006	1.011	1.009	1.002	0.994	0.986	0.996	0.980
NGM-BVAR-SVO-TVM	1.015	1.024	1.012	1.015	1.000	1.021	0.983	1.024
SVD-GP-NC	1.002	1.003	1.024*	1.022	1.033	1.019	1.007	0.971
1985:Q1–2017:Q4								
UC-SV	1.034	1.034	1.078*	1.053	1.082*	1.063**	1.092*	1.065*
NM-BVAR-SV	0.995	1.008	1.007	1.015	1.001	1.004	0.978	0.976
NM-BVAR-SVO	1.004	1.023	1.015	1.035	0.999	1.050*	0.980	1.091**
NM-BVAR-SV-TVM	0.996	1.009	1.017	1.019	1.003	0.997	0.983	0.975
NM-BVAR-SVO-TVM	1.008	1.022	1.019	1.032*	1.011	1.045**	0.977	1.058
AM-BVAR-SVO	1.004	1.011***	1.005*	1.020***	0.996	1.044***	1.002	1.100***
AM-BVAR-SV-TVM	1.003	1.008	1.015	1.013	0.995	0.986	0.983	0.969*
AM-BVAR-SVO-TVM	1.012	1.020	1.020	1.029*	0.999	1.027*	0.965*	1.035
NGM-BVAR-SV	1.011*	1.010**	1.003	1.009**	1.002	1.009***	1.010**	1.018***
NGM-BVAR-SVO	1.013	1.020***	1.008	1.028***	0.997	1.057***	1.013*	1.124***
NGM-BVAR-SV-TVM	1.014	1.020	1.009	1.004	0.989	0.980	0.982	0.964**
NGM-BVAR-SVO-TVM	1.025	1.034**	1.013	1.020	0.994	1.021	0.967	1.025
SVD-GP-NC	0.998	1.001	1.021	1.017	1.040	1.013	1.038	0.974

(continued)

Table 3. (Continued)

Model/Horizon	h = 1		h = 2		h = 4		h = 8	
	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS
2018:Q1–2023:Q4								
UC-SV	1.107	1.066	1.150	1.139	1.087	1.119	1.051	1.094
NM-BVAR-SV	1.001	1.010	0.990	0.986	1.010	1.003	1.032**	1.034*
NM-BVAR-SVO	0.996	1.009	0.996	0.995	1.017	1.017	1.035**	1.019
NM-BVAR-SV-TVM	0.990	0.995	1.000	0.995	1.012	1.011	1.021	1.018
NM-BVAR-SVO-TVM	0.990	0.997	0.992	0.990	1.018	1.022	1.015	1.003
AM-BVAR-SVO	0.996	1.001	1.004	1.007	1.005	1.012	1.000	0.988
AM-BVAR-SV-TVM	0.989	0.985	0.999	0.989	0.998	0.994	1.020	1.024
AM-BVAR-SVO-TVM	0.988	0.987	0.991	0.986	1.007	1.007	1.013	1.008
NGM-BVAR-SV	1.005	1.004	0.998	0.999	1.001	1.001	0.998	0.996**
NGM-BVAR-SVO	1.003	1.005	1.003	1.007	1.004	1.013	0.997	0.985
NGM-BVAR-SV-TVM	0.976	0.972	1.008	0.994	1.005	1.005	1.020	1.033
NGM-BVAR-SVO-TVM	0.974	0.977	1.005	0.995	1.015	1.021	1.012	1.017
SVD-GP-NC	1.015	1.008	1.037	1.047	1.017	1.042	0.948*	0.958

Note: Values below 1 indicate improvement over the AM-BVAR-SV model. Significance assessed by Diebold-Mariano-West test using Newey-West standard errors with $h + 1$ lags. *, **, and *** represent the .10, .05, and 0.1 significance levels, respectively. Entries in bold denote ratios below 1 that are the lowest for a given forecasting horizon and measure. In columns with no bold entries, the benchmark model is best.

Table 4. Forecast Ratios for RMSEs and CRPSs, PCE Housing Inflation

Model/Horizon	h = 1		h = 2		h = 4		h = 8	
	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS
1985:Q1–2023:Q4								
UC-SV	1.001	0.998	1.044	1.047	1.129**	1.119**	1.175**	1.155**
NM-BVAR-SV	1.018	1.018	1.008	1.009	1.015	1.014	1.026*	1.039*
NM-BVAR-SVO	1.016	1.019	1.008	1.011	1.017	1.023	1.027**	1.070**
NM-BVAR-SV-TVM	0.997	0.989	0.980	0.977	0.982	0.989	1.030	1.035
NM-BVAR-SVO-TVM	0.997	0.993	0.975	0.978	0.979	0.993	1.027	1.056
AM-BVAR-SVO	0.997**	0.999	1.000	1.001	1.000	1.006	0.998	1.028
AM-BVAR-SV-TVM	0.995	0.986	0.986	0.983	0.997	1.003	1.053	1.049
AM-BVAR-SVO-TVM	0.995	0.989	0.982	0.984	0.996	1.007	1.051	1.067
NGM-BVAR-SV	0.997	0.996*	0.998**	0.998	0.999	0.999	0.998	0.999
NGM-BVAR-SVO	0.994*	0.995	0.996**	0.999	0.998	1.006	0.997	1.029
NGM-BVAR-SV-TVM	1.236**	1.205***	1.056	1.043	1.023	1.017	1.050	1.047
NGM-BVAR-SVO-TVM	1.239**	1.209***	1.054	1.043	1.022	1.019	1.051	1.060
SVD-GP-NC	1.005	1.002	1.039	1.029	1.072*	1.070**	1.026	1.040
1985:Q1–2017:Q4								
UC-SV	1.013	1.002	1.041	1.036	1.103*	1.097*	1.202**	1.167**
NM-BVAR-SV	1.016	1.019*	1.003	1.010	1.004	1.012	1.040	1.053*
NM-BVAR-SVO	1.015	1.020*	1.003	1.015	1.006	1.031	1.044*	1.108***
NM-BVAR-SV-TVM	1.018	1.003	1.022	1.002	1.033	1.021	1.075	1.065
NM-BVAR-SVO-TVM	1.018	1.006	1.015	1.003	1.027	1.031	1.072	1.105*
AM-BVAR-SVO	0.997**	1.000	0.999	1.004	0.999	1.015	0.999	1.052**
AM-BVAR-SV-TVM	1.018	1.001	1.032	1.010	1.054	1.037	1.096	1.072
AM-BVAR-SVO-TVM	1.017	1.004	1.028	1.013	1.050	1.047	1.097	1.110*
NGM-BVAR-SV	0.997	0.996	0.998	0.998	0.998	0.999	0.997	0.999
NGM-BVAR-SVO	0.994	0.996	0.995*	1.001	0.997	1.015	0.996	1.054**
NGM-BVAR-SV-TVM	1.188**	1.174***	1.066	1.048	1.051	1.033	1.086	1.058
NGM-BVAR-SVO-TVM	1.186**	1.177***	1.063	1.049	1.048	1.040	1.091	1.087
SVD-GP-NC	1.000	0.998	1.023	1.026	1.065	1.068	1.050	1.057

(continued)

Table 4. (Continued)

Table 5. Forecast Ratios for RMSEs and CRPSs, Core PCE Goods Inflation

Model/Horizon	h = 1		h = 2		h = 4		h = 8	
	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS
1985:Q1–2023:Q4								
UC-SV	1.057	1.015	1.075	1.023	1.081	1.050	1.035	0.992
NM-BVAR-SV	1.023*	1.021**	1.023	1.020	0.994	0.992	1.009	0.993
NM-BVAR-SVO	1.019	1.022**	1.018	1.023	0.992	1.013	1.007	1.037
NM-BVAR-SV-TVM	1.026	1.014	1.037	1.014	0.993	0.982	1.004	0.974
NM-BVAR-SVO-TVM	1.026	1.013	1.039	1.021	0.988	0.992	0.998	0.998
AM-BVAR-SVO	0.997*	1.000	0.996***	1.003	0.998	1.017**	0.998	1.033**
AM-BVAR-SV-TVM	1.013	0.999	1.024	0.995	0.989	0.970	0.982	0.949
AM-BVAR-SVO-TVM	1.011	0.996	1.020	0.997	0.982	0.976	0.974	0.967
NGM-BVAR-SV	1.004*	1.006***	1.004*	1.006**	1.001	1.005**	1.003*	1.009***
NGM-BVAR-SVO	1.002	1.007***	1.000	1.010***	0.998	1.023***	1.000	1.044**
NGM-BVAR-SV-TVM	1.021	1.001	1.024	0.988	0.981	0.960	0.977	0.942
NGM-BVAR-SVO-TVM	1.016	0.997	1.019	0.990	0.973	0.965	0.967	0.958
SVD-GP-NC	1.021**	1.020	1.028	1.030	1.032*	1.047	1.058***	1.052*
1985:Q1–2017:Q4								
UC-SV	0.967	0.967	0.949	0.949	0.995	0.964	0.949	0.937
NM-BVAR-SV	1.042**	1.030**	1.044	1.032	0.998	0.993	0.978	0.974
NM-BVAR-SVO	1.039**	1.032**	1.037	1.038*	0.993	1.025	0.974	1.045
NM-BVAR-SV-TVM	1.037	1.016	1.034	1.012	1.003	0.981	1.003	0.960
NM-BVAR-SVO-TVM	1.035	1.014	1.038	1.022	0.998	0.996	0.996	1.001
AM-BVAR-SVO	0.998	1.002	0.993**	1.005	0.955	1.026***	1.000	1.055***
AM-BVAR-SV-TVM	1.005	0.993	0.990	0.977	0.971	0.955	0.956	0.924*
AM-BVAR-SVO-TVM	0.998	0.987	0.985	0.982	0.954	0.966	0.944	0.957
NGM-BVAR-SV	1.008***	1.008***	1.009**	1.008***	1.005	1.008***	1.011***	1.015***
NGM-BVAR-SVO	1.006*	1.009***	1.001	1.015***	0.999	1.036***	1.009**	1.072***
NGM-BVAR-SV-TVM	0.999	0.988	0.976	0.966	0.952	0.939*	0.936	0.909**
NGM-BVAR-SVO-TVM	0.990	0.981	0.967	0.969	0.935*	0.948	0.920	0.938
SVD-GP-NC	1.029*	1.020	1.044	1.033	1.044	1.032	1.091**	1.051

(continued)

Table 5. (Continued)

Model/Horizon	$h = 1$		$h = 2$		$h = 4$		$h = 8$	
	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS
2018:Q1–2023:Q4								
UC-SV	1.141	1.134	1.189*	1.205	1.139*	1.243*	1.093*	1.116***
NM-BVAR-SV	1.004	1.001	1.000	0.990	0.991	0.989	1.032*	1.036*
NM-BVAR-SVO	0.998	0.998	0.999	0.988	0.991	0.986	1.030*	1.018
NM-BVAR-SV-TVM	1.015	1.008	1.040	1.019	0.986	0.984	1.004	1.007
NM-BVAR-SVO-TVM	1.017	1.011	1.041	1.017	0.989	0.984	0.999	0.991
AM-BVAR-SVO	0.996*	0.997	0.998	0.997	0.999	0.996	0.996**	0.984
AM-BVAR-SV-TVM	1.022	1.014	1.057	1.038	1.002	1.004	1.000	1.005
AM-BVAR-SVO-TVM	1.024	1.017	1.055	1.034	1.002	0.999	0.995	0.988
NGM-BVAR-SV	1.001	1.002	0.999	0.999	0.999	0.997	0.998***	0.997***
NGM-BVAR-SVO	0.998	1.001	0.999	0.997	0.998	0.993	0.993**	0.981
NGM-BVAR-SV-TVM	1.042	1.034	1.070	1.045	1.001	1.008	1.006	1.017
NGM-BVAR-SVO-TVM	1.042	1.037	1.069	1.043	0.999	1.002	1.000	1.001
SVD-GP-NC	1.014	1.019	1.011	1.023	1.023	1.079	1.033***	1.054**

Note: Values below 1 indicate improvement over the AM-BVAR-SV model. Significance assessed by Diebold-Mariano-West test using Newey-West standard errors with $h + 1$ lags. *, **, and *** represent the .10, .05, and 0.1 significance levels, respectively. Entries in bold denote ratios below 1 that are the lowest for a given forecasting horizon and measure. In columns with no bold entries, the benchmark model is best.

in which inclusion of time variation does not necessarily improve forecasting performance.

The addition of volatility outliers to the baseline BVAR-SV specification has essentially no benefit in the 1985–2017 sample of forecasts. For aggregate inflation, under the baseline AM prior, SVO typically harms accuracy a little, more sharply for density forecasts than point (as might be expected, as volatility outlier modeling will affect the spreads of densities more than the means or medians of predictive distributions), at longer horizons more than shorter. For example, at the horizon $h = 8$, the AM-BVAR-SVO forecast is less accurate than baseline by 15 percent for CRPS and less than 1 percent for RMSE. The same broad pattern also applies with inflation disaggregates. It also applies under the NM and NGM priors with SVO as compared with the SV specification.

The performance of the AM-BVAR-SV baseline as compared with univariate models is more mixed. While adding time-varying means to the BVAR models at least occasionally yields some gains in the accuracy of 1985–2017 forecasts, adding richer time variation in parameters in the form of the SVD-GP-NC specification is generally less successful, although with variation across inflation measures and horizons. Forecasts from this model are often—although not always—less accurate than the AM-BVAR-SV baseline and its AM-BVAR-SV-TVM extension that allows time-varying means.

On the other hand, stripping the model down to the UC-SV specification in which inflation is the sum of a time-varying trend and noise can help forecast accuracy for some inflation measures. Consistent with the literature, for forecasting aggregate core inflation, the UC-SV model consistently and modestly (statistically significantly in a majority of cases) improves on the accuracy of the AM-BVAR-SV baseline, by roughly 2 to 5 percent in RMSE and 3 to 7 percent in CRPS. The same applies to forecasts of core goods inflation, with comparable UC-SV gains but without any statistical significance. However, for both point and density forecasts of core services excluding housing and housing inflation, the UC-SV model is less accurate than the AM-BVAR-SV baseline. For these measures, the BVAR's advantage tends to increase with the horizon, peaking at about 10 percent for supercore and 20 percent for housing.

In the shorter subsample of 2018–23 that includes the volatility of the pandemic period, the patterns in forecast accuracy—reported

in the bottom panel of each table—broadly resemble those for the longer subsample, but with some differences. Considering the choice of prior for the BVAR-SV model, it continues to be the case that the AM prior is as good as or better than the NM and NGM priors, although not without some exceptions. In this shorter sample, for both point and density forecast accuracy, the NM and NGM priors perform very similarly to the AM baseline for aggregate inflation, and sometimes slightly better but other times slightly worse for the disaggregate inflation measures (rarely with statistical significance, except at the horizon $h = 8$ for which the sample includes very few independent observations).

As to the additional model features of time-varying means and outlier volatility, in this shorter subsample these features have modestly more benefit than in the longer subsample, but not on a consistent basis. In the shorter subsample, the addition of time-varying means to the baseline AM-BVAR-SV specification slightly reduces the accuracy of forecasts of aggregate and core goods inflation, whereas, at horizons of $h = 1$ through 4 quarters ahead, adding TVM very slightly improves accuracy for supercore inflation and notably improves accuracy (by up to 13 percent) for housing. These same patterns obtained with the baseline AM prior also apply under the NM and NGM priors.

In the case of volatility outliers (SVO), the addition of this feature to the baseline AM-BVAR-SV model has little impact on the accuracy of point forecasts but offers small improvements to the accuracy of density forecasts at longer horizons. At the horizon $h = 8$, the SVO specification improves on the baseline CRPS accuracy by roughly 2 to 4 percent. The finding that the outlier modeling improves forecasts in samples including the pandemic but not generally before is consistent with the macroeconomic forecast results of Carriero et al. (2024).

As to the accuracy of BVAR models compared with univariate specifications, the SVD-GP-NC model that allows rich time variation in parameters has less success in improving forecast accuracy in the shorter subsample than in the longer subsample. Except in the case of supercore forecasts at the horizon $h = 8$, this model is dominated by the AM-BVAR-SV baseline in the recent subsample.

Perhaps more notably, our baseline AM-BVAR-SV model performs better compared with the UC-SV model in the shorter

subsample than the longer subsample. For most inflation measures and horizons, point and density forecasts from the UC-SV model are less accurate than forecasts from the AM-BVAR-SV baseline, by notable degrees (e.g., by roughly 20 percent in both RMSE and CRPS of core PCE forecasts at $h = 2$). In this subsample, the UC-SV model only offers gains for shorter horizon forecasts of housing inflation and otherwise falls short of the baseline forecasts. As we will show below, the better performance of the BVAR baseline in the more recent sample reflects some tendency for its forecasts to more quickly pick up the surge of inflation that began in 2021.

As might be expected, forecast accuracy results for the full sample of 1985–2023 (provided in the top panel of the tables) contain elements of the patterns in the subsample results. In the full-sample results on point and density forecast accuracy, as in the subsamples, among the priors for the BVAR-SV model, the AM prior is in most cases slightly better than the NM and NGM priors, but occasionally slightly worse (with very few instances of statistical significance). Given the AM prior, extending the BVAR-SV model to include time-varying means (TVM) most commonly makes forecasts slightly less accurate but occasionally makes them more accurate (e.g., housing inflation for some horizons, and with gains peaking at 5 percent for core goods at $h = 8$), not by enough to be statistically significant. Similarly, extending the BVAR-SV model to include volatility outliers (SVO) has no consistent benefit; in a few cases the inclusion of volatility outliers very slightly improves accuracy but in most cases reduces it.

As to univariate models compared with the AM-BVAR-SV specification, over the full sample the accuracy of forecasts from the SVD-GP-NC model are nearly always worse than baseline, for all inflation measures and horizons and both point and density forecasts (with one exception, for the CRPS of supercore inflation at $h = 8$). In the case of the UC-SV specification, for both point and density forecasts of core PCE inflation, this univariate trend plus noise model is (except in the CRPS result at $h = 8$) a little less accurate than the AM-BVAR-SV baseline, without any differences that are statistically significant. The baseline AM-BVAR-SV has a larger advantage over the UC-SV model in accuracy of inflation components (supercore and housing inflation modestly more so than goods inflation) than in accuracy of core PCE inflation.

In order to shed some light on time-varying forecast performance, Figure 3 presents time series of aggregate core PCE inflation forecasts from the UC-SV, AM-BVAR-SV, and SVD-GP-NC models.¹⁴ Dates in this figure represent the date of each forecast origin, t . As the figure shows, in broad terms these forecasts generally track together for much of the sample. Near the end of the sample, the AM-BVAR-SV forecast appears to react more quickly to increases in past inflation, rising more quickly than the UC-SV forecast, and subsequently declining more quickly than the UC-SV forecast. That reaction likely helps account for the better performance—for aggregate inflation—of the BVAR-SV model compared with UC-SV in the 2018–23 subsample relative to the 1985–2017 subsample. Future research will be needed to determine whether postpandemic forecast accuracy for these models returns to prepandemic patterns.

Figure 4 focuses on the period of 2020–23, and presents time series of aggregate core PCE inflation forecasts from the UC-SV, AM-BVAR-SV, and the SVD-GP-NC models along with median forecasts from the Survey of Professional Forecasters (SPF). SPF forecasts are aligned with model forecasts so that, for a given SPF release, the last available quarterly inflation date is equivalent to the date of our forecast origin, t . The figure shows that our model forecasts, particularly the AM-BVAR-SV model, have tended to be more accurate than the SPF forecasts. Both our model forecasts and the SPF forecasts have tended to lag behind inflation’s rise and subsequent fall.


Overall, we view these results as indicating the BVAR-SV model with an adaptive Minnesota-style prior to be an effective tool for forecasting inflation in the core services excluding housing, housing, and core goods components of core PCE inflation.

4.2 “Bottom-Up” Approach to Aggregate Inflation Forecasting

A natural question for research is whether aggregate inflation is best forecast directly or by aggregating forecasts of the three components that we consider. In the results above, the aggregate inflation


¹⁴In the interest of readability we only present forecasts for this specific BVAR-SV model. However, other BVAR-SV models exhibit largely the same dynamics and patterns as this particular BVAR-SV model.

Figure 3. Time Series of Forecasts of (aggregate) Core PCE Inflation



Note: Dates on the x-axis represent the date of each forecast origin, t . Realized and forecasted values represent the n -step-ahead value from that given forecast origin.

Figure 4. Time Series of Forecasts of (aggregate) Core PCE Inflation, Focusing on 2020 Onward



Note: Dates on the x-axis represent the date of each forecast origin, t . Realized and forecasted values represent the n -step-ahead value from that given forecast origin. Realized inflation is represented as a quarter-on-quarter annualized rate through the transformation of $400 \cdot \log(\frac{x_t}{x_{t-1}})$, as are forecasts from the UC-SV, AM-BVAR-SV, and SVD-GP-NC models. SPF forecasts represent medians and are represented as the annualized quarterly percent change. SPF forecasts are aligned with model forecasts so that, for a given SPF release, the last available quarterly inflation date is equivalent to the date of our forecast origin, t .

forecast is directly taken from the models that include aggregate inflation. But it is possible that near-singularity arising from collinearity of inflation components with aggregate inflation could adversely affect forecast accuracy due to issues with numerical precision and with distortion in stochastic volatility computations. In our experience, we have found that some models (but less so with the baseline BVAR-SV specifications) that feature this collinearity tend to be vulnerable to explosive forecasts, especially as the model size grows and complexity is added. As a result, an alternative method to forecasting aggregate inflation is to instead drop aggregate inflation from the estimated model, forecast the disaggregate components, and then create an aggregate forecast by using the inflation weights of each component. This approach is commonly known as the “bottom-up” approach and has been widely applied to inflation forecasting (e.g., Hubrich 2005), although not widely with success (e.g., Faust and Wright 2013).

Accordingly, we consider forecasts of aggregate inflation in which we drop core inflation from each model and instead forecast the disaggregates and aggregate inflation as a weighted average of the component forecasts. We approximate the chain-weighted inflation weights by taking, for each pseudo-real-time period, the latest expenditure shares of each component. We then use forecasts of the disaggregates and the approximated weights to produce aggregate forecasts for core inflation. In our application we take just the latest expenditure shares and use them for all forecasting horizons, but alternatively one could forecast expenditure shares.¹⁵ However, this would come at the cost of introducing extra parameters with their associated estimation uncertainty. So we abstain from this procedure, since we find that expenditure shares tend to move little over our forecast horizons. In practice, we produce aggregate UC-SV forecasts by producing univariate UC-SV forecasts for each component of inflation and then aggregating those forecasts. For the BVAR-SV aggregate forecast, we can simply take the component forecasts from the vector y_{t+h} . Finally, for the SVD-WN model we estimate three

¹⁵An alternative approach to approximating chain weights is to follow the approach used in the replication files of Stock and Watson (2016), which calculates an average expenditure share weight for component i at time t as $w_t^i = \frac{e_t^i + e_{t-1}^i}{\sum_i e_t^i + \sum_i e_{t-1}^i}$, where e_t^i represents the total expenditure on component i at time t .

separate models, one for each component of inflation, and aggregate their respective component forecasts to create an aggregate forecast.

Table 6 compares the bottom-up approach to the direct approach to forecasting aggregate inflation used in earlier tables of results and summarized above. For each model and prior specification, we report the ratio of the RMSE or CRPS from the bottom-up approach to the RMSE or CRPS from the direct approach; entries below (above) 1 mean that the bottom-up approach beats (falls short of) the direct approach. These results fairly consistently show little advantage to a bottom-up approach to forecasting aggregate inflation. For the UC-SV and baseline AM-BVAR-SV specifications, the bottom-up forecasts are less accurate than the direct forecasts, most noticeably for UC-SV density forecasts (less so for UC-SV point forecasts). For other Bayesian prior and model feature choices, the bottom-up approach is sometimes slightly better than the direct, and sometimes slightly worse. For these models, there are no consistent gains to a bottom-up approach. However, at some horizons, the bottom-up approach does slightly improve forecasts from the SVD-GP-NC model. In the 2018–23 subsample, as in the longer subsample, there is no consistent benefit to the bottom-up approach. For example, as regards forecast accuracy from the UC-SV and baseline models, the bottom-up approach is relatively not as poor in the short subsample as in the long subsample, but it is not clearly better, either. On balance, there seems to be little reason to use a bottom-up approach, although in some settings, doing so has small positive impacts on forecast accuracy.

5. Looking Ahead: 2024:Q2–2026:Q4

The last section's results on historical forecast accuracy establish that the BVAR-SV model with an adaptive Minnesota-style prior is an effective tool for forecasting inflation in the core services excluding housing, housing, and core goods components of core PCE inflation, as well as aggregate core inflation. Accordingly, we use the AM-BVAR-SV model to assess the inflation outlook going forward.¹⁶ This BVAR specification is useful for capturing

¹⁶The BVAR-SV model with time-varying means (i.e., AM-BVAR-SV-TVM) generates similar forecasts for the period ahead; therefore, for the sake of brevity, we do not report them.

Table 6. Forecast Ratios for RMSEs and CRPSs When Using Bottom-Up Approach, Core PCE Inflation

Model/Horizon	h = 1		h = 2		h = 4		h = 8	
	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS
1985:Q1–2023:Q4								
UC-SV	1.018	1.070***	1.024	1.075***	1.025	1.092***	1.022	1.112***
NM-BVAR-SV	1.003	1.004	1.002	1.002	1.004	1.003	0.998	0.997
NM-BVAR-SVO	1.000	0.998	1.000	0.998	1.002	0.997	0.992	0.989
NM-BVAR-SV-TVM	1.015	1.013*	1.008	1.005	1.012	1.007	1.008	1.001
NM-BVAR-SVO-TVM	1.009	1.009	1.013	1.005	1.013	1.001	1.011	0.988
AM-BVAR-SV	1.004	1.008	1.010	1.017	1.017	1.020	1.010	1.005
AM-BVAR-SVO	1.003	1.003	1.010	1.012	1.015	1.015	1.005	1.001
AM-BVAR-SV-TVM	1.011	1.013	1.009	1.012	1.020	1.017	1.013	1.006
AM-BVAR-SVO-TVM	1.008	1.012	1.016	1.013	1.024	1.015	1.023	1.001
NGM-BVAR-SV	1.000	1.002	1.009	1.010	1.016	1.012	1.010	0.994
NGM-BVAR-SVO	0.998	0.998	1.006	1.003	1.015	1.003	1.005	0.984
NGM-BVAR-SV-TVM	1.020	1.015	1.018	1.014	1.025	1.013	1.016	0.996
NGM-BVAR-SVO-TVM	1.014	1.009	1.023	1.008	1.027	0.999	1.027	0.980
SVD-GP-NC	0.990	0.994	0.996	1.001	1.014	1.033	0.998	1.012
1985:Q1–2017:Q4								
UC-SV	1.012	1.088***	1.023*	1.108***	1.013	1.139***	1.008	1.175***
NM-BVAR-SV	1.008	1.006	1.011	1.007	1.015	1.005	1.015	1.000
NM-BVAR-SVO	0.999	0.997	1.007	1.000	1.014	0.999	1.008	0.991
NM-BVAR-SV-TVM	1.011	1.011	0.997	0.999	0.994	0.995	0.988	0.989
NM-BVAR-SVO-TVM	1.000	1.005	1.000	0.997	0.994	0.986	0.989	0.969*
AM-BVAR-SV	1.008	1.009	1.031	1.028*	1.041	1.029	1.018	1.005
AM-BVAR-SVO	1.003	1.003	1.028	1.020	1.042	1.023	1.012	1.001
AM-BVAR-SV-TVM	1.008	1.012	1.010	1.011	1.013	1.010	0.998	0.996
AM-BVAR-SVO-TVM	1.003	1.009	1.015	1.011	1.022	1.007	1.015	0.985
NGM-BVAR-SV	0.999	1.001	1.024	1.016	1.038	1.016	1.015	0.987
NGM-BVAR-SVO	0.994	0.996	1.019	1.007	1.040	1.004	1.006	0.976*
NGM-BVAR-SV-TVM	1.002	1.003	1.013	1.008	1.007	0.997	0.991	0.978
NGM-BVAR-SVO-TVM	0.990	0.995	1.018	1.000	1.013	0.980	1.008	0.952**
SVD-GP-NC	0.981*	0.986	1.001	1.001	1.061*	1.056*	1.036	1.031

(continued)

Table 6. (Continued)

Model/Horizon	h = 1		h = 2		h = 4		h = 8	
	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS	RMSE	CRPS
2018:Q1–2023:Q4								
UC-SV	1.024	1.024	1.024	1.002	1.032	1.003	1.031*	0.996
NM-BVAR-SV	0.997	0.998	0.990	0.990	0.995	0.999	0.986	0.988
NM-BVAR-SVO	1.002	1.000	0.990	0.992	0.991	0.993	0.981*	0.984*
NM-BVAR-SV-TVM	1.020	1.018	1.023	1.022	1.029	1.038	1.022	1.029
NM-BVAR-SVO-TVM	1.020	1.018	1.029	1.030	1.031	1.038	1.028**	1.037*
AM-BVAR-SV	1.000	1.002	0.984*	0.986	0.996	0.999	1.004	1.006**
AM-BVAR-SVO	1.004	1.004	0.986	0.990	0.993	0.994	1.001	1.002
AM-BVAR-SV-TVM	1.015	1.015	1.009	1.013	1.026	1.035	1.022***	1.029**
AM-BVAR-SVO-TVM	1.016	1.019	1.016	1.020	1.026	1.035	1.029***	1.040***
NGM-BVAR-SV	1.002	1.005	0.989	0.991	0.998	1.004	1.006	1.011***
NGM-BVAR-SVO	1.005	1.006	0.990	0.994	0.995	0.998	1.004	1.005
NGM-BVAR-SV-TVM	1.044	1.050	1.024	1.030	1.038	1.049	1.032**	1.036***
NGM-BVAR-SVO-TVM	1.045	1.049	1.028	1.033	1.038	1.047	1.038***	1.046***
SVD-GP-NC	1.003	1.018	0.990	0.998	0.974	0.984	0.967*	0.974

Note: Values below 1 indicate improvement over the corresponding core PCE inflation forecast that is taken directly from the model that includes aggregate inflation. The UC-SV model in the table represents the aggregate forecast from three separate disaggregate UC-SV models. *, **, and *** represent the .10, .05, and 0.1 significance levels, respectively. Entries in bold denote ratios below 1 that are the lowest for a given forecasting horizon and measure. In columns with no bold entries, no bottom-up approach beats its direct counterpart.

the historical comovement of inflation, its components, and wage growth and the resulting implications for the inflation outlook over the next few years. We estimate the model using the latest available data, i.e., 1959:Q2 through 2024:Q1, and use the estimated model to forecast inflation, its components, and ECI growth from 2024:Q2 through 2026:Q4.

This section begins by presenting baseline unconditional forecasts and then examines the inflation outlook under some alternative scenarios, one conditioning on a path of wage growth returning to 3 percent and the other conditioning on a path of aggregate core inflation returning to 2 percent.

5.1 *Baseline Forecasts*

Following standard practice in the literature to incorporate high-frequency intraquarterly information, we entropically tilt (condition) the current quarter, i.e., 2024:Q2, unconditional forecast of core PCE inflation toward the Q2 nowcast estimate (of core PCE inflation, not the other variables in the model) taken from the Federal Reserve Bank of Cleveland's (Cleveland Fed's) website.¹⁷ We refer to the resulting conditional forecasts as the baseline forecasts and characterize them as unconditional forecasts because the only condition is one that simply accounts for the most recent information on aggregate core PCE inflation. Table 7 reports these forecasts (point forecasts computed as posterior means) under the column heading labeled "Base." The top panel in the table reports quarterly annualized growth rates, and the bottom panel reports the corresponding forecasts expressed in four-quarter growth rates (i.e., averages of annualized quarterly log growth rates over four quarters, which in the fourth quarter of each year represent Q4/Q4 inflation rates for the year).

The table shows that the baseline forecast has core PCE inflation slowly decelerating to a quarterly annualized rate of 2.6 percent by

¹⁷As of May 1, 2024, the Q2 nowcast of core PCE inflation is estimated to be 3.07 percent (quarterly annualized rate), which is nearly identical to the unconditional forecast of 3.11 percent. Cleveland Fed inflation nowcasts are based on the model developed in Knotek and Zaman (2017). Knotek and Zaman (2023) document the competitive accuracy of Cleveland Fed nowcasts compared with nowcasts from professional forecasters and other competing approaches.

Table 7. Looking Ahead, 2024:Q2–2026:Q4: AM-BVAR-SV

Date	Nominal Wage		Supercore			Housing			Goods			Core PCE			
	Base	Wage	Core Sc.	Base	Wage	Core Sc.	Base	Wage	Core Sc.	Base	Wage	Core Sc.	Base	Wage	Core Sc.
		Sc.			Sc.			Sc.			Sc.			Sc.	
Quarterly Annualized (QoQ) Forecasts															
2024:Q2	4.0	3.9	3.9	3.8	3.8	3.9	5.2	-0.3	-0.3	3.1	3.1	3.1	3.1	3.1	3.1
2024:Q3	3.8	3.8	3.8	3.6	3.5	3.5	5.0	-0.1	-0.4	2.9	2.8	2.9	2.8	2.9	2.9
2024:Q4	3.8	3.7	3.8	3.6	3.3	3.4	4.9	0.0	-0.4	2.9	2.7	2.9	2.7	2.8	2.8
2025:Q1	3.8	3.6	3.8	3.5	3.1	3.3	4.7	0.2	-0.4	2.9	2.5	2.9	2.5	2.7	2.7
2025:Q2	3.7	3.5	3.6	3.4	3.0	3.1	4.6	0.2	-0.5	2.8	2.4	2.8	2.4	2.5	2.5
2025:Q3	3.7	3.4	3.5	3.4	2.8	3.0	4.5	0.3	-0.6	2.8	2.2	2.8	2.2	2.4	2.4
2025:Q4	3.6	3.2	3.4	3.3	2.7	2.8	4.4	0.2	-0.6	2.8	2.1	2.8	2.1	2.2	2.2
2026:Q1	3.6	3.0	3.3	3.3	2.6	2.6	4.3	0.3	-0.5	2.7	2.1	2.7	2.1	2.0	2.0
2026:Q2	3.6	3.0	3.2	3.3	2.6	2.6	4.2	0.4	-0.4	2.7	2.1	2.7	2.1	2.0	2.0
2026:Q3	3.5	3.0	3.2	3.2	2.7	2.6	4.1	0.3	-0.3	2.7	2.1	2.7	2.1	2.0	2.0
2026:Q4	3.5	3.0	3.1	3.2	2.7	2.6	4.1	0.4	-0.2	2.6	2.2	2.6	2.2	2.0	2.0
Corresponding Four-Quarter Trailing Rate (4Q) Forecasts															
2024:Q2	4.0	4.0	4.0	3.6	3.6	3.6	5.5	-1.1	-1.1	2.7	2.7	2.7	2.7	2.7	2.7
2024:Q3	4.0	4.0	4.0	3.7	3.7	3.7	5.3	-0.6	-0.7	2.9	2.9	2.9	2.9	2.9	2.9
2024:Q4	4.0	4.0	4.0	4.0	3.9	3.9	5.2	-0.2	-0.4	3.1	3.0	3.1	3.0	3.1	3.1
2025:Q1	3.8	3.8	3.8	3.6	3.4	3.5	5.0	-0.1	-0.4	3.0	2.8	3.0	2.8	2.8	2.8
2025:Q2	3.8	3.6	3.7	3.5	3.2	3.3	4.8	0.1	-0.4	2.9	2.6	2.9	2.6	2.7	2.7
2025:Q3	3.7	3.5	3.7	3.5	3.1	3.2	4.7	0.2	-0.5	2.9	2.5	2.9	2.5	2.6	2.6
2025:Q4	3.7	3.4	3.6	3.4	2.9	3.1	4.5	0.3	-0.5	2.8	2.3	2.8	2.3	2.5	2.5
2026:Q1	3.6	3.3	3.5	3.3	2.8	2.9	4.4	0.3	-0.5	2.8	2.2	2.8	2.2	2.3	2.3
2026:Q2	3.6	3.2	3.4	3.3	2.7	2.8	4.3	0.2	-0.5	2.7	2.1	2.7	2.1	2.2	2.2
2026:Q3	3.6	3.1	3.3	3.3	2.6	2.7	4.2	0.3	-0.5	2.7	2.1	2.7	2.1	2.1	2.1
2026:Q4	3.5	3.0	3.2	3.2	2.7	2.6	4.2	0.4	-0.4	2.7	2.1	2.7	2.1	2.0	2.0
Note: Numbers reported in the top (bottom) panel are the quarterly annualized (four-quarter trailing rate) forecasts from the AM-BVAR-SV benchmark model for the forecast period spanning 2024:Q2 through 2026:Q4. The columns denoted “Base” refer to baseline forecasts, “Wage Sc.” refers to forecasts generated under the wage scenario, and “Core Sc.” refers to forecasts generated under the core PCE 2 percent scenario. In the top panel, the numbers in bold indicate conditions imposed from outside the model: 2024:Q2 condition (of 3.1 percent) on core PCE; 2026:Q1–2026:Q4 conditions (of 3.0 percent) on nominal wages for wage scenario; 2026:Q1–2026:Q4 conditions (of 2.0 percent) on core PCE for core PCE scenario.															

the end of 2026. The slow decline of Q4/Q4 core inflation from 3.1 percent in 2024 to 2.7 percent in 2026 is associated with a projected gradual moderation in the growth rate of nominal wages, from the current annualized quarterly rate of 4.0 percent in 2024:Q2 to 3.6 percent in 2026:Q1 (and, on Q4/Q4 basis, from 4.0 percent in 2024 to 3.5 percent in 2026).

The baseline forecast projects moderation in both supercore and housing components of inflation, but at a differing pace. Supercore inflation is projected to decline from an annualized quarterly rate of 3.8 percent in 2024:Q2 to 3.2 percent in 2026:Q4, and housing inflation slows from 5.2 percent to 4.1 percent. In contrast, core goods inflation is projected to increase from -0.3 percent to an annualized rate of 0.4 percent at the end of 2026. On a four-quarter basis, smoothing through quarterly fluctuations, both housing and supercore components are forecast to decline from 2024:Q2 to 2026:Q4, most sharply in the case of housing inflation (from 5.5 to 4.2 percent) and significantly in the case of supercore inflation (from 3.6 to 3.2 percent), whereas goods inflation is forecast to increase (from -1.1 to 0.4 percent).

The forecasted Q4/Q4 inflation rates in 2026 are somewhat higher than their pre-COVID (2017–19) averages of 3.3 percent (housing), 2.1 percent (supercore), and -0.6 percent (core goods). Hence, not surprisingly, in 2026 aggregate core PCE inflation is 70 basis points above the Federal Reserve's longer-run inflation target of 2 percent and 100 basis points above its 2017–19 average rate. Note that the solidly negative prepandemic rate of goods inflation reflects a pattern that emerged in 1996. In an accounting sense, sustained modest deflation in goods prices contributed to keeping aggregate core inflation low from the mid-1990s up until the COVID-19 pandemic. However, in recent years, some public commentary has raised some concern that increased global fragmentation in trade and deglobalization could contribute to higher goods inflation and in turn slow progress toward returning aggregate inflation to the Federal Reserve's 2 percent goal. Results below will shed some additional light on the possible importance of returning to below-zero goods inflation.


While the baseline forecast just described refers to the means of predictive distributions, there is of course considerable uncertainty around the outlook captured in the model's (posterior) forecast

distributions. We turn now to considering some probabilistic assessments of the outlook using these distributions. Some estimates consider particular cutoff values for inflation, its components, and wage growth. We use cutoffs based on prepandemic (2017–19 average) rates of inflation and wage growth, with some adjustment to reflect the fact that, on average from 2017 through 2019, core PCE inflation was 30 basis points below the Federal Open Market Committee’s (FOMC’s) longer-run inflation target of 2 percent. For cutoffs used in probabilistic assessments, we take a simple approach of using conditions or scenarios that mark up each component and wage growth by a bit more or a bit less than 30 basis points, consistent with the shortfall in average core PCE inflation for 2017–19. The bit more or bit less is chosen to make for nice round numbers in the conditions. (Weighted by spending shares, the component cutoffs correspond to aggregate core inflation of 2 percent.) For simplicity, below we will refer to “prepandemic averages” of inflation, its components, and wage growth even though the values used in these calculations are by design a little higher than the actual 2017–19 average values.

Panel A of Figure 5 plots the marginal probability of core PCE inflation (the four-quarter average rate) being less than or equal to 2 percent over the forecast horizon. Panels B, C, and D provide the marginal probabilities of the components’ inflation rates being less than their respective prepandemic averages (adjusted as noted above). Based on the baseline forecast density, the probability that core PCE inflation averages 2 percent or less is below 35 percent through 2026:Q4. So while the baseline forecast has a mean core inflation rate well above 2 percent in 2026, the (35 percent) probability of reaching 2 percent is not small. As can be seen, contributing factors to this moderate probability are comparable probabilities of negative core goods inflation (probability below 50 percent in most quarters), supercore inflation below 2.5 percent (probability less than 35 percent), and housing inflation below 3.5 percent (probability less than 35 percent).


Figure 6 plots *joint* probabilities of observing both core PCE inflation ≤ 2 percent and component inflation rates above or below the cutoff values specified as adjustments of 2017–19 averages. The left panels plot the joint probabilities that are based on the components’ inflation rates below the respective thresholds, and the

Figure 5. Marginal Probabilities
Based on Baseline Forecasts



Note: The plots are marginal probabilities of core PCE inflation and its components constructed from their baseline density forecasts (four-quarter growth rates). Panel A plots the probability of core PCE inflation at or below 2 percent, panel B probability that goods inflation is negative, panel C probability of housing inflation below 3.5 percent, and panel D probability of supercore inflation less than 2.5 percent.

Figure 6. Joint Probabilities Based on Baseline Forecasts




Note: The plots are joint probabilities of core PCE inflation and its components constructed from their baseline density forecasts (four-quarter growth rates).

right panels at or above the thresholds (note that the right panels use an axis scaling different from the left panels). For both the core goods and supercore components, the joint probability of seeing core PCE inflation ≤ 2 percent and the component's inflation rate below the respective threshold is significantly higher (by more than 20 percentage points) than the joint probability of core inflation at or below 2 percent and the component's inflation rate above its respective threshold. In the case of housing, the differences in the joint probabilities of housing above or below the threshold and seeing core inflation at or below 2 percent are relatively modest. For example, in 2026:Q4, the joint probability of PCE inflation at or below 2 percent and housing inflation below 3.5 percent is a little below 20 percent, whereas the corresponding probability with housing inflation at or above 3.5 percent is nearly 16 percent, little different.

Figure 7 reports *conditional* probabilities of core PCE inflation at or below 2 percent, given component inflation rates either below or above the cutoffs representing rough measures of prepandemic levels. The left panels plot the conditional probabilities of core PCE given inflation in the components below the indicated threshold, and the right panels (with axes scaled differently than the left panels) provide results for component inflation rates at or above the threshold. The results in the left panels indicate that, conditional on seeing inflation in the components at levels below their adjusted prepandemic levels, the probabilities of core PCE inflation at or below 2 percent are significantly higher than 50 percent, especially for the goods and supercore components. For example, in 2026:Q4, conditional on core goods inflation below zero, the probability of core PCE inflation at or below 2 percent is 68 percent, and conditional on supercore inflation below 3.5 percent, the probability of core PCE inflation at or below 2 percent is 82 percent. The probability of PCE inflation of 2 percent or less is still high conditioning on housing inflation returning to its prepandemic level, but at 58 percent in 2026:Q4 it is not as high as when conditioning on prepandemic levels of other components. Conversely, conditional probabilities of core PCE inflation at or below 2 percent in 2026 are low if goods inflation is persistently positive, housing inflation stays above 3.5 percent, or supercore inflation is above 2.5 percent, respectively. These probabilities are especially low—between 6 and 10 percent—in the cases of

Figure 7. Conditional Probabilities
Based on Baseline Forecasts



Note: The plots are conditional probabilities of core PCE inflation, conditional on the components inflation rates above or below certain thresholds. The probabilities are constructed from their baseline density forecasts (four-quarter growth rates).

goods inflation remaining positive and supercore inflation remaining at or above 2.5 percent, compared with a probability of 23 percent with housing inflation remaining at or above 3.5 percent.

The fact that housing inflation getting back to a prepandemic level is projected to matter for returning inflation to 2 percent but somewhat less sharply than estimated for the other two inflation components is likely tied to the respective weights of the components in aggregate core PCE inflation. Supercore spending is by far the largest component (of these three components) of consumption, with a nominal spending share (in 2024:Q1) of 57 percent. Goods are second in importance, with a spending share of 25 percent, and housing is third, at a spending share of 18 percent.

5.2 *Conditional Forecast: Wage Scenario*

The baseline forecast has nominal wage growth moderating to an annualized rate of 3.6 percent in 2025 and 3.5 percent on a Q4/Q4 basis in 2026. This pace of wage growth is about 0.5 percentage point higher than 3 percent, the pace we have taken—in view of prepandemic averages—to be roughly consistent with core PCE inflation of 2 percent. There are reasons to believe that the projected slowing in the economy in response to restrictive monetary policy may lead to quicker wage moderation than indicated in the baseline forecast. DeLuca and Van Zandweghe (2023), among others, project nominal wage growth to decline to 3 percent by 2025. If that were to occur, would the slower wage growth be associated with a faster moderation in core inflation than implied in the baseline? To answer this question, we build a scenario using our model, assuming that by the start of 2026 and beyond, nominal wage growth averages 3 percent. The resulting conditional forecasts are shown in Table 7 under the column labeled “Wage Sc.”

Conditioning the model with a nominal wage path that calls for wage growth to decline to 3 percent by early 2026 is associated with a quicker decline in core PCE inflation than implied by the baseline. The forecast has core PCE inflation moving down to 2.1 percent in 2026 (slightly higher than the median forecast of 2.0 percent in the Federal Reserve’s Summary of Economic Projections released following the March 2024 FOMC meeting). The forecasts of the components reveal that the steeper fall in aggregate core PCE

inflation in the scenario as compared with the baseline projection is associated with declines in the inflation rates of all three components that are also steeper in the scenario than in the baseline. This scenario also reveals that of the three components, supercore and core goods are relatively more sensitive to wage developments than is housing inflation: In 2026, the Q4/Q4 inflation rates for supercore and goods are 50–80 basis points lower in the wage scenario than in the baseline forecast, whereas the inflation rate for housing is just 30 basis points lower in the scenario than in the baseline.

5.3 Conditional Forecast: Core PCE 2 Percent Scenario

The previous scenario illustrated that, according to our benchmark AM-BVAR-SV model, a moderation in nominal wage growth to 3 percent is consistent with returning core inflation closer to 2 percent, but not quite with getting core inflation there. Alternatively, we can assess what paths of inflation components and wage growth will be consistent with core inflation reaching 2 percent in 2026 by explicitly conditioning on such a lower path of core PCE inflation. The resulting conditional forecasts are reported in Table 7 under the column labeled “Core Sc.”

Under this scenario, our baseline BVAR projects that moderation of core PCE inflation to 2 percent in 2026 would be associated with more slowing of inflation in the supercore and goods components than is projected under the baseline case and the wage growth scenario, more so when compared with the baseline forecast than the wage growth scenario. On a Q4/Q4 basis, goods (supercore) inflation is projected to slow to -0.6 (2.6) percent in the 2 percent inflation scenario, compared with -0.4 (2.7) percent in the wage scenario and 0.4 (3.2) percent in the baseline. Housing inflation and wage growth show relatively less change in this scenario as compared with the wage scenario and baseline forecast. The Q4/Q4 housing inflation forecast for 2026 is 3.9 percent in the 2 percent inflation scenario, the same as in the wage scenario in which core PCE inflation remains a little above 2 percent, and 30 basis points below the baseline. The Q4/Q4 forecast of wage growth in this scenario is very close to the 3 percent pace considered in the wage scenario.

Together, the results of the slower wage growth and 2 percent core PCE inflation scenarios may be seen as pointing to the

likelihood that, for core PCE inflation to near or reach the FOMC's longer-run target of 2 percent over the next few years, the three components of core inflation and wage growth will need to slow to rates close to, although a little above, their prepandemic averages.

6. Conclusions

Motivated by the recent emphasis that some economic commentators and policymakers have placed on a three-component breakdown of core inflation into goods excluding food and energy, services excluding energy and housing, and housing, this paper examines the efficacy and implications of forecasts of the three-component breakdown of core inflation in the U.S.

We first examine a range of models and establish the historical efficacy of a BVAR with stochastic volatility in five variables, consisting of core PCE inflation, its three components, and wage growth. An examination of the historical accuracy of point and density forecasts indicate that this model estimated with an adaptive Minnesota-style prior is an effective tool for forecasting inflation in the core services excluding housing, housing, and core goods components of core PCE inflation, as well as aggregate core inflation (although less so for aggregate inflation than its components). Adding features such as time-varying parameters or more variables either fails to improve or harms forecast accuracy. Our historical results also indicate that aggregate inflation is best forecast directly; no consistent benefit can be achieved by aggregating disaggregate components.

We also examine medium-term forecasts for the period ahead, 2024 through 2026. In particular, we consider the likelihood of inflation returning to the Federal Reserve's 2 percent target under our model's baseline forecast distribution. Our model's baseline (mean) projection puts core inflation at 2.7 percent in 2026, well below its level in 2023 but still elevated relative to the Federal Reserve's longer-run target of 2 percent. But the model's forecast puts a 35 percent probability on hitting 2 percent core inflation in 2026. In these baseline predictive distributions, a return of core PCE inflation to 2 percent or less over the next few years is much higher when conditioning on goods or supercore inflation slowing to near prepandemic levels than when conditioning on these components remaining above the same thresholds.

We also analyze alternative scenarios. One features a gradual decline of wage growth to near its prepandemic level to assess the likelihood that such a development may be key to continued disinflation. The second scenario conditions on a decline of aggregate core inflation to 2 percent to provide a joint assessment of the paths of component inflation rates and wage growth most likely associated with a return of aggregate inflation to target. These scenario analyses indicate that slower wage growth will likely be associated with reduced inflation in all three components, not just supercore, but housing less so than goods and supercore inflation. In addition, based on the historical dynamics of inflation and its components as captured by the BVAR, it is most likely the case that a return of core inflation to the Federal Reserve's 2 percent target will be associated with all three components returning to near prepandemic levels.

References

- Antolín-Díaz, J., I. Petrella, and J. F. Rubio-Ramirez. 2021. "Structural Scenario Analysis with SVARs." *Journal of Monetary Economics* 117 (January): 798–815. <https://doi.org/10.1016/j.jmoneco.2020.06.001>.
- Arias, J. E., J. F. Rubio-Ramirez, and M. Shin. 2023. "Macroeconomic Forecasting and Variable Ordering in Multivariate Stochastic Volatility Models." *Journal of Econometrics* 235 (2): 1054–86. <https://doi.org/10.1016/j.jeconom.2022.04.013>.
- Banbura, M., and A. van Vlodrop. 2018. "Forecasting with Bayesian Vector Autoregressions with Time Variation in the Mean." Discussion Paper No. 2018-025/IV, Tinbergen Institute. <https://doi.org/10.2139/ssrn.3145055>.
- Bognanni, M. 2018. "A Class of Time-Varying Parameter Structural VARs for Inference under Exact or Set Identification." Working Paper No. 18-11, Federal Reserve Bank of Cleveland. <https://doi.org/10.26509/frbc-wp-201811>.
- . 2022. "Comment on 'Large Bayesian Vector Autoregressions with Stochastic Volatility and Non-conjugate Priors'." *Journal of Econometrics* 227 (2): 498–505. <https://doi.org/10.1016/j.jeconom.2021.10.008>.
- Carriero, A., J. Chan, T. E. Clark, and M. Marcellino. 2022. "Corrigendum to 'Large Bayesian Vector Autoregressions

- with Stochastic Volatility and Non-conjugate Priors'." *Journal of Econometrics* 227 (2): 506–12. <https://doi.org/10.1016/j.jeconom.2021.11.010>.
- Carriero, A., T. E. Clark, and M. Marcellino. 2019. "Large Bayesian Vector Autoregressions with Stochastic Volatility and Non-conjugate Priors." *Journal of Econometrics* 212 (1): 137–54. <https://doi.org/10.1016/j.jeconom.2019.04.024>.
- Carriero, A., T. E. Clark, M. Marcellino, and E. Mertens. 2024. "Addressing COVID-19 Outliers in BVARs with Stochastic Volatility." *Review of Economics and Statistics* 106 (5): 1403–17. https://doi.org/10.1162/rest_a_01213.
- Chan, J. C. C. 2017. "Notes on Bayesian Macroeconometrics." Version 1.4. https://joshuachan.org/notes_BayesMacro.html.
- . 2021. "Minnesota-Type Adaptive Hierarchical Priors for Large Bayesian VARs." *International Journal of Forecasting* 37 (3): 1212–26. <https://doi.org/10.1016/j.ijforecast.2021.01.002>.
- . 2023. "Large Hybrid Time-Varying Parameter VARs." *Journal of Business and Economic Statistics* 41 (3): 890–905. <https://doi.org/10.1080/07350015.2022.2080683>.
- Chan, J. C. C., and E. Eisenstat. 2018. "Bayesian Model Comparison for Time-Varying Parameter VARs with Stochastic Volatility." *Journal of Applied Econometrics* 33 (4): 509–32. <https://doi.org/10.1002/jae.2617>.
- Chan, J. C. C., and I. Jeliazkov. 2009. "Efficient Simulation and Integrated Likelihood Estimation in State Space Models." *International Journal of Mathematical Modelling and Numerical Optimisation* 1 (1/2): 101–20. <https://doi.org/10.1504/IJMMNO.2009.030090>.
- Clark, T. E. 2011. "Real-Time Density Forecasts from Bayesian Vector Autoregressions with Stochastic Volatility." *Journal of Business and Economic Statistics* 29 (3): 327–41. <https://doi.org/10.1198/jbes.2010.09248>.
- Clark, T. E., and F. Ravazzolo. 2015. "Macroeconomic Forecasting Performance under Alternative Specifications of Time-Varying Volatility." *Journal of Applied Econometrics* 30 (4): 551–75. <https://doi.org/10.1002/jae.2379>.
- Cogley, T., S. Morozov, and T. J. Sargent. 2005. "Bayesian Fan Charts for U.K. Inflation: Forecasting and Sources

- of Uncertainty in an Evolving Monetary System.” *Journal of Economic Dynamics and Control* 29 (11): 1893–1925. <https://doi.org/10.1016/j.jedc.2005.06.005>.
- D’Agostino, A., L. Gambetti, and D. Giannone. 2013. “Macroeconomic Forecasting and Structural Change.” *Journal of Applied Econometrics* 28 (1): 82–101. <https://doi.org/10.1002/jae.1257>.
- DeLuca, M., and W. Van Zandweghe. 2023. “Postpandemic Nominal Wage Growth: Inflation Pass-Through or Labor Market Imbalance?” Economic Commentary No. 2023-13, Federal Reserve Bank of Cleveland. <https://doi.org/10.26509/frbc-ec-202313>.
- Diebold, F. X., and R. S. Mariano. 1995. “Comparing Predictive Accuracy.” *Journal of Business and Economic Statistics* 13 (3): 253–63. <https://doi.org/10.2307/1392185>.
- Eo, Y., L. Uzeda, and B. Wong. 2023. “Understanding Trend Inflation through the Lens of the Goods and Services Sectors.” *Journal of Applied Econometrics* 38 (5): 751–66. <https://doi.org/10.1002/jae.2975>.
- Faust, J., and J. H. Wright. 2013. “Forecasting Inflation.” In *Handbook of Economic Forecasting*, Vol. 2, Part A, ed. G. Elliott and A. Timmermann, 2–56 (chapter 1). Elsevier. <https://doi.org/10.1016/B978-0-444-53683-9.00001-3>.
- Hauzenberger, N., F. Huber, G. Koop, and L. Onorante. 2022. “Fast and Flexible Bayesian Inference in Time-Varying Parameter Regression Models.” *Journal of Business and Economic Statistics* 40 (4): 1904–18. <https://doi.org/10.1080/07350015.2021.1990772>.
- Hubrich, K. 2005. “Forecasting Euro Area Inflation: Does Aggregating Forecasts by HICP Component Improve Forecast Accuracy?” *International Journal of Forecasting* 21 (1): 119–36. <https://doi.org/10.1016/j.ijforecast.2004.04.005>.
- Joseph, A., G. Potjagailo, E. Kalamara, C. Chakraborty, and G. Kapetanios. 2022. “Forecasting UK Inflation Bottom Up.” Staff Working Paper No. 915, Bank of England.
- Karlsson, S. 2013. “Forecasting with Bayesian Vector Autoregression.” In *Handbook of Economic Forecasting*, Vol. 2, Part B, ed. G. Elliott and A. Timmermann, 791–897 (chapter 15). Elsevier. <https://doi.org/10.1016/B978-0-444-62731-5.00015-4>.

- Kim, S., N. Shephard, and S. Chib. 1998. "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models." *Review of Economic Studies* 65 (3): 361–93. <https://doi.org/10.1111/1467-937X.00050>.
- Knotek, E. S., and S. Zaman. 2017. "Nowcasting U.S. Headline and Core Inflation." *Journal of Money, Credit and Banking* 49 (5): 931–68. <https://doi.org/10.1111/jmcb.12401>.
- . 2019. "Financial Nowcasts and Their Usefulness in Macroeconomic Forecasting." *International Journal of Forecasting* 35 (4): 1708–24. <https://doi.org/10.1016/j.ijforecast.2018.10.012>.
- . 2023. "A Real-Time Assessment of Inflation Nowcasting at the Cleveland Fed." Economic Commentary No. 2023-06, Federal Reserve Bank of Cleveland. <https://doi.org/10.26509/frbc-ec-202306>.
- Krüger, F., T. E. Clark, and F. Ravazzolo. 2017. "Using Entropic Tilting to Combine BVAR Forecasts with External Nowcasts." *Journal of Business and Economic Statistics* 35 (3): 470–85. <https://doi.org/10.1080/07350015.2015.1087856>.
- Krüger, F., S. Lerch, T. L. Thorarinsdottir, and T. Gneiting. 2021. "Predictive Inference Based on Markov Chain Monte Carlo Output." *International Statistical Review* 89 (2): 274–301. <https://doi.org/10.1111/insr.12405>.
- McGregor, T., and F. Toscani. 2022. "A Bottom-Up Reduced Form Phillips Curve for the Euro Area." IMF Working Paper No. 22/260. <https://doi.org/10.5089/9798400219108.001>.
- Powell, J. H. 2022. "Inflation and the Labor Market." Speech at the Hutchins Center on Fiscal and Monetary Policy, Brookings Institution, Washington, DC, November 30.
- Primiceri, G. E. 2005. "Time Varying Structural Vector Autoregressions and Monetary Policy." *Review of Economic Studies* 72 (3): 821–52. <https://doi.org/10.1111/j.1467-937X.2005.00353.x>.
- Ravazzolo, F., and S. P. Vahey. 2014. "Forecast Densities for Economic Aggregates from Disaggregate Ensembles." *Studies in Nonlinear Dynamics and Econometrics* 18 (4): 367–81. <https://doi.org/10.1515/snde-2012-0088>.
- Robertson, J. C., E. W. Tallman, and C. H. Whiteman. 2005. "Forecasting Using Relative Entropy." *Journal of Money, Credit and Banking* 37 (3): 383–401. <https://doi.org/10.1353/mcb.2005.0034>.

- Stock, J. H., and M. W. Watson. 2002. "Macroeconomic Forecasting Using Diffusion Indexes." *Journal of Business and Economic Statistics* 20 (2): 147–62. <https://doi.org/10.1198/073500102317351921>.
- . 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39 (s1): 3–33. <https://doi.org/10.1111/j.1538-4616.2007.00014.x>.
- . 2016. "Core Inflation and Trend Inflation." *Review of Economics and Statistics* 98 (4): 770–84. https://doi.org/10.1162/REST_a_00608.
- Tallman, E. W., and S. Zaman. 2017. "Forecasting Inflation: Phillips Curve Effects on Services Price Measures." *International Journal of Forecasting* 33 (2): 442–57. <https://doi.org/10.1016/j.ijforecast.2016.10.004>.
- . 2020. "Combining Survey Long-Run Forecasts and Nowcasts with BVAR Forecasts Using Relative Entropy." *International Journal of Forecasting* 36 (2): 373–98. <https://doi.org/10.1016/j.ijforecast.2019.04.024>.
- Verbrugge, R., and S. Zaman. 2024. "Post-COVID Inflation Dynamics: Higher for Longer." *Journal of Forecasting* 43 (4): 871–93. <https://doi.org/10.1002/for.3070>.
- Waggoner, D. F., and T. Zha. 1999. "Conditional Forecasts in Dynamic Multivariate Models." *Review of Economics and Statistics* 81 (4): 639–51. <https://doi.org/10.1162/003465399558508>.
- West, K. D. 1996. "Asymptotic Inference about Predictive Ability." *Econometrica* 64 (5): 1067–84. <https://doi.org/10.2307/2171956>.
- Zellner, A. 1986. "On Assessing Prior Distributions and Bayesian Regression Analysis with g-Prior Distributions." In *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti*, ed. P. Goel and A. Zellner, 223–43. Elsevier.