Bank of Finland Research Discussion Papers
1 • 2026

Fabio Verona

Forecasting inflation: The sum of the cycles outperforms the whole

Bank of Finland
Research

Bank of Finland Research Discussion Papers
Editor-in-Chief Esa Jokivuolle

Bank of Finland Research Discussion Papers 1/2026
7 January 2026

Fabio Verona
Forecasting inflation: The sum of the cycles outperforms the whole

ISSN 1456-6184, online

Bank of Finland
Research Unit

PO Box 160
FIN-00101 Helsinki

Phone: +358 9 1831

Email: research@bof.fi
Website: www.suomenpankki.fi/en/research/research-unit/

Forecasting inflation:

The sum of the cycles outperforms the whole*

Fabio Verona*

Abstract

Inflation dynamics reflect forces operating at different cycles, from short-lived shocks to long-term structural trends. We introduce the sum-of-the-cycles (SOC) method, which exploits this multi-frequency structure of inflation for forecasting. SOC decomposes inflation into cyclical components, applies forecasting models suited to their persistence, and recombines them into an aggregate forecast. Across U.S. inflation measures and horizons, SOC consistently outperforms leading time-series benchmarks, reducing forecast errors by about 25 percent at short horizons and nearly 50 percent at long horizons. During the 2020-21 inflation surge, when many models – including advanced machine-learning methods – struggled, SOC retained strong performance by incorporating shortage indicators. Beyond accuracy, SOC enhances interpretability: financial variables dominate high- and business-cycle frequencies, Phillips Curve models are most informative at medium frequencies, and factor-based methods, forecast combinations, and shortage indices prevail at low frequencies. This combination of accuracy and transparency makes SOC a practical complement to existing tools for inflation forecasting and policy analysis.

Keywords: inflation forecasting; frequency decomposition; cycles; forecast combination; shortage indicators; Phillips curve; macro-finance.

JEL codes: C22, C53, E31, E32, E37

* I thank Francesco Ravazzolo for useful comments, and Paweł Skrzypczyński for providing the time series of the jobs-workers gap. The views expressed are mine and do not necessarily reflect those of the Bank of Finland.

* Bank of Finland – Monetary Policy and Research Department, and University of Porto – *Cef.up* (fabio.verona@bof.fi)

1 Introduction

Inflation dynamics are driven by multiple forces operating at different frequencies. Short-run fluctuations arise from transitory shocks such as energy prices or exchange rates, while medium- and long-run movements reflect more persistent drivers, including labor-market slack, monetary and fiscal policy, and demographic trends. This perspective – captured in Yellen's (2016) discussion of inflation and later termed the "Fed's view" by Hasenzagl, Pellegrino, Reichlin and Ricco (2022) – aligns closely with the New Keynesian Phillips Curve (à la Coibion and Gorodnichenko, 2015), in which expectations anchor long-run inflation, resource utilization drives business-cycle dynamics, and cost-push shocks generate short-term volatility.

Traditional forecasting models, beginning with Gordon (1970), compress this heterogeneous structure into a single aggregate process. This helps explain why simple benchmarks such as the Atkeson-Ohanian (2001) random-walk model remain difficult to outperform – a result repeatedly documented in the inflation-forecasting literature (e.g., Faust and Wright, 2013 and Rossi, 2025). Even the New Keynesian Phillips Curve, while conceptually consistent with a multi-frequency view, has delivered limited forecasting gains in practice. What remains missing is a method that forecasts inflation by explicitly exploiting frequency-specific dynamics in a systematic and transparent way.

To fill this gap, we develop the sum-of-the-cycles (SOC) method. The key idea is straightforward: rather than forecasting inflation directly, the series is decomposed into frequency components. Each component is forecast with the model best suited to its behavior, and the resulting forecasts are recombined into an aggregate inflation forecast. The novelty lies not only in decomposing inflation into frequency components but also in aligning predictors with the persistence levels they capture, following a bottom-up logic inspired by sum-of-the-parts approaches in finance (e.g., Ferreira and Santa-Clara, 2011 and Faria and Verona, 2018). Applied here for the first time to inflation forecasting, SOC provides a practical and interpretable tool for analyzing inflation dynamics.

Our results show that the SOC method delivers substantial and economically meaningful forecasting gains. Across U.S. inflation measures and horizons, SOC consistently outperforms leading time-series

benchmarks, with improvements especially pronounced at medium and long horizons, where forecast errors fall by nearly 50 percent. During the 2020-21 inflation surge – when most models, including advanced machine-learning methods (Naghi, O'Neill and Danielova Zaharieva, 2024), struggled – SOC retained strong performance by incorporating shortage indicators at the relevant frequencies. This robustness underscores its practical value for both forecasting and policy analysis.

Beyond accuracy, SOC enhances interpretability. Because the decomposition is explicit, movements in inflation forecasts can be attributed to specific predictors operating at specific frequencies. Financial variables dominate high- and business-cycle dynamics; Phillips Curve models are most informative at medium frequencies; and factor-based approaches, forecast-combination models, and shortage indicators prevail at low frequencies. This transparency is particularly valuable for policy institutions seeking to distinguish between temporary and persistent inflation pressures.

Relative to existing approaches, SOC advances the literature in three main ways. First, unlike trend-cycle decompositions that focus primarily on estimating long-run inflation (e.g., Stock and Watson, 2007, 2016, and Hasenzagl et al., 2022), SOC produces forecasts across all frequencies, preserving the richness of inflation dynamics. Second, unlike Martins and Verona (2024), who apply frequency-specific forecasting within a Phillips Curve setting, SOC generalizes this logic by combining a broad set of econometric models and predictors tailored to each frequency. Third, unlike machine-learning, wavelet-factor, and other data-rich models (e.g., Stock and Watson, 1999, 2002, Banbura, Giannone and Reichlin, 2010, Rua, 2011, 2017, Giannone, Lenza and Primiceri, 2015, Medeiros, Vasconcelos, Veiga and Zilberman, 2021, and Naghi et al., 2024), SOC remains transparent and modular, allowing the contribution of each predictor and frequency to be traced directly. At the same time, SOC is aligned with other bottom-up approaches in the forecasting literature (e.g., Byrne, O'Gorman, Scally and Zekaite, 2024 and Joseph, Potjagailo, Chakraborty and Kapetanios, 2024), but it differs in a crucial respect: the disaggregation here is not across sectors, variables, or countries, but across frequencies, which map directly onto economic persistence and horizon-specific predictability.

This frequency-specific perspective has a natural theoretical underpinning in models of information rigidities. In sticky-information frameworks (Mankiw and Reis, 2002, Reis, 2006a,b, and Verona, 2014),

agents update information infrequently, causing transitory shocks to be largely ignored while persistent shocks accumulate and shape expectations. Similarly, in noisy-information models (Sims, 2003 and Mackowiak and Wiederholt, 2015), agents filter information imperfectly, dampening high-frequency noise while allowing persistent components to dominate. Both mechanisms support the view that inflation predictability varies systematically across frequencies.

The rest of the paper is organized as follows. Section 2 illustrates the potential gains from exploiting frequency-specific information. Section 3 describes the data, the frequency decomposition, and provides a first look at the frequency dependence between inflation and its predictors. Section 4 outlines the forecasting methodology and models. Section 5 presents the empirical results, and section 6 concludes.

2 Potential gains from frequency-specific forecasting

Reliable inflation forecasts must account for the slowly evolving local mean of inflation. Consequently, accurate prediction of the low-frequency component is essential (see, e.g., Stock and Watson, 2007, 2016, del Negro and Schorfheide, 2013, Faust and Wright, 2013, and Chan, Clark and Koop, 2018). For instance, Martins and Verona (2024) show that a low-frequency New Keynesian Phillips Curve model significantly outperforms the random-walk benchmark. However, focusing exclusively on low-frequency dynamics provides only limited improvements, as it neglects the richer multi-frequency structure of inflation. Capturing these additional components – particularly short- and medium-term fluctuations – is crucial for achieving substantial forecasting gains.

To assess the potential benefits of the SOC method, we consider a hypothetical case in which each frequency component of inflation is assumed to be predicted perfectly. This exercise provides an upper bound on the accuracy that could be achieved by exploiting frequency-specific information, even though such performance is unattainable in practice. Table 1 reports the resulting root mean squared forecast error (RMSE) of Consumer Price Index (CPI) inflation relative to the random walk at horizons of one quarter (h=1), one year (h=4), and two years (h=8). The frequency components are labeled D_1 through

D_6, ordered from highest to lowest frequency.[1]

Across all forecasting horizons, only the lowest-frequency component D_6, which captures cycles longer than 16 years, outperforms the random walk on its own (see the diagonal of the table). This confirms the well-documented importance of low-frequency dynamics in inflation forecasting. However, when higher-frequency components are progressively added to D_6, forecast accuracy improves markedly. Reductions in RMSE reach up to 60 percent relative to the random walk. By contrast, aggregating higher-frequency components while excluding D_6 performs poorly, underscoring the central role of the low-frequency component as the foundation for accurate forecasts.

These results highlight the value of a forecasting method that systematically exploits the entire spectrum of inflation dynamics – a role played by the SOC method.

3 Data, filtering, and frequency-domain properties

We now turn to the empirical setting, introducing the data, the frequency-decomposition methodology, and the descriptive properties of inflation and its predictors across cycles.

3.1 Measures and predictors of inflation

We use quarterly U.S. time series spanning 1978:Q1-2024:Q4. Let P_t denote either the CPI or the Personal Consumption Expenditures (PCE) price index in quarter t. Our forecasting exercise targets the annualized h-period average inflation rate, computed as

$$\pi_t^h = \frac{400}{h} \ln\left(\frac{P_t}{P_{t-h}}\right) \ ,$$

for horizons h=1, 4, and 8, corresponding to one quarter, one year, and two years, respectively – intervals particularly relevant for monetary policy deliberations.

[1] Details of the frequency decomposition are provided in section 3.2.

PCE inflation is the Federal Open Market Committee's (FOMC) preferred measure because it aligns with its policy target and incorporates a broader expenditure base. CPI inflation, while more volatile, provides a sharper signal of prices actually paid by households and is particularly informative for short-term dynamics. As shown in table 2, panel A, PCE inflation is slightly less volatile and more persistent than CPI inflation, primarily due to lower weights on housing, energy, and food. The correlation between the two measures is high – 0.96, 0.98, and 0.99 for h=1,4, and 8 – indicating that, despite their differences, the two series capture broadly similar dynamics across horizons.

Our predictor set combines traditional macro-financial variables widely used in the inflation forecasting literature (see, e.g., Stock and Watson, 2003 and Ang, Bekaert and Wei, 2007) with novel news-based measures of supply constraints. We use 20 predictors in total.

Monetary and financial indicators. Real M2 growth is included as a monetary measure, consistent with the quantity-theory tradition. Short-term interest rates, including the three-month Treasury bill rate (TBL) and the shadow federal funds rate (SHR), transmit monetary policy conditions. The term spread (TMS), also known as the slope of the yield curve, is a well-established leading indicator of the business cycle and inflation. Long-term government bond returns (LTR) and the default return spread (DFR) capture risk premia and inflation hedging motives. The dividend-price ratio (DP) is a valuation measure often used to forecast long-horizon returns and inflation.

Expectations. Inflation expectations from the Michigan Survey of Consumers (MSC) provide forward-looking information that enters standard Phillips Curve models.

Real activity and slack. Unemployment (U) and the jobs-workers gap (JWG) proxy labor market slack. The Chicago Fed National Activity Index (CFNAI) summarizes a broad range of economic activity and has a long tradition in inflation forecasting, while the more recently proposed Sahm rule (SAHM) serves as a timely recession indicator.

Commodity and energy indicators. The CPI energy component (ENERGY) and crude petroleum producer prices (OIL) capture cost shocks that directly pass through to headline inflation at short horizons and may also shape longer-term inflation expectations.

Shortage indices. Finally, we include the Caldara, Iacoviello and Yu (2025) news-based shortage indices, which capture supply-side constraints usually not captured by standard macro-financial predictors. These include an overall index (SH-all) and disaggregated measures for industries (SH-ind), energy (SH-en), food (SH-food), and labor (SH-lab), as well as an aggregate U.S. shortage index (SH-USA).

A detailed description of the predictors is provided in appendix A. Figure 4 plots the time series of all 20 predictors, illustrating their heterogeneity in persistence and volatility over the sample. Table 2, panel B summarizes their descriptive statistics. The table highlights substantial cross-sectional variation: some predictors, such as DP and MSC, are highly persistent and smooth, while others, such as OIL and short-term interest rates, exhibit greater volatility and rapid adjustments. This diversity suggests that these variables may provide complementary information at different frequencies and forecast horizons.

Our analysis relies on final-vintage data rather than real-time vintages. Although some series – such as PCE inflation and monetary aggregates – are subject to revisions, a large body of evidence (e.g., Croushore and Stark, 2001, Clark and McCracken, 2010, and the survey by Faust and Wright, 2013) shows that revisions to several inflation measures are generally small and concentrated in short-lived, high-frequency adjustments, and forecast-error statistics are typically not sensitive to the distinction between real-time and latest-available data. Since the SOC method decomposes inflation and predictors across frequencies, such revisions have limited influence on the components that drive most of its forecast performance, namely business-cycle and lower-frequency movements. For these reasons, we base our assessment on the latest available vintages.

3.2 Frequency decomposition methodology

To extract the frequency components from inflation and its predictors, we employ wavelet multiresolution analysis (MRA). This method decomposes a time series into multiple frequency bands, enabling a more granular analysis of the short-, medium-, and long-run dynamics embedded in each series (Percival and Walden, 2000).

For a generic series X_t, the multiresolution representation is

$$X_t = \sum_{j=1}^{J} X_t^{D_j} + X_t^{S_J} , \tag{1}$$

where $X_t^{D_j}$ $(j = 1,2,\ldots,J)$ represent the detail (higher-frequency) components and $X_t^{S_J}$ denotes the smooth (low-frequency) component. Lower values of j correspond to higher-frequency fluctuations, while higher values of j capture slower, persistent movements.

We implement the maximal overlap discrete wavelet transform (MODWT) with the Haar wavelet filter, a standard choice in macrofinance applications (Bandi, Perron, Tamoni and Tebaldi, 2019, Kilponen and Verona, 2022, Martins and Verona, 2023, 2024, Stein, 2024, Canova, 2025, and Faria and Verona, 2025a,b). The detail and smooth components are computed as:

$$X_t^{D_j} = \frac{1}{2^{2j}} \left[\sum_{i=0}^{2^{(j-1)}-1} - \sum_{i=2^{(j-1)}}^{2^j-1} \right] \left[\sum_{p=0}^{2^{(j-1)}-1} - \sum_{p=2^{(j-1)}}^{2^j-1} \right] X_{t+i-p \bmod T} \tag{2}$$

and

$$X_t^{S_J} = \frac{1}{2^{2J}} \sum_{i=0}^{2^J-1} \sum_{p=0}^{2^J-1} X_{t+i-p \bmod T} , \tag{3}$$

where *mod* denotes the modulo operator for boundary correction. To minimize filtering distortions at the sample edges, we follow Gallegati, Gallegati, Ramsey and Semmler (2011), Kang, In and Kim (2017), and Stein (2024) and apply reflecting boundary conditions.

Given the sample length, we set *J*=5, which yields six components: five detail components (D_1 through D_5) and one smooth component (S_5, hereafter denoted D_6). These correspond to fluctuations with periodicities of 2-4 quarters (D_1), 1-2 years (D_2), 2-4 years (D_3), 4-8 years (D_4), 8-16 years (D_5), and more than 16 years (D_6).[2] By construction, the sum of these frequency components exactly reconstructs the original series. In other words, adding the detail and smooth components together yields the full series

[2] In the MODWT, each wavelet filter at frequency j approximates an ideal high-pass filter with passband $f \in \left[1/2^{j+1}, 1/2^j\right]$, while the smooth component is associated with frequencies $f \in \left[0, 1/2^{j+1}\right]$. As regards the choice of J, the number of observations dictates the maximum number of frequency bands that can be used. In particular, if t_0 is the number of observations in the in-sample period, then J has to satisfy the constraint $J \leq \log_2 t_0$.

without loss of information. Robustness checks with alternative filters are reported in section 5.4.2.

Figure 1 illustrates this decomposition for CPI and PCE inflation at h=1. Lower-frequency components capture the smooth, persistent trends, while higher-frequency components reflect the rapid, short-term fluctuations often associated with transitory shocks. Similar patterns are observed at h=4 and h=8, as shown in figures 2 and 3.

Because the SOC method is designed to exploit differences in persistence and horizon-specific predictability, we next analyze how inflation and its predictors behave across frequencies. This descriptive evidence provides the empirical foundation for the forecasting analysis that will follow.

3.3 Frequency-domain properties of inflation and predictors

A key advantage of the SOC method is its ability to exploit the multi-frequency structure of inflation and its predictors. In this section we explore how inflation and its predictors behave across different frequency bands. Subsection 3.3.1 summarizes the descriptive frequency characteristics of the series, while subsection 3.3.2 examines the frequency relationship between inflation and its predictors.

3.3.1 Variance decomposition by frequency

Table 3 reports the spectral decomposition of the target variables and predictors. The results reveal pronounced heterogeneity in how variance is distributed across frequencies.

For both CPI and PCE inflation, variance spans the entire spectrum but becomes increasingly concentrated in the low-frequency bands (D_5 and D_6) as the inflation measure becomes more persistent. At the two-year average (h=8), more than 60 percent of inflation's variance is captured by the D_6 component, consistent with its strong persistence and alignment with the Federal Reserve's long-run inflation target. This highlights the central role of low-frequency dynamics in shaping medium- and long-horizon forecasts.

Predictors display equally rich frequency-specific patterns. For the predictors that exhibit high persistence (see table 2), such as DP, TBL, SHR, and OIL, more than 60 percent of their variance is concentrated in

the low-frequency band D_6. In contrast, long-term government bond returns (LTR) and the default return spread (DFR) are dominated by high-frequency movements, providing timely information for short-run inflation dynamics. Other variables, such as the term spread (TMS) and unemployment (U), display a more balanced variance distribution across short-, medium-, and long-term components, highlighting their relevance across the frequency spectrum. Finally, the shortage indices (SH-) also concentrate much of their variance in the D_6 band, though SH-en is relatively more tilted toward higher frequencies, suggesting that energy-related shortages contribute both transitory and persistent components to inflation dynamics.

Overall, this complementary structure across predictors motivates the next subsection, where we examine directly how the frequency components of predictors relate to those of inflation.

3.3.2 Frequency dependence between inflation and predictors

As a preliminary step, we examine whether inflation and its predictors display systematic relationships across frequencies. This exercise is descriptive rather than a forecasting test: our aim is to illustrate the alignment of predictors with specific inflation components and to motivate the SOC method developed in section 4.

To this end, we estimate by ordinary least squares (OLS) a set of bivariate regressions of the form

$$\pi_t^{D_j,h} = c_p^{D_j,h} + \alpha_p^{D_j,h} x_{p,t}^{D_j} + \varepsilon_t^{p,D_j,h} \ ,$$

where $\pi_t^{D_j,h}$ denotes the component of inflation at frequency band D_j (j=1, ... ,6) for inflation measure h, and $x_{p,t}^{D_j}$ is the corresponding frequency component of predictor p. All predictors are standardized to unit variance prior to filtering. We focus on results for CPI inflation at h=4 (one-year average inflation), since the patterns for quarterly (h=1) and two-year-average (h=8) CPI inflation, as well as for PCE inflation, are broadly similar.

Figure 5 displays the regression results in heatmap form. Each cell corresponds to a specific predictor–frequency combination. The color indicates the explanatory power (R^2) of the regression, with lighter

shades denoting higher values, while the numerical entries report the estimated coefficient (first line) and the R^2 (second line), together with conventional significance stars. Predictors are listed along the vertical axis, with DP placed at the top, and frequency bands are shown on the horizontal axis. This format provides a compact and intuitive overview of frequency-specific relationships, facilitating comparisons across predictors and frequencies at a glance.

Several systematic patterns emerge. M2, MSC, U, CPI energy, OIL, TBL, and SAHM exhibit explanatory power across the entire frequency spectrum, indicating that these variables embed both short-run and long-run information for inflation. TMS and SHR display their strongest links at business-cycle and medium frequencies (D_3-D_5), with some spillovers into D_6, consistent with their role in capturing cyclical demand and supply pressures. DFR and JWG relate mainly to inflation at short-run and business-cycle frequencies, where their predictive content is strongest. By contrast, DP emerges primarily as a low-frequency signal, with strong predictive strength concentrated in the D_6 band. Finally, most shortage indicators load significantly at business-cycle and medium frequencies (D_3-D_5), while SH-en also has strong explanatory power in D_6, suggesting that persistent energy-related shortages can influence both cyclical and long-run inflation dynamics.

This regression analysis reveals substantial heterogeneity in how different predictors align with inflation across frequencies. Some predictors carry information across the full spectrum, others are concentrated in business-cycle or short-run bands, and a few are anchored in the low-frequency band. This diversity of frequency-specific relationships is valuable, as it provides complementary signals at different horizons. Taken together, these findings provide preliminary support for the idea that predictors should be matched to the persistence levels they capture. This motivates the need for a forecasting method that can systematically exploit such frequency-specific alignment – precisely the role of the SOC method introduced in the next section.

4 Forecasting methodology and models

4.1 Forecasting methodology

Out-of-sample forecasts are generated using a direct forecasting approach with expanding estimation windows. The first forecasts are produced using the sample 1978:Q1-1999:Q4, after which the sample is expanded by one observation and the models are re-estimated, producing a new set of forecasts. We focus on the annualized h-quarter average inflation rates π_t^h, with forecasts denoted $\hat{\pi}_{t+h}^h$.

Although rolling-window forecasts could also be considered, we adopt expanding windows because the accuracy of the wavelet decomposition benefits from the longest possible sample. Importantly, the MODWT is a two-sided filter, but the decomposition is updated recursively at each point in time using only data available up to that quarter. The use of the modulo operator in the decomposition ensures that no forward-looking information contaminates the forecasts, thereby guaranteeing real-time feasibility.

4.2 Time-series models

We next turn to the suite of forecasting models. The heterogeneous persistence and volatility of inflation and its predictors across frequencies motivate a broad range of econometric specifications. Broadly speaking, we consider standard linear models such as bivariate, multivariate, and Phillips Curve regressions; dimensionality-reduction and shrinkage methods; and forecast-combination models with fixed and time-varying weights. This set of models provides a comprehensive overview of approaches commonly used in the literature, though it is not exhaustive (see, e.g., Rossi, 2025). For example, we do not include ARIMA or (Bayesian) VARs models, as they typically do not outperform simpler time-series approaches (Faust and Wright, 2013). We also exclude theoretical models such as agent-based and DSGE models due to their mixed forecasting performance (e.g., Poledna, Miess, Hommes and Rabitsch, 2023, del Negro, Dogra, Gleich, Gundam, Lee, Nallamotu and Pacula, 2024, and Martinez-Martin, Morris, Onorante and Piersanti, 2024).

4.2.1 Benchmark model

The most widely used benchmark model in the literature on inflation forecasting is the random-walk model of Atkeson and Ohanian (2001) (denoted AO). According to this model, the h-period-ahead forecast is an average of inflation in the past four quarters: $\hat{\pi}^h_{t+h} = \frac{1}{4}\sum_{\tau=0}^{3}\pi^h_{t-\tau}$.

Another commonly used benchmark is the unobserved components stochastic volatility (UCSV) model of Stock and Watson (2007). However, this model has not consistently outperformed simple univariate benchmarks (Jarocinski and Lenza, 2018 and Banbura, Lenza and Paredes, 2024), and its forecasts are often close to those of the random-walk model (Verbrugge, 2024). Given this evidence, we leave the UCSV model out and consider the AO model as the most informative and challenging benchmark for U.S. inflation.

4.2.2 Other time-series models

We begin by considering a suite of standard time-series (TS) models widely used in the inflation forecasting literature. The first is a simple autoregressive specification, the AR(p) model:

$$\pi^h_t = \alpha + \varphi_1\pi^h_{t-1} + \varphi_2\pi^h_{t-2} + \ldots + \varphi_p\pi^h_{t-p} + \varepsilon_t \ ,$$

where p denotes the number of lags. The h-quarter-ahead forecasts are given by

$$\hat{\pi}^h_{t+h} = \hat{\alpha} + \hat{\varphi}_1\pi^h_{t-1} + \hat{\varphi}_2\pi^h_{t-2} + \ldots + \hat{\varphi}_p\pi^h_{t-p} \ .$$

The lag length is selected using both the AIC and SIC criteria, with a maximum of six lags allowed; the resulting models are denoted AR-AIC and AR-SIC.

Next, we estimate a New Keynesian Phillips Curve (NKPC) model, following the framework of Coibion and Gorodnichenko (2015). The model is specified as:

$$\pi^h_t = c^{h,PC} + \alpha^h_1 MSC_{t-1} + \alpha^h_2 slack_{t-1} + \alpha^h_3 ENERGY_{t-1} + \varepsilon^{h,PC}_t \ ,$$

where MSC represents inflation expectations, *slack* is alternatively measured by the unemployment rate (U), the JWG index, or the Sahm rule (SAHM), and *ENERGY* is energy inflation. Forecasts are computed as:

$$\hat{\pi}^h_{t+h} = \hat{c}^{h,PC} + \hat{\alpha}^h_1 MSC_t + \hat{\alpha}^h_2 slack_t + \hat{\alpha}^h_3 ENERGY_t \ .$$

The three variants of this model are labeled PC1 (U as the slack measure), PC2 (JWG), and PC3 (SAHM).

In addition, we estimate 20 bivariate models of the form:

$$\pi^h_t = c^{h,p} + \beta^{h,p} x_{p,t-1} + \varepsilon^{h,p}_t \ ,$$

where x_p is one of the 20 predictors described in section 3.1, and the corresponding forecasts are:

$$\hat{\pi}^h_{p,t+h} = \hat{c}^{h,p} + \hat{\beta}^{h,p} x_{p,t} \ .$$

Models are labeled with the acronym of the predictor used (e.g., MSC, DP, SH).

We further consider shrinkage methods, including LASSO (Tibshirani, 1996), elastic net (ENET; Zou and Hastie, 2005), and ridge regression (RIDGE; Hoerl and Kennard, 1970). These models handle high-dimensional predictor sets by shrinking coefficients toward zero, which helps prevent overfitting while retaining relevant predictors.[3]

To exploit dimensionality reduction, we apply principal component analysis (PCA) and partial least squares (PLS). For PCA, the first principal component extracted from the predictor set is used as a single predictor in the bivariate regression. While PCA effectively reduces dimensionality, it does so without exploiting the information contained in the dependent variable, potentially limiting its predictive power. To address this, we also implement a PLS regression, which incorporates information from the dependent

[3] For the LASSO forecast, we estimate a penalized regression that shrinks some coefficients exactly to zero. We let the strength of the penalty vary over a fine grid of values between 0.01 and 1, and we select the preferred specification using five-fold cross-validation. For RIDGE, we estimate a regression that shrinks coefficients toward zero without eliminating predictors; here the penalty strength ranges from 0.1 to 1, and we choose the value that yields the lowest in-sample prediction error. For the ENET, which blends the two types of shrinkage, we evaluate several mixtures between LASSO-type and RIDGE-type penalization, and for each mixture we again search over the same grid of penalty strengths as in the LASSO case. The final ENET model corresponds to the combination that minimizes the cross-validation error.

variable when extracting the factor, thereby improving predictive performance (Kelly and Pruitt, 2013, 2015 and Dai, Jiang, Kang and Xue, 2025). For out-of-sample forecasts, we include only the first PLS factor as a predictor.[4] In both these dimensionality reduction models, the predictors are standardized to mean zero and unit variance before extracting the components.

The last group of models consists of forecast combinations, which pool information across the previous 30 models following the approach of Stock and Watson (2004) and Rapach, Strauss and Zhou (2010). Letting $\hat{\pi}^h_{m,t+h}$ denote the forecast from model m, the combined forecast is:

$$\hat{\pi}^h_{c,t+h} = \sum_{m=1}^{m_{max}} \omega_{m,t} \hat{\pi}^h_{m,t+h} \ ,$$

where $\omega_{m,t}$ represents the combining weights and m_{max} is the number of models combined. We consider two broad classes of combinations. The first applies fixed weighting schemes: a simple average across forecasts (C-MEAN), the median (C-MEDIAN), and a trimmed mean (C-Tr.Mean) that excludes the lowest and highest forecasts before averaging. The second class employs time-varying weights based on the historical out-of-sample performance of individual models, using the discounted mean squared prediction error (DMSPE) method. In this approach, the weights are given by:

$$\omega_{m,t} = \frac{\phi_{m,t}^{-1}}{\sum_{m=1}^{30} \phi_{m,t}^{-1}} \quad , \quad \phi_{m,t} = \sum_{s=n}^{t-1} \theta^{t-1-s} \left(\pi^h_{t+h} - \hat{\pi}^h_{m,t+h} \right)^2 \ ,$$

where n denotes the last observation before the holdout period and θ is the discount factor controlling how much weight is placed on more recent performance. When $\theta = 1$, there is no discounting, replicating the optimal combination rule of Bates-Granger (1969) under uncorrelated forecast errors. Lower values of θ place greater emphasis on recent performance, producing more dynamic weights. We implement four specifications with $\theta = 0.25$, 0.50, 0.75, and 1, denoted C-DMSPE 0.25, C-DMSPE 0.50, C-DMSPE 0.75, and C-DMSPE 1, respectively. The holdout period used to initialize the performance weights spans

[4] For both PCA and PLS, it is possible to include multiple factors – as well as linear or non-linear combinations of these factors – within a multivariate regression framework, as in Ludvigson and Ng (2007) and Lin, Wu and Zhou (2018). Huang, Jiang, Li, Tong and Zhou (2022) propose a scaled principal component analysis, which, in spirit, is conceptually similar to PLS, as it assigns greater weight to predictors that are more strongly correlated with the target variable during factor extraction.

1995:Q1-1999:Q4.[5]

In total, we consider 37 TS models: two AR models, three Phillips Curve models, 20 bivariate models, PCA, PLS, three shrinkage models (LASSO, ENET, RIDGE), and seven forecast combinations (of the previous 30 models). Ultimately, for each inflation measure and horizon, we report the TS specification with the lowest out-of-sample RMSE over the full evaluation period.

4.3 The SOC method

For each frequency component of inflation, we estimate 38 models: the same 37 TS models described in section 4.2.2 plus one additional variant of the PLS model.

To illustrate how the forecasts in the frequency domain are made, consider the Phillips Curve specification. For each frequency D_j (j=1,...,6) of inflation, we estimate:

$$\pi_t^{D_j,h} = c^{D_j,h,PC} + \alpha_1^{D_j,h} MSC_{t-1}^{D_j} + \alpha_2^{D_j,h} slack_{t-1}^{D_j} + \alpha_3^{D_j,h} ENERGY_{t-1}^{D_j} + \varepsilon_t^{D_j,h,PC} \ ,$$

where $\pi_t^{D_j,h}$ is frequency D_j of inflation, $MSC_{t-1}^{D_j}$ is frequency D_j of MSC inflation expectations, and so on. The forecasts with this model are then computed as

$$\hat{\pi}_{t+h}^{D_j,h} = \hat{c}^{D_j,h,PC} + \hat{\alpha}_1^{D_j,h} MSC_t^{D_j} + \hat{\alpha}_2^{D_j,h} slack_t^{D_j} + \hat{\alpha}_3^{D_j,h} ENERGY_t^{D_j} \ .$$

This frequency-by-frequency approach allows the model to exploit the alignment between predictors and inflation at the same frequency, ensuring that high-frequency shocks in the predictors are mapped to high-frequency components of inflation, while persistent predictors inform the lower-frequency components of inflation.

[5] Alternative ways to model instability in combination weights include Bayesian dynamic model averaging (e.g., Koop and Korobilis, 2012 and Groen, Paap and Ravazzolo, 2013) and model-selection procedures based on the model confidence set (Hansen, Lunde and Nason, 2011 and Baumeister, Huber, Lee and Ravazzolo, 2026). Implementing these approaches within our multi-frequency framework would substantially increase computational burden with limited expected forecasting gains relative to the DMSPE combinations used here.

The additional specification, PLS-2, builds upon the standard PLS model (hereafter PLS-1) by modifying the way the factors are extracted. Whereas PLS-1 computes the first factor by correlating each frequency component of inflation with the original, unfiltered set of predictors, PLS-2 instead constructs the factor from the frequency-decomposed predictors that match the frequency of the inflation component being forecast. This alignment of information sets aims to enhance the predictive content of the PLS factor by ensuring that the dimensionality-reduction step respects the multi-frequency structure of the data.

For each inflation measure, frequency band, and forecast horizon, we select the model with the lowest out-of-sample RMSE among the 38 candidates. The corresponding forecasts of the six frequency components (D_1–D_6) are then summed to obtain the aggregate inflation forecast according to the SOC method.

By construction, this frequency-domain approach delivers a richer and more granular modeling environment, allowing the forecasts to incorporate heterogeneous dynamics across temporal horizons. The importance of this alignment will become evident in the next section, where we show that the frequency-domain approach yields systematic gains over time-domain models, particularly at medium and long horizons.

5 Empirical results

Having introduced the forecasting framework, we now present the empirical evidence on how frequency-specific information improves inflation forecasts. We first examine which predictors and models drive performance across frequencies and horizons, then assess overall and dynamic accuracy, before turning to the optimized SOC version, the role of shortage indicators, and robustness checks.

5.1 Baseline performance of the SOC method

We assess the baseline performance of the SOC method by examining both its overall accuracy and its dynamic behavior over time. These analyses show whether and when the SOC method improves on

standard approaches. Before turning to forecast accuracy, we first identify which models and predictors drive performance across frequencies and horizons, providing the economic intuition behind the results.

5.1.1 Models and predictors by frequency and horizon

This subsection identifies which forecasting models and predictors generate SOC's gains, linking forecast performance to underlying economic signals. Table 5 reports, for each forecast horizon, the best-performing models for the aggregate inflation series and for each of its frequency components in the SOC method. The first column lists the forecast horizons, while the second column shows the best TS model for overall inflation. Columns 3 through 8 identify the best models for each frequency component (D_1 to D_6, where D_1 denotes high-frequency and D_6 low-frequency variation). The upper panel refers to forecasts of CPI inflation, while the lower panel presents the corresponding results for PCE inflation.

The aggregate TS forecasts highlight the importance of model combinations. At the one-quarter horizon (h=1), the best performer is the C-DMSPE 0.25 combination, which assigns time-varying weights to individual models based on their very recent performance. This emphasizes the benefits of quickly adaptive weighting in short-term forecasts, where responsiveness to changing conditions is critical. At the four- and eight-quarter horizons (h=4 and h=8), the best performer shifts to the C-MEAN combination, which equally averages forecasts across the candidate models, suggesting that stability and diversification gain importance at medium and longer horizons. This confirms the well-known finding that combining forecasts – even in a simple way – shields against instabilities in individual models and improves forecast performance.

For the frequency-specific components, results show systematic variation across horizons. In the higher-frequency bands (D_1 to D_4), C-DMSPE and PLS forecasts dominate at h=1. As the horizon lengthens to four and eight quarters, individual predictor models become the top performers. Among these, the dividend-price ratio (DP) stands out for its consistent selection, indicating that it provides valuable information for predicting short-run and cyclical components of inflation, particularly at longer horizons. This likely reflects the ability of the DP ratio to capture forward-looking financial conditions that drive aggregate demand and, hence, inflation dynamics.

For the medium-frequency band (D_5), results indicate that individual predictor models perform best at h=1 – specifically, the DP for CPI inflation and the Michigan Survey of Consumers (MSC) for PCE inflation. At longer horizons (h=4 and h=8), the Phillips Curve model becomes the leading specification, with the SAHM cyclical indicator being particularly effective for capturing these medium-term inflation dynamics.

For the low-frequency band (D_6), results underscore its stability and persistence. At shorter horizons (h=1 and h=4), the C-DMSPE 1 model – with time-varying combination weights – highlights the benefits of adaptive pooling for slow-moving inflation trends. At the eight-quarter horizon (h=8), the leading models change: for CPI inflation, the PCA specification dominates, while for PCE inflation, the overall shortage index (SH-all) emerges as the top predictor. This shift implies that structural indicators capturing persistent economic pressures become increasingly informative when forecasting the low-frequency component of inflation over longer horizons.

Overall, these findings confirm the central insight of the paper: the SOC method enhances forecast accuracy by flexibly matching each frequency component and horizon with its most suitable model. This frequency-predictor mapping provides the economic intuition for the empirical forecast gains discussed next.

5.1.2 Overall forecast accuracy

We now compare aggregate forecast accuracy across models and horizons relative to benchmark models. Forecast accuracy is evaluated using the RMSE, calculated for each model and compared with the AO benchmark model. Statistical significance is assessed using the Diebold and Mariano (1995) and West (1996) tests, with Newey-West standard errors to account for autocorrelation and heteroscedasticity.

We begin by evaluating the performance of each forecasting model individually. Table 4 summarizes the root mean squared errors (RMSEs) for the different model specifications. Panel A reports the RMSEs of the Atkeson-Ohanian (AO) benchmark model, while panel B presents the RMSEs of the time-series (TS) and SOC models, expressed relative to the AO benchmark. Values below one indicate that the

alternative model outperforms AO. Statistical significance is marked with one, two, and three stars for the 10 percent, 5 percent, and 1 percent levels, respectively. The final row of panel B reports the RMSE of the SOC method relative to that of the TS model, allowing a direct comparison between the SOC and aggregate time-series models.

The AO model performs better for PCE than for CPI, making it a more challenging benchmark to outperform for the former. This reflects the fact that PCE inflation filters out the most volatile components and better captures underlying trends. For example, during the Global Financial Crisis, CPI inflation plunged sharply from 6.2 percent to -9.2 percent between 2008:Q3 and 2008:Q4, whereas PCE inflation exhibited a more moderate decline from 4.3 percent to -6.4 percent over the same period.

For the TS model, the short-run gains are clear: at h=1, the relative RMSE falls to 0.889 for CPI and 0.896 for PCE, both statistically significant at the 1 percent level. By h=4, the advantage becomes mixed: CPI remains better than AO at 0.848 (significant), whereas PCE is 0.859, not statistically different from AO. At h=8, the TS model no longer improves on AO.

The SOC method, by contrast, improves in a statistically significant way on AO across all horizons and both inflation measures: 0.853 / 0.858 at h=1, 0.789 / 0.770 at h=4, and 0.658 / 0.699 at h=8 for CPI / PCE, respectively. Relative to the best TS model, SOC shows no significant gain at h=1 (0.959 for both CPI and PCE), a modest but significant edge for PCE at h=4 (0.897), and substantial, statistically significant gains at h=8 (0.717 for CPI; 0.761 for PCE).

Taken together, the TS model delivers modest improvements at the very short horizon, but these gains fade as the horizon lengthens. By contrast, SOC's frequency-specific aggregation delivers forecasting gains that are robust at h=1 and expand markedly at h=4 and h=8, underscoring its advantage over the aggregate TS model. Next we assess how forecast accuracy evolves through time and across major inflation episodes.

5.1.3 Dynamic performance

Graphical evidence illustrates the economic significance of forecast differences and how model performance evolves over time.

The left panels in figure 6 display the realized path of CPI inflation (black solid line) together with forecasts from the AO benchmark (green line), the best time-series (TS) model (red line), and the SOC method (cyan line). The figure is organized by forecast horizon, with separate panels for one-quarter, four-quarter, and eight-quarter-ahead forecasts, allowing direct comparison of short-, medium-, and long-term predictive performance over time.

At the one-quarter horizon, the forecasts of all three models closely track actual CPI inflation during relatively calm periods, such as 2012-2016, when inflation remained stable. Nonetheless, the SOC method shows a slight edge in turbulent periods. During the 2008-09 Global Financial Crisis, SOC captures the sharp drop in CPI inflation more accurately than AO or TS, both of which adjust too slowly. Similarly, during the early stages of the post-pandemic surge (2020-21), SOC is quicker to reflect the upward pressure on prices, while AO and TS remain too conservative.

At the four- and eight-quarter horizons, differences between models become clearer. The AO benchmark systematically lags major turning points, offering little anticipation of sharp movements. The TS model provides moderate improvement but still tends to smooth medium-term swings. By contrast, the SOC method better anticipates medium-run inflation dynamics: it more accurately predicts the depth of the 2008-09 downturn and the sustained build-up of inflation pressures during 2021-22.

These time-series patterns are consistent with the cumulative forecast error profiles shown in the right panels in figure 6. We plot the cumulative sum of squared forecast-error differences, defined as "AO minus alternative" for TS (red) and SOC (cyan); an upward slope therefore indicates periods when the alternative model outperforms AO. At the one-quarter horizon, the cumulative error differences show modest but consistent gains for SOC during volatile episodes such as 2008-09 and 2020-21, while SOC's performance is similar to that of the AO and TS models in tranquil periods. At the four-quarter horizon, the upward drift in the cumulative error line is more pronounced, underscoring SOC's growing

advantage as the horizon lengthens. At the eight-quarter horizon, the slope steepens markedly during the post-pandemic inflation shock, signaling that SOC's frequency-based decomposition provides substantial improvements in capturing the persistence of inflationary pressures.

Taken together, figure 6 highlights a consistent pattern: the SOC method delivers robust gains at all horizons but is particularly effective at medium- and longer-term forecasts. The TS model performs reasonably well at short horizons, but its performance deteriorates as the forecast horizon extends, leaving SOC as the only approach that consistently adapts to both short-run fluctuations and persistent low-frequency movements in CPI inflation.

For PCE inflation, the left panels in figure 7 plot realized PCE inflation (black solid line) together with AO (green), TS (red), and SOC (cyan) forecasts across horizons. The patterns are broadly similar. Consistent with the CPI results, at h=1, all models track PCE inflation relatively well in calm periods, though SOC responds more quickly during turbulent episodes like the 2008-09 crisis and the onset of the 2020-21 surge. At longer horizons, both AO and TS increasingly underpredict persistent movements, while SOC remains closer to the realized path, particularly during the sustained run-up after 2020 and the gradual disinflation following 2008-09. The right panels confirm this: SOC's cumulative gains grow steadily with horizon, reflecting superior capture of low-frequency components in PCE. Consistent with the CPI results, the evidence confirms that SOC is the only approach that consistently adapts to short-run fluctuations while accurately capturing the persistent low-frequency dynamics that drive PCE inflation.

Taken together, the average and dynamic results confirm that the SOC method delivers systematic gains over both benchmarks. We now analyze whether SOC's gains reflect improved efficiency or bias reductions.

5.1.4 Bias-variance decomposition

To clarify the mechanism behind SOC's gains, we use the Theil (1971) mean squared error (MSE) decomposition, following the implementation of Rapach et al. (2010). This splits the MSE into two components: squared forecast bias and a remainder term that reflects forecast efficiency (variance). Points closer to the

origin indicate superior performance, combining low bias (x-axis) and high efficiency (y-axis).

The scatterplots in figure 8 present the results. Across horizons and inflation measures, the SOC method consistently lies closer to the origin than both the AO benchmark and the best TS model, highlighting its ability to balance the bias-variance trade-off more effectively. In particular, the results indicate that the SOC method achieves superior accuracy primarily by reducing bias (especially versus TS models), while maintaining similar or higher forecast efficiency compared with the TS and AO models, respectively.

5.2 The optimized SOC method

Building on the baseline SOC results, we next examine whether forecast accuracy can be further improved by selectively aggregating the most informative frequency components. The SOC method is uniquely suited to this optimization because its modular structure treats each frequency band separately; aggregate time-series models cannot target specific bands in this way. By exploiting this flexibility, we can evaluate whether excluding noise-dominated components enhances predictive performance, while still retaining the transparency of the SOC method.

The SOC method constructs the aggregate inflation forecast as the sum of the best individual forecasts of the six frequency components (D_1 to D_6). Recent studies (Faria and Verona, 2021 and Martins and Verona, 2024) suggest that including all frequency components may not always be optimal, as some higher-frequency bands can introduce noise, particularly at medium and longer horizons, and thus hurt the forecast performance. Furthermore, it may be the case that summing other frequency forecasts might yield better inflation forecasts than the sum of the best individual frequency forecasts (as in the baseline SOC method).

Guided by this evidence, we implement an optimized SOC method that, for each horizon, constructs the forecast of inflation based on the sum of the four frequency-component forecasts that minimizes the RMSE. This retains the benefits of frequency decomposition while improving accuracy by excluding less informative components. We call this the optimized SOC method.[6]

[6] No additional gains were found by summing five individual frequency forecasts.

Table 6 reports, for each inflation measure and forecast horizon, the individual component forecasts used in the optimized SOC method. Several clear patterns emerge. First, as expected, the low-frequency component (D_6) is consistently retained, typically using PCA or SH-ind forecasts, reflecting its importance for capturing slow-moving inflation trends. Second, the high-frequency component (D_1), dominated by noise, is never included. Third, aside from PCA – which pools information from the full predictor set – the optimized SOC method relies almost exclusively on individual predictors, favoring more parsimonious and targeted signals. Fourth, none of the combination forecasts (that are often used in the baseline SOC method) are selected, highlighting that the optimization process systematically prefers sharper, component-specific information over broad averaging. Finally, the specific forecasts selected by the optimized SOC method differ from those used in the baseline SOC method reported in table 5, underscoring that this optimization step meaningfully reshapes the aggregation across cycles.

Results in panel C of table 4, together with the blue lines in figures 6 and 7, show that the optimized SOC method delivers substantial gains relative to both the AO benchmark and the aggregate TS model at all horizons. Relative RMSEs indicate reductions up to about 50 percent, with the largest improvements at h=4 and h=8 for both CPI and PCE inflation. Even at the one-quarter horizon (h=1), where high-frequency noise complicates forecasting, the optimized SOC method consistently improves accuracy, though the gains are smaller. These improvements are also visible in the left panels of figures 6 and 7, which show a tighter alignment of the forecasts with realized inflation, and in the right panels of those figures, where the blue lines lie consistently above those of the other models, signaling superior cumulative performance.

Looking at the bias-variance decomposition in figure 8, the improved performance of the optimized SOC method, compared to the baseline SOC method, reflects substantial improvements in forecast efficiency while maintaining low bias.

Comparing results in panel C with those in panel B highlights the incremental value of the optimization: the optimized SOC method reduces forecast errors by an additional 10 to 20 percent relative to the baseline SOC, demonstrating the added benefit of selectively including the most informative components.

Table 7 quantifies the incremental contribution of each included frequency-forecast component. We de-

note by M4 the forecast using only the lowest-frequency component; M3 augments this with the second-lowest component; M2 adds the third-lowest; and M1 adds the fourth-lowest. Hence, "M1-M4" refers to the cumulative sum of the four frequency components selected by the optimization. Across horizons and for both CPI and PCE inflation, M4 (the low-frequency model) delivers the largest individual gains relative to the AO benchmark, particularly at longer horizons (h=4 and h=8), where reductions in RMSE are economically meaningful and statistically significant. Adding additional forecasts – M3, M2, and M1, separately or together – yields incremental improvements, with many of these forecasting gains significant at the 5 percent or even 1 percent level. The full optimized SOC method (sum of the M1-M4 forecasts, reported in the last row) achieves the lowest RMSEs and the strongest statistical significance across all horizons. Moving from M4 alone to the full optimized combination typically results in an additional reduction of around 5 to 10 percent in RMSE and increased statistical significance in most cases.

Taken together, these results show that the optimized SOC method delivers substantial accuracy gains by including the most informative frequency components. A natural question is whether this flexibility also enhances performance in periods of heightened inflation volatility, such as the 2020-21 inflation surge, to which we turn next.

5.3 Shortage predictors and the 2020-21 inflation surge

The inflation surge of 2020-21 exposed the central role of supply bottlenecks and shortages in shaping inflation dynamics. To quantify their contribution, we compare forecasts from the optimized SOC method with and without shortage-related predictors.

Figure 9 plots realized inflation (black lines), the forecasts from the optimized SOC method including shortages (blue solid lines), and the forecasts from the same method excluding shortages (blue dashed lines). For both CPI (top panels) and PCE (bottom panels), and across all horizons, incorporating shortage indicators markedly improves real-time tracking of the inflation upswing. SOC without shortages persistently underpredicts the magnitude and persistence of the surge, especially at medium and long

horizons. This demonstrates that timely indicators of supply constraints provide valuable information that becomes particularly relevant when inflation is driven by bottlenecks and production delays. This interpretation is consistent with recent micro-evidence showing that supply-side bottlenecks generated large price increases in a small set of product categories, with these granular forces accounting for a substantial share of the 2020-21 inflation surge (Alvarez-Blaser, Auer, Lein and Levchenko, 2025).

The figure also includes the forecasts from the best-performing time-series model (red lines). This model also uses shortage indicators but does not exploit the SOC decomposition. Comparing the TS model with the optimized SOC method including shortages isolates the value of the SOC structure itself. Despite using similar predictors, the TS model fails to anticipate the speed and persistence of the inflation acceleration and displays larger forecast errors throughout the episode. In counterfactual terms, relying on the TS model – rather than SOC – would have provided policymakers with substantially weaker foresight during a period of rapidly shifting inflation dynamics.

Taken together, these results illustrate two distinct but complementary points. First, shortage indicators are essential for capturing inflation movements driven by supply-side disruptions. Second, even when predictors are held constant, the SOC method delivers superior real-time performance relative to standard time-series models. This highlights the usefulness of SOC in turbulent periods, precisely when traditional approaches may struggle to adapt.

5.4 Robustness checks

We assess robustness along three dimensions: the definition of the inflation measure, the choice of filter used for the frequency decomposition, and comparisons with professional forecasts.

5.4.1 Alternative inflation measures

We first examine robustness to alternative measures of inflation. In particular, we use the SOC method and its optimized variant to forecast core CPI and the GDP price deflator. The results, available upon request, indicate that the main findings carry over. For the GDP price deflator, performance is broadly in

line with that for headline CPI and PCE. For core CPI, the baseline SOC performs somewhat less well, but the optimized SOC method delivers even larger gains at long horizons (h=8). These additional results confirm that the forecasting gains of the SOC method are not specific to headline measures but extend to both core and output-based indicators of inflation.

5.4.2 Filter choice

We next investigate the sensitivity of the SOC method to the choice of filter used for the frequency decomposition. While widely used, the Haar filter is only an approximation to an ideal band-pass filter, which raises the question of whether the results are sensitive to this choice. To address this concern, we compare outcomes across alternative frequency decompositions. Using different wavelet filters – Daubechies-2 and Fejér-Korovkin-4 – yields slightly lower accuracy than our baseline, but the differences are modest, indicating that the results are not overly sensitive to the specific wavelet choice. In contrast, replacing the wavelet decomposition with the Christiano and Fitzgerald (2003) asymmetric band-pass filter – calibrated to extract frequency components that match those obtained with the Haar filter – produces a marked deterioration in forecast accuracy. These findings mirror those in Faria and Verona (2020) and Martins and Verona (2024). The Christiano-Fitzgerald filter is less adept at handling non-stationarities and structural breaks, making it less effective at capturing the time-varying dynamics of inflation.

Overall, these results confirm that SOC's gains are robust to filter choice and that wavelet filters perform better than alternative band-pass methods.

5.4.3 Comparison with professional forecasts

Finally, we compare the SOC method with the Survey of Professional Forecasters (SPF), a widely used and policy-relevant benchmark in empirical forecasting work. In our baseline specifications, expectations enter only through the New Keynesian Phillips Curve, which incorporates the Michigan Survey of Consumers (MSC). Prior work shows that SPF expectations contain useful information about future inflation

(e.g., Ang et al., 2007 and Berge, 2018).

To benchmark SOC against this reference, we compare its performance with the SPF forecasts. The SPF is not directly comparable to the mechanical models considered in this paper, as its projections incorporate expert judgment, large-scale institutional models, and information available at the time of each survey round. Moreover, because the SPF does not elicit 1-quarter-ahead or 8-quarter-ahead expectations, we follow Berge (2018) and use the 4-quarter-ahead expectation as a proxy for all horizons.

The results show that SOC performs favorably relative to the SPF. For CPI inflation, the optimized SOC reduces RMSEs to 0.84, 0.73, and 0.62 at horizons h=1, 4, and 8, with a statistically significant improvement (at the 1 percent level) at the short horizon. For PCE inflation, the optimized SOC achieves relative RMSEs of 0.82, 0.73, and 0.65, each statistically significant (at either the 5 percent or 1 percent level). These patterns indicate that the frequency-domain approach delivers gains even relative to a benchmark that embeds institutional judgment and non-mechanical information.

Taken together, these findings reinforce the earlier robustness exercises and support the central result: decomposing inflation into frequency components and aligning predictors accordingly systematically improves forecasting performance and yields economically interpretable gains.

6 Conclusion

This paper has shown that decomposing inflation into its frequency components provides a powerful way to improve both the accuracy and interpretability of inflation forecasts. The sum-of-the-cycles (SOC) method systematically exploits the multi-frequency structure of inflation by matching each component with the model best suited to its persistence. Across U.S. inflation measures and forecast horizons, SOC consistently outperforms standard benchmarks, with particularly large gains at medium and long horizons. These improvements reflect the method's ability to harness complementary information across frequencies, thereby reducing forecast errors in both tranquil and turbulent periods, including the 2020-21 inflation surge.

From a policy perspective, the gains at medium and long horizons are especially relevant. During episodes such as the 2020-21 surge, more reliable real-time assessments of the magnitude and persistence of inflation would have altered the perceived policy trade-offs. By providing clearer signals at these horizons, the SOC method reduces the risk of underestimating persistent inflation pressures and supports more timely and informed policy decisions.

For policy institutions, SOC can serve as a frequency-specific cross-check alongside DSGE, semi-structural, and judgmental forecasts, helping to diagnose whether forecast errors reflect temporary disturbances or persistent shifts. Because the decomposition is explicit, SOC clarifies which predictors drive inflation dynamics at different horizons, offering policymakers a transparent lens through which to interpret and communicate inflation developments.

In sum, the evidence shows that decomposing inflation into its frequency components improves forecasting accuracy while deepening economic interpretation. The sum of the cycles is indeed more than the whole: SOC not only sharpens forecasts but also illuminates the underlying forces shaping inflation, providing valuable insights for both policy and future research.

References

Alvarez-Blaser, Santiago, Raphael Auer, Sarah M. Lein, and Andrei A. Levchenko, "The Granular Origins of Inflation," NBER Working Papers 33404 2025.

Ang, Andrew, Geert Bekaert, and Min Wei, "Do macro variables, asset markets, or surveys forecast inflation better?," *Journal of Monetary Economics*, 2007, *54* (4), 1163–1212.

Atkeson, Andrew and Lee E. Ohanian, "Are Phillips curves useful for forecasting inflation?," *Federal Reserve Bank of Minneapolis Quarterly Review*, 2001, *25* (2), 2–11.

Banbura, Marta, Domenico Giannone, and Lucrezia Reichlin, "Large Bayesian Vector Auto Regressions," *Journal of Applied Econometrics*, 2010, *25* (1), 71–92.

_ , Michele Lenza, and Joan Paredes, "Forecasting Inflation in the US and in the Euro Area," in Michael P. Clements and Ana Beatriz Galvao, eds., *Handbook of Research Methods and Applications*

in Macroeconomic Forecasting, Edward Elgar Publishing, 2024, chapter 9, pp. 218–245.

Bandi, Federico, Bernard Perron, Andrea Tamoni, and Claudio Tebaldi, "The Scale of Predictability," *Journal of Econometrics*, 2019, *208* (1), 120–140.

Bates, J. M. and C. W. J. Granger, "The Combination of Forecasts," *Operational Research Quarterly*, 1969, *20* (4), 451–468.

Baumeister, Christiane, Florian Huber, Thomas K. Lee, and Francesco Ravazzolo, "Forecasting Natural Gas Prices in Real Time," *Journal of Applied Econometrics*, 2026, *forthcoming.*

Berge, Travis J., "Understanding survey-based inflation expectations," *International Journal of Forecasting*, 2018, *34* (4), 788–801.

Byrne, Stephen, Paraic O'Gorman, John Scally, and Zivile Zekaite, "Inflation Forecasting at the Central Bank of Ireland," *Quarterly Bulletin Articles*, June 2024, pp. 71–107.

Caldara, Dario, Matteo Iacoviello, and David Yu, "Measuring shortages since 1900," International Finance Discussion Papers 1407 2025.

Canova, Fabio, "FAQ: How do I estimate the output gap?," *Economic Journal*, 2025, *135* (665), 59–80.

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

Christiano, Lawrence J. and Terry J. Fitzgerald, "The Band Pass Filter," *International Economic Review*, 2003, *44* (2), 435–465.

Clark, Todd E. and Michael W. McCracken, "Averaging forecasts from VARs with uncertain instabilities," *Journal of Applied Econometrics*, 2010, *25* (1), 5–29.

Coibion, Olivier and Yuriy Gorodnichenko, "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation," *American Economic Journal: Macroeconomics*, 2015, *7* (1), 197–232.

Croushore, Dean and Tom Stark, "A real-time data set for macroeconomists," *Journal of Econometrics*, 2001, *105* (1), 111–130.

Dai, Zhifeng, Fuwei Jiang, Jie Kang, and Bowen Xue, "Stock return predictability in the frequency domain," *International Journal of Forecasting*, 2025, *41* (3), 1126–1147.

del Negro, Marco and Frank Schorfheide, "DSGE Model-Based Forecasting," in G. Elliott, C. Granger, and A. Timmermann, eds., *Handbook of Economic Forecasting*, Vol. 2 of *Handbook of Economic Forecasting*, Elsevier, 2013, chapter 2, pp. 57–140.

_ , Keshav Dogra, Aidan Gleich, Pranay Gundam, Donggyu Lee, Ramya Nallamotu, and Brian Pacula, "The NY Fed DSGE Model: A Post-COVID Assessment," *AEA Papers and Proceedings*, 2024, *114*, 95–100.

Diebold, Francis X. and Roberto S. Mariano, "Comparing Predictive Accuracy," *Journal of Business & Economic Statistics*, 1995, *13* (3), 253–263.

Faria, Gonçalo and Fabio Verona, "Forecasting stock market returns by summing the frequency-decomposed parts," *Journal of Empirical Finance*, 2018, *45*, 228 – 242.

_ **and** _ , "The yield curve and the stock market: Mind the long run," *Journal of Financial Markets*, 2020, *50* (C).

_ **and** _ , "Out-of-sample time-frequency predictability of the equity risk premium," *Quantitative Finance*, 2021, *21* (12), 2119–2135.

_ **and** _ , "The economic value of frequency-domain information," *Journal of Portfolio Management*, 2025, *51* (4), 128–143.

_ **and** _ , "Unlocking predictive potential: the frequency-domain approach to equity premium forecasting," *Journal of Empirical Finance*, 2025, *83*, 101648.

Faust, Jon and Jonathan H. Wright, "Forecasting Inflation," in G. Elliott, C. Granger, and A. Timmermann, eds., *Handbook of Economic Forecasting*, Vol. 2 of *Handbook of Economic Forecasting*, Elsevier, 2013, chapter 1, pp. 2–56.

Ferreira, Miguel A. and Pedro Santa-Clara, "Forecasting stock market returns: the sum of the parts is more than the whole," *Journal of Financial Economics*, 2011, *100* (3), 514–537.

Gallegati, Marco, Mauro Gallegati, James B. Ramsey, and Willi Semmler, "The US Wage Phillips Curve across Frequencies and over Time," *Oxford Bulletin of Economics and Statistics*, 2011, *73* (4), 489–508.

Giannone, Domenico, Michele Lenza, and Giorgio E. Primiceri, "Prior Selection for Vector Autore-

gressions," *Review of Economics and Statistics*, 2015, *97* (2), 436–451.

Gordon, Robert J., "The Recent Acceleration of Inflation and Its Lessons for the Future," *Brookings Papers on Economic Activity*, 1970, *1* (1), 8–47.

Goyal, Amit, Ivo Welch, and Athanasse Zafirov, "A Comprehensive 2022 Look at the Empirical Performance of Equity Premium Prediction," *Review of Financial Studies*, 2024, *37* (11), 3490–3557.

Groen, Jan J.J., Richard Paap, and Francesco Ravazzolo, "Real-Time Inflation Forecasting in a Changing World," *Journal of Business & Economic Statistics*, 2013, *31* (1), 29–44.

Hansen, Peter R., Asger Lunde, and James M. Nason, "The Model Confidence Set," *Econometrica*, 2011, *79* (2), 453–497.

Hasenzagl, Thomas, Filippo Pellegrino, Lucrezia Reichlin, and Giovanni Ricco, "A Model of the Fed's View on Inflation," *Review of Economics and Statistics*, 2022, *104* (4), 686–704.

Hatzius, Jan, "Inflation: What we have learned and what we need to know," *Journal of Monetary Economics*, 2024, *148* (S).

Hoerl, Arthur E. and Robert W. Kennard, "Ridge Regression: Biased Estimation for Nonorthogonal Problems," *Technometrics*, 1970, *12* (1), 55–67.

Huang, Dashan, Fuwei Jiang, Kunpeng Li, Guoshi Tong, and Guofu Zhou, "Scaled PCA: A New Approach to Dimension Reduction," *Management Science*, 2022, *68* (3), 1678–1695.

Jarocinski, Marek and Michele Lenza, "An Inflation-Predicting Measure of the Output Gap in the Euro Area," *Journal of Money, Credit and Banking*, 2018, *50* (6), 1189–1224.

Joseph, Andreas, Galina Potjagailo, Chiranjit Chakraborty, and George Kapetanios, "Forecasting UK inflation bottom up," *International Journal of Forecasting*, 2024, *40* (4), 1521–1538.

Kang, Byoung Uk, Francis In, and Tong Suk Kim, "Timescale betas and the cross section of equity returns: Framework, application, and implications for interpreting the Fama-French factors," *Journal of Empirical Finance*, 2017, *42*, 15–39.

Kelly, Bryan and Seth Pruitt, "Market Expectations in the Cross-Section of Present Values," *Journal of Finance*, 2013, *68* (5), 1721–1756.

_ **and** _ , "The three-pass regression filter: A new approach to forecasting using many predictors,"

Journal of Econometrics, 2015, *186* (2), 294–316.

Kilponen, Juha and Fabio Verona, "Investment dynamics and forecast: Mind the frequency," *Finance Research Letters*, 2022, *49* (C).

Koop, Gary and Dimitris Korobilis, "Forecasting inflation using dynamic model averaging," *International Economic Review*, 2012, *53* (3), 867–886.

Lin, Hai, Chunchi Wu, and Guofu Zhou, "Forecasting Corporate Bond Returns with a Large Set of Predictors: An Iterated Combination Approach," *Management Science*, 2018, *64* (9), 4218–4238.

Ludvigson, Sydney C. and Serena Ng, "The empirical risk-return relation: A factor analysis approach," *Journal of Financial Economics*, 2007, *83* (1), 171–222.

Mackowiak, Bartosz and Mirko Wiederholt, "Business Cycle Dynamics under Rational Inattention," *Review of Economic Studies*, 2015, *82* (4), 1502–1532.

Mankiw, N. Gregory and Ricardo Reis, "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics*, 2002, *117* (4), 1295–1328.

Martinez-Martin, Jaime, Richard Morris, Luca Onorante, and Fabio Massimo Piersanti, "Merging Structural and Reduced-Form Models for Forecasting," *The B.E. Journal of Macroeconomics*, 2024, *24* (1), 399–437.

Martins, Manuel M.F. and Fabio Verona, "Inflation dynamics in the frequency domain," *Economics Letters*, 2023, *231*, 111304.

_ **and** _ , "Forecasting inflation with the new Keynesian Phillips curve: frequencies matter," *Oxford Bulletin of Economics and Statistics*, 2024, *86* (4), 811–832.

Medeiros, Marcelo C., Gabriel F. R. Vasconcelos, Alvaro Veiga, and Eduardo Zilberman, "Forecasting Inflation in a Data-Rich Environment: The Benefits of Machine Learning Methods," *Journal of Business & Economic Statistics*, 2021, *39* (1), 98–119.

Naghi, Andrea A., Eoghan O'Neill, and Martina Danielova Zaharieva, "The benefits of forecasting inflation with machine learning: New evidence," *Journal of Applied Econometrics*, 2024, *39* (7), 1321–1331.

Percival, Donald B. and Andrew T. Walden, *Wavelet Methods for Time Series Analysis*, Cambridge

University Press, 2000.

Poledna, Sebastian, Michael Gregor Miess, Cars Hommes, and Katrin Rabitsch, "Economic forecasting with an agent-based model," *European Economic Review*, 2023, *151*, 104306.

Rapach, David E., Jack K. Strauss, and Guofu Zhou, "Out-of-Sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy," *Review of Financial Studies*, 2010, *23* (2), 821–862.

Reis, Ricardo, "Inattentive Consumers," *Journal of Monetary Economics*, 2006, *53* (8), 1761–1800.

_ , "Inattentive Producers," *Review of Economic Studies*, 2006, *73* (3), 793–821.

Rossi, Barbara, "Inflation Forecasting," in Guido Ascari and Riccardo Trezzi, eds., *Research Handbook on Inflation*, Edward Elgar Publishing, 2025, chapter 20, pp. 440–474.

Rua, Antonio, "A wavelet approach for factor-augmented forecasting," *Journal of Forecasting*, 2011, *30* (7), 666–678.

_ , "A wavelet-based multivariate multiscale approach for forecasting," *International Journal of Forecasting*, 2017, *33* (3), 581–590.

Sahm, Claudia, "Direct Stimulus Payments to Individuals," in Heather Boushey, Ryan Nunn, and Jay Shambaugh, eds., *Recession Ready: Fiscal Policies to Stabilize the American Economy*, Brookings Institution and The Hamilton Project, 2019, chapter 3, pp. 15–38.

Sims, Christopher A., "Implications of Rational Inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665–690.

Stein, Tobias, "Forecasting the equity premium with frequency-decomposed technical indicators," *International Journal of Forecasting*, 2024, *40* (1), 6–28.

Stock, James H. and Mark W. Watson, "Forecasting inflation," *Journal of Monetary Economics*, 1999, *44* (2), 293–335.

_ **and** _ , "Macroeconomic Forecasting Using Diffusion Indexes," *Journal of Business & Economic Statistics*, 2002, *20* (2), 147–162.

_ **and** _ , "Forecasting Output and Inflation: The Role of Asset Prices," *Journal of Economic Literature*, 2003, *41* (3), 788–829.

_ **and** _ , "Combination forecasts of output growth in a seven-country data set," *Journal of Forecasting*, 2004, *23* (6), 405–430.

_ **and** _ , "Why Has U.S. Inflation Become Harder to Forecast?," *Journal of Money, Credit and Banking*, 2007, *39* (s1), 3–33.

_ **and** _ , "Core Inflation and Trend Inflation," *Review of Economics and Statistics*, 2016, *98* (4), 770–784.

Theil, Henri, *Applied Economic Forecasting*, North-Holland, Amsterdam, 1971.

Tibshirani, Robert, "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society: Series B (Methodological)*, 1996, *58* (1), 267–288.

Verbrugge, Randal J., "Inflation's Last Half Mile: Higher for Longer?," *Economic Commentary*, 2024, (2024-09), 1–8.

Verona, Fabio, "Investment Dynamics with Information Costs," *Journal of Money, Credit and Banking*, 2014, *46* (8), 1627–1656.

West, Kenneth D, "Asymptotic Inference about Predictive Ability," *Econometrica*, 1996, *64* (5), 1067–1084.

Wu, Jing Cynthia and Fan Dora Xia, "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," *Journal of Money, Credit and Banking*, 2016, *48* (2-3), 253–291.

Yellen, Janet, "Macroeconomic Research After the Crisis," 2016. speech at the 60th Annual Economic Conference sponsored by the Federal Reserve Bank of Boston, October 14, 2016.

Zou, Hui and Trevor Hastie, "Regularization and Variable Selection via the Elastic Net," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2005, *67* (2), 301–320.

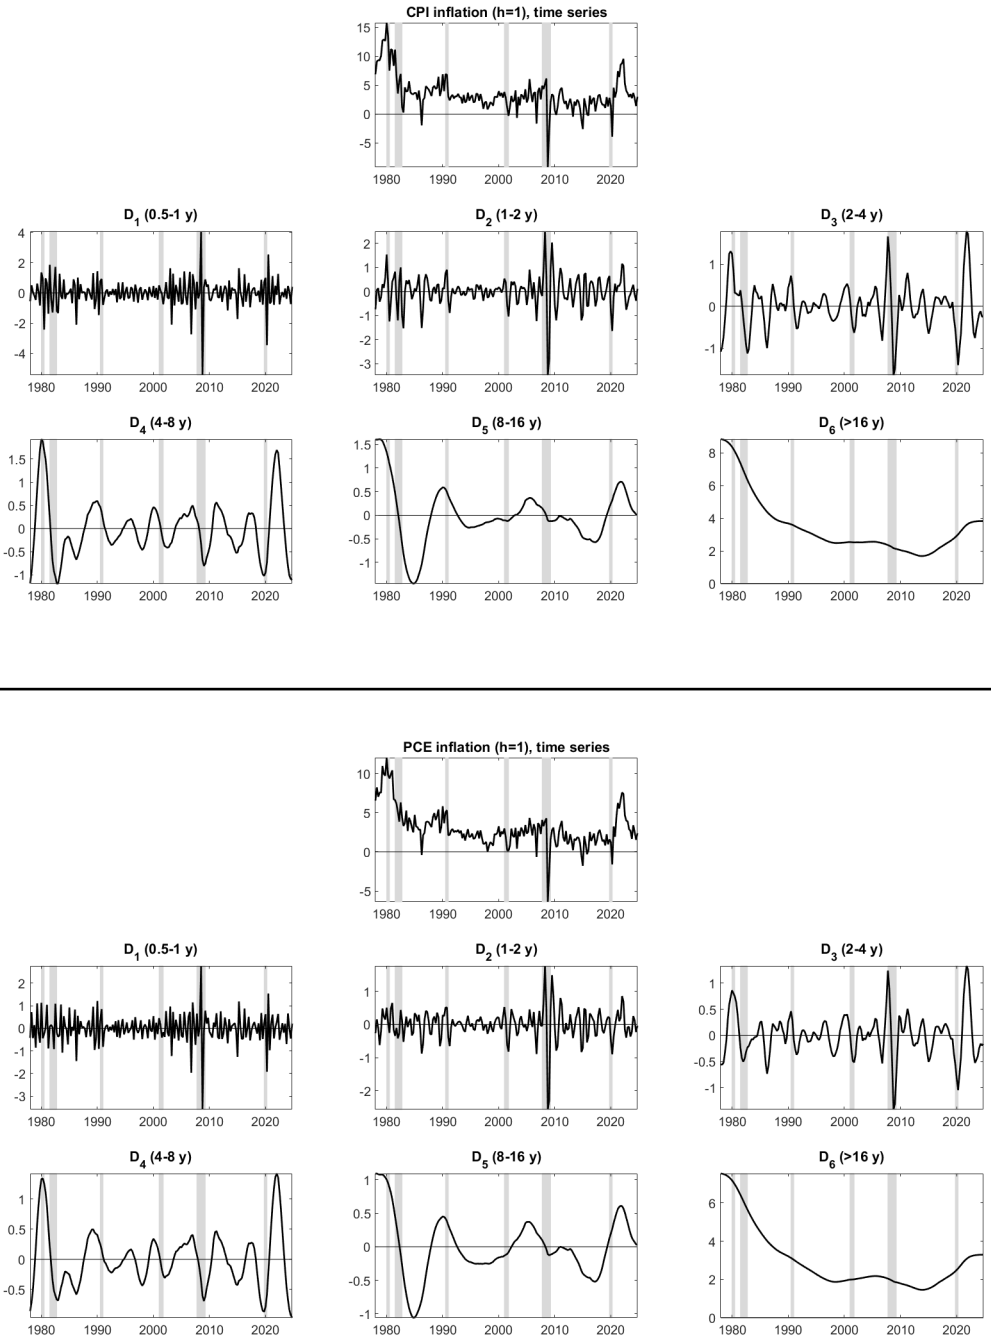

Figure 1: U.S. one-quarter inflation and frequency components

Notes: The figure shows annualized quarterly inflation rates (upper graphs) and their six frequency components obtained from the Haar filter (lower graphs). Upper panel: CPI inflation. Lower panel: PCE inflation. Sample period: 1978:Q1-2024:Q4. Shaded areas denote NBER recessions.

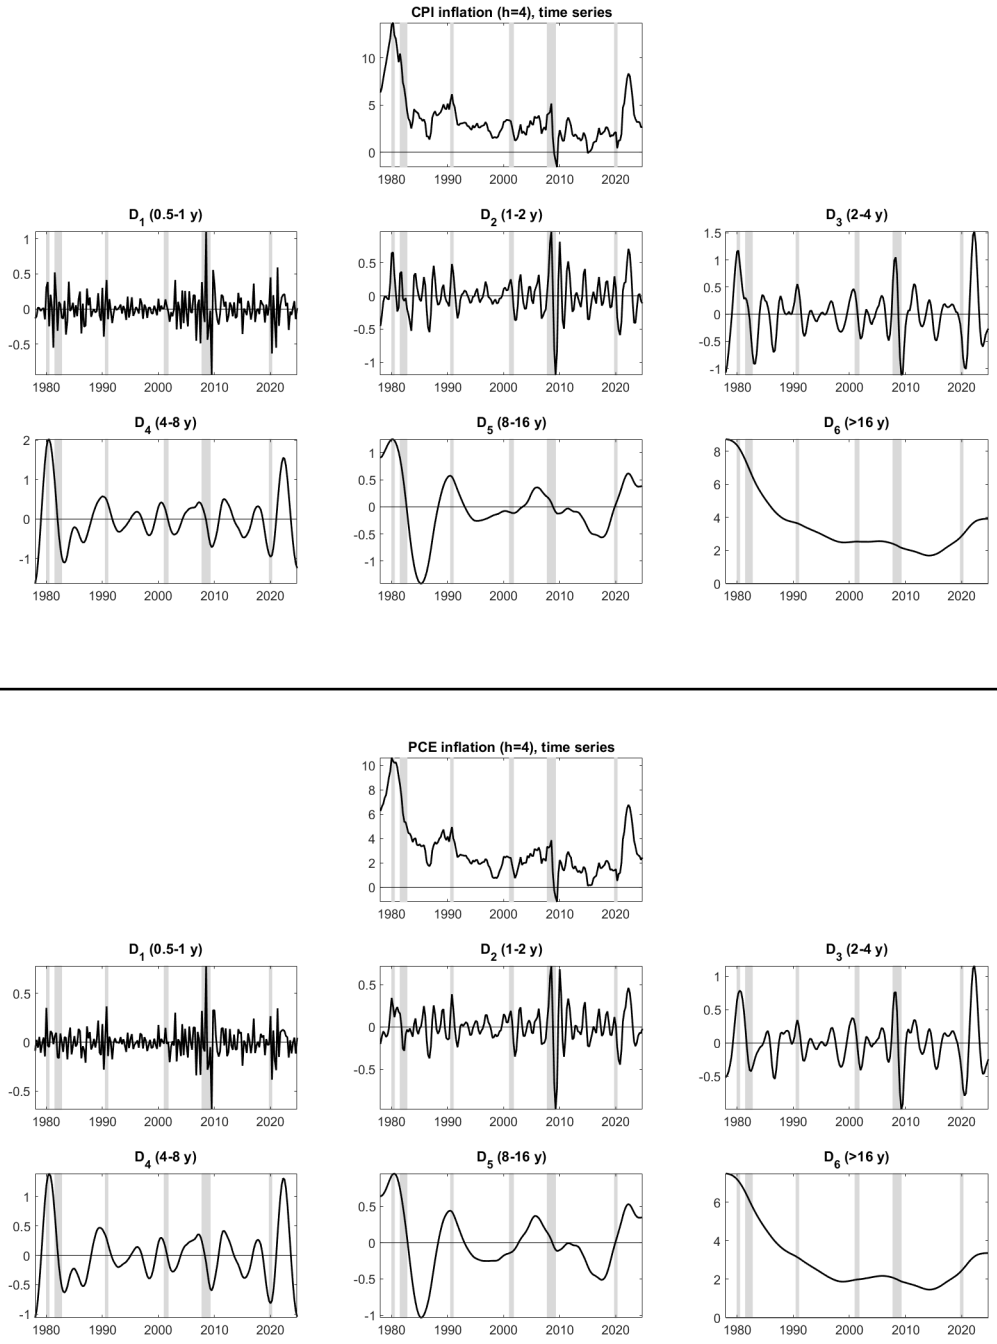

Figure 2: U.S. one-year average inflation and frequency components

Notes: The figure shows annualized one-year average inflation rates (upper graphs) and their six frequency components obtained from the Haar filter (lower graphs). Upper panel: CPI inflation. Lower panel: PCE inflation. Sample period: 1978:Q1-2024:Q4. Shaded areas denote NBER recessions.

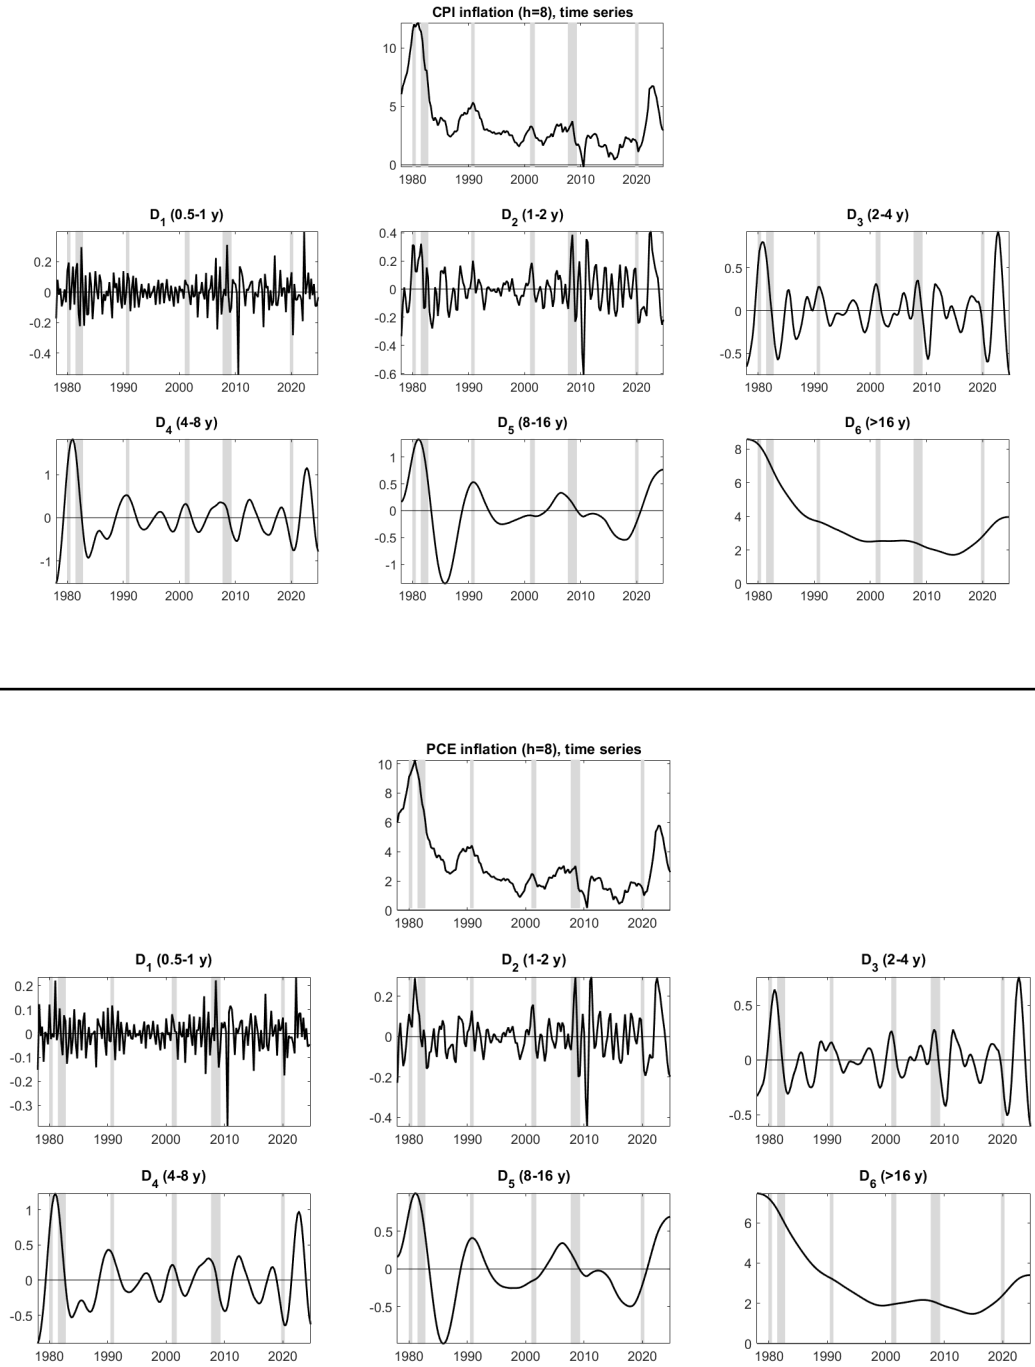

Figure 3: U.S. two-year average inflation and frequency components

Notes: The figure shows annualized two-year average inflation rates (upper graphs) and their six frequency components obtained from the Haar filter (lower graphs). Upper panel: CPI inflation. Lower panel: PCE inflation. Sample period: 1978:Q1-2024:Q4. Shaded areas denote NBER recessions.

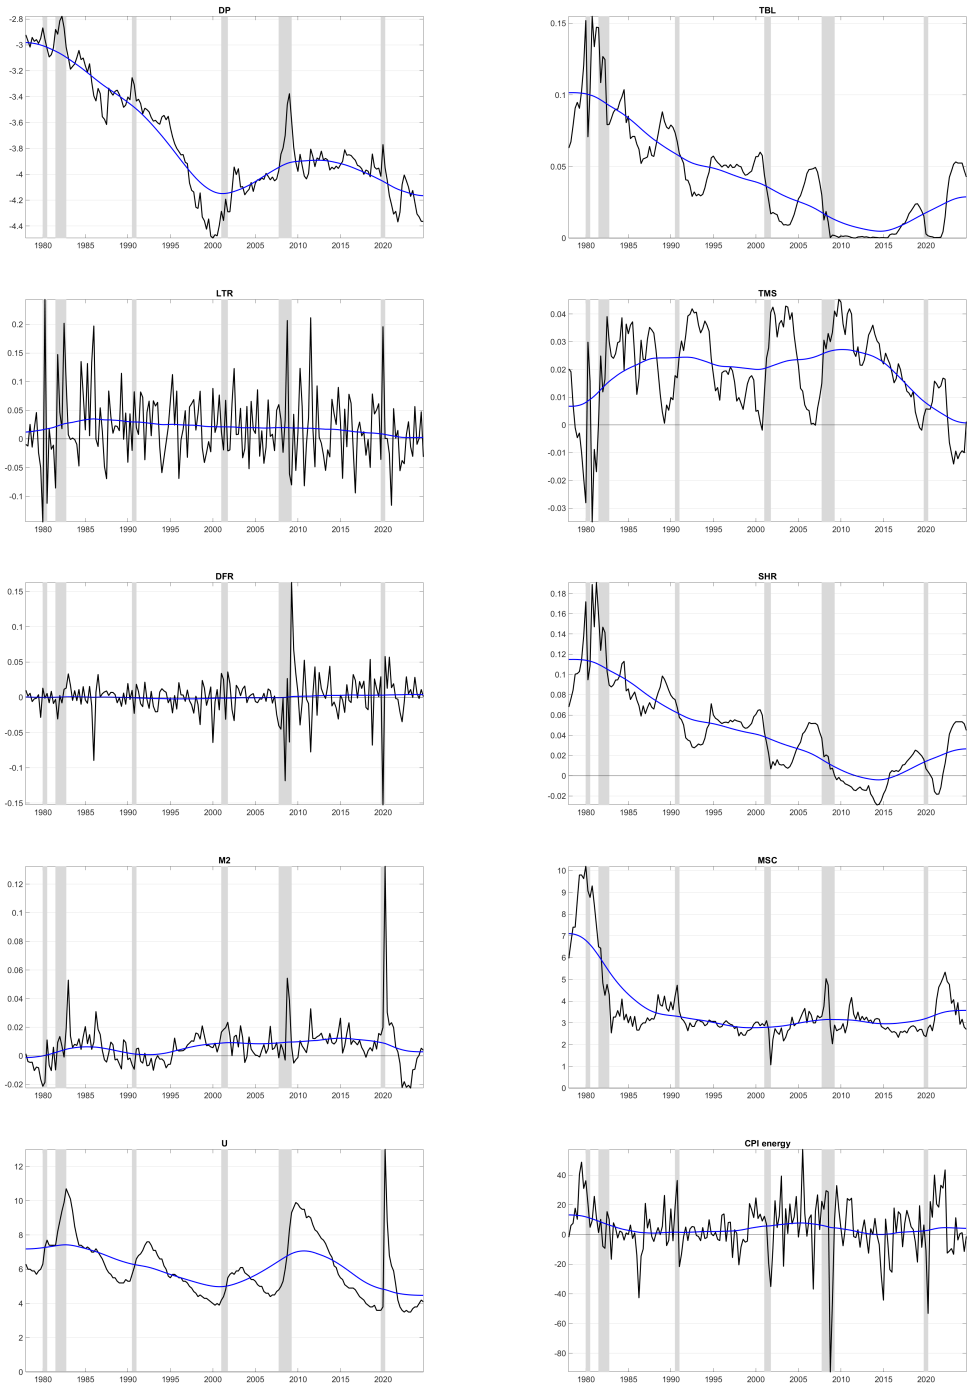

Figure 4: Predictors of inflation and their long-run components

Notes: The figure plots the 20 predictors used in the forecasting analysis (black lines) together with their long-run components (blue lines). The long-run components are extracted using the Haar filter and capture fluctuations with periodicities longer than 16 years (D_6). Sample period: 1978:Q1-2024:Q4. Shaded areas denote NBER recessions.

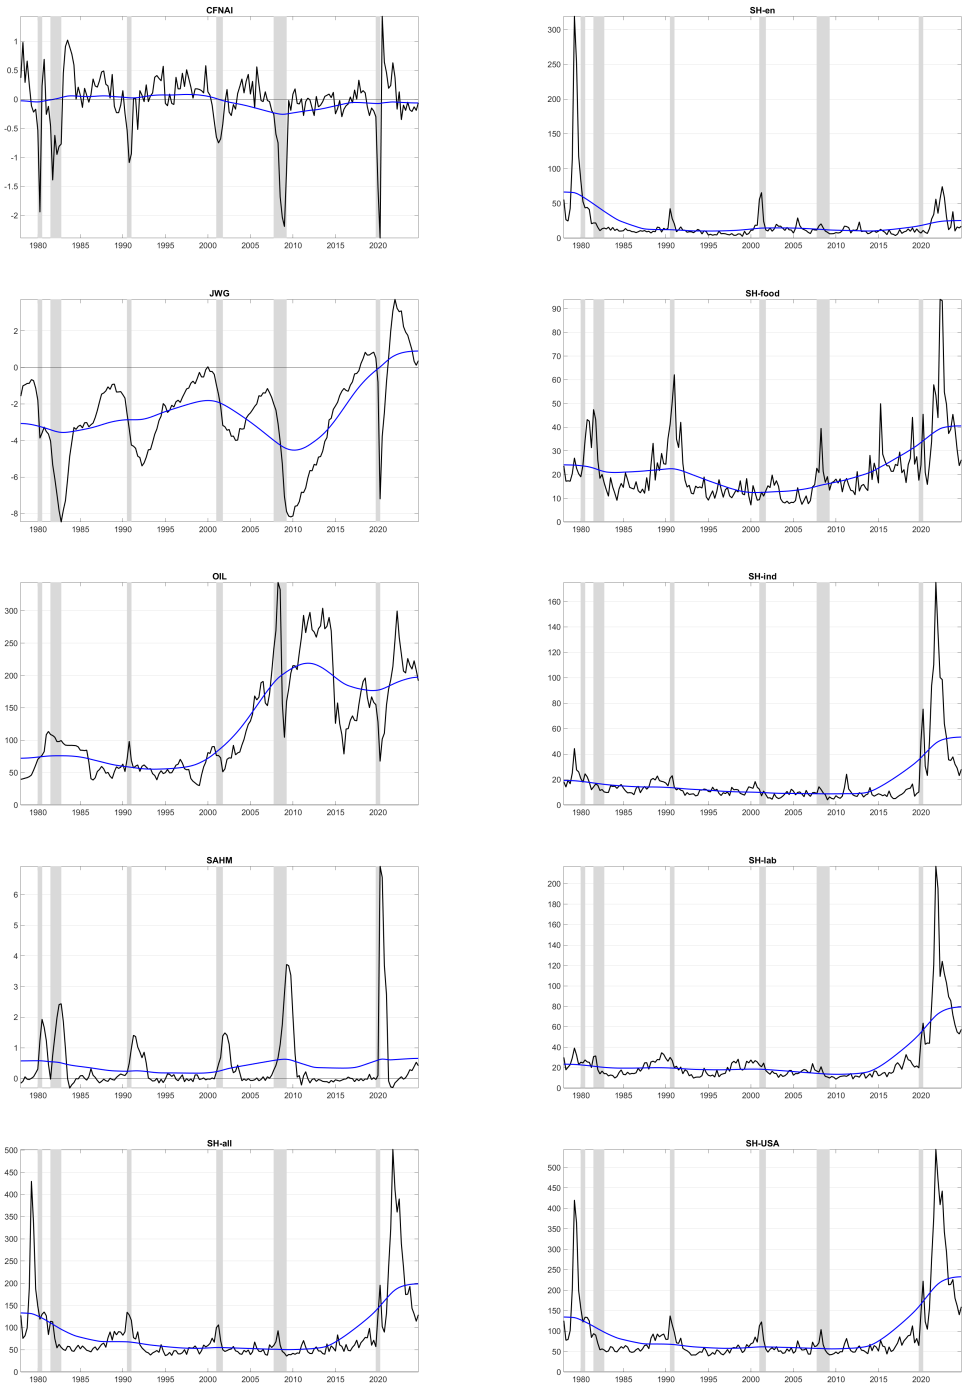

Figure 4 - cont

Notes: The figure plots the 20 predictors used in the forecasting analysis (black lines) together with their long-run components (blue lines). The long-run components are extracted using the Haar filter and capture fluctuations with periodicities longer than 16 years (D_6). Sample period: 1978:Q1-2024:Q4. Shaded areas denote NBER recessions.

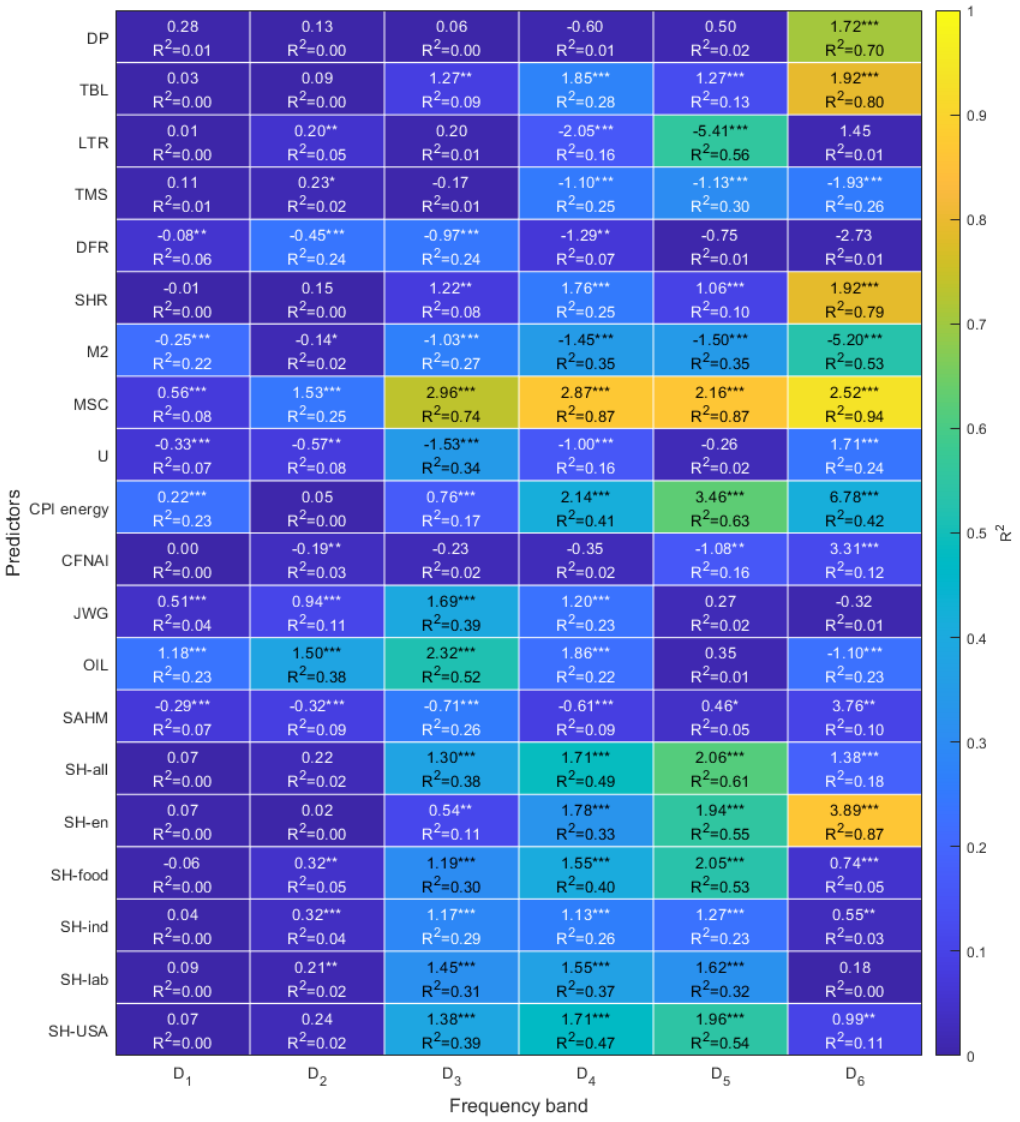

Figure 5: Bivariate regressions of frequency components of CPI inflation on predictors (h=4)

Notes: U.S. data, 1978:Q1-2024:Q4. Each cell corresponds to a bivariate OLS regression of the frequency component of CPI inflation (one-year average, h=4) on the corresponding component of each predictor. The color indicates the explanatory power (R^2) of the regression, with darker (blue) cells representing lower R^2 and lighter (yellow) cells higher R^2. Numerical entries report the estimated coefficient (first line) and the R^2 value (second line), along with significance levels indicated by asterisks: 10 % (*), 5 % (**), and 1 % (***). Frequency definitions as in table 3.

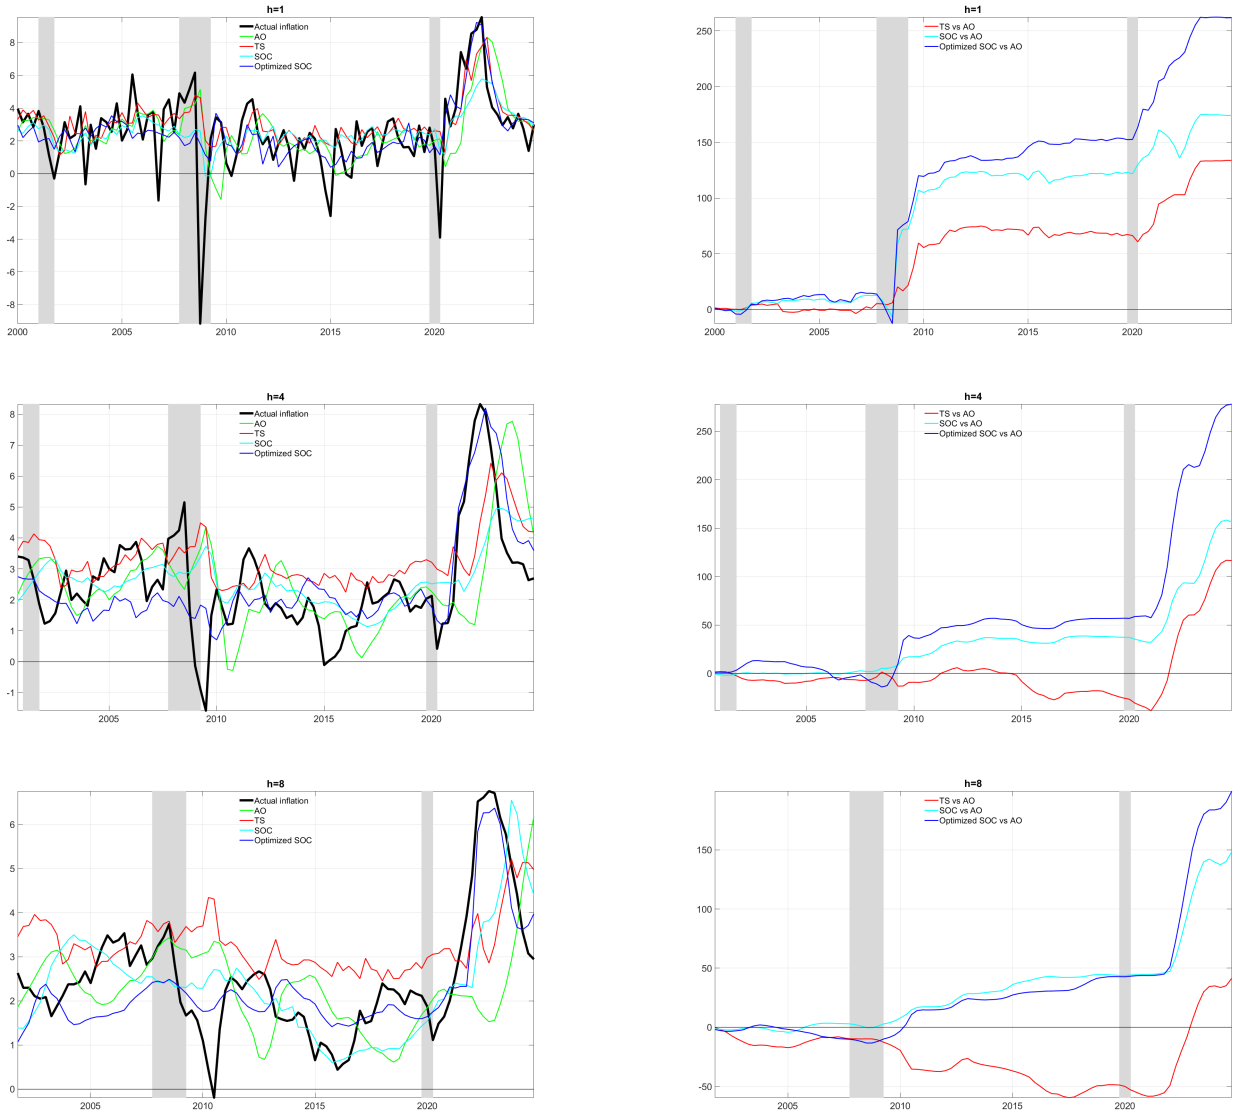

Figure 6: CPI inflation forecasts and cumulative forecast error differences

Notes: The left panels plot realized CPI inflation (black lines) together with forecasts from the AO benchmark (green), the best-performing time-series (TS) model (red), the baseline SOC method (cyan), and the optimized SOC method (blue). The right panels plot the cumulative differences in squared forecast errors relative to the AO model (red: TS-AO, cyan: baseline SOC-AO, blue: optimized SOC-AO). Rows correspond to forecast horizons: one quarter (h=1, 2000:Q1-2024:Q4), one year (h=4, 2000:Q4-2024:Q4), and two years (h=8, 2001:Q4-2024:Q4). Shaded areas denote NBER recessions.

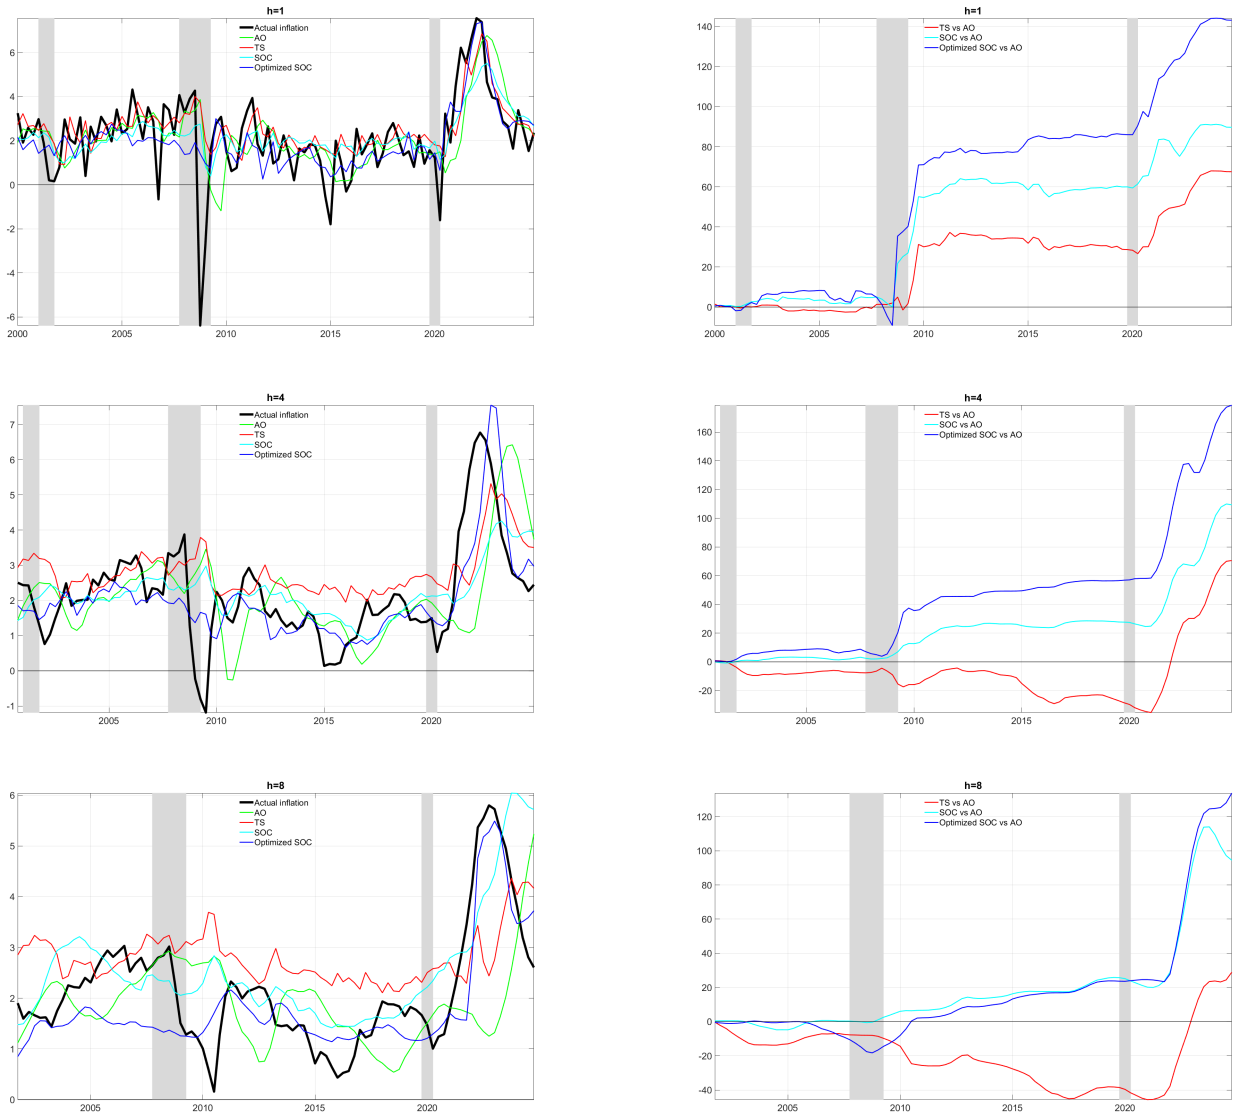

Figure 7: PCE inflation forecasts and cumulative forecast error differences

Notes: The left panels plot realized PCE inflation (black lines) together with forecasts from the AO benchmark (green), the best-performing time-series (TS) model (red), the baseline SOC method (cyan), and the optimized SOC method (blue). The right panels plot the cumulative differences in squared forecast errors relative to the AO model (red: TS-AO, cyan: baseline SOC-AO, blue: optimized SOC-AO). Rows correspond to forecast horizons: one quarter (h=1, 2000:Q1-2024:Q4), one year (h=4, 2000:Q4-2024:Q4), and two years (h=8, 2001:Q4-2024:Q4). Shaded areas denote NBER recessions.

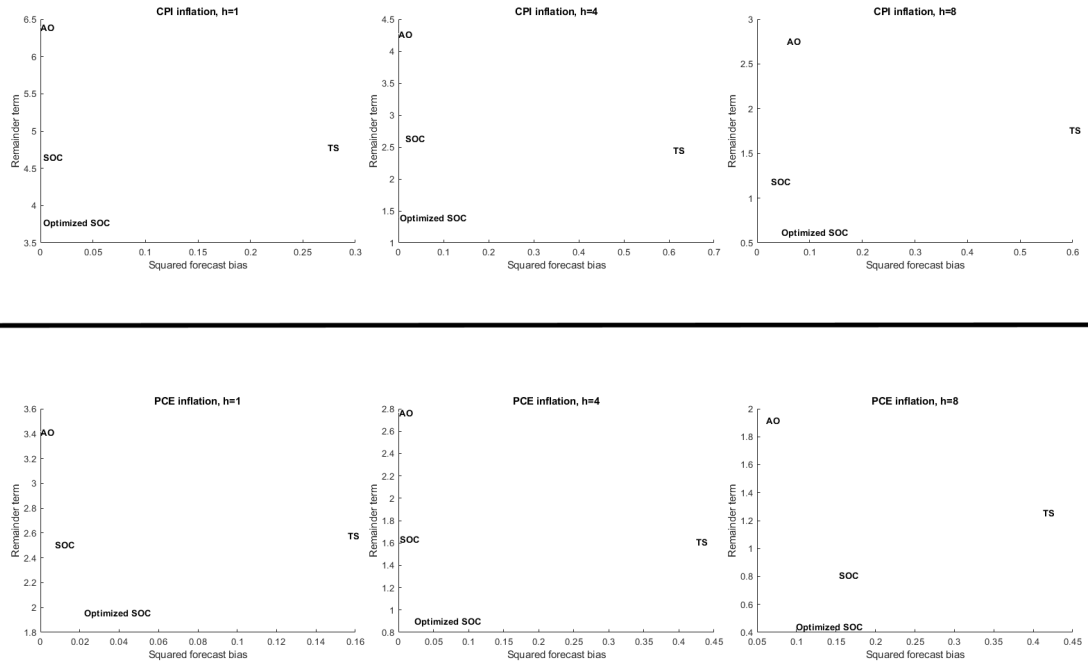

Figure 8: Bias-variance decomposition of forecast errors

Notes: This figure shows the Theil (1971) mean squared forecast error (MSFE) decomposition into squared forecast bias (x-axis) and a remainder term reflecting efficiency (y-axis). Points closer to the origin indicate better overall performance. Upper panels: CPI inflation. Lower panels: PCE inflation. First column: h=1. Second column: h=4. Third column: h=8.

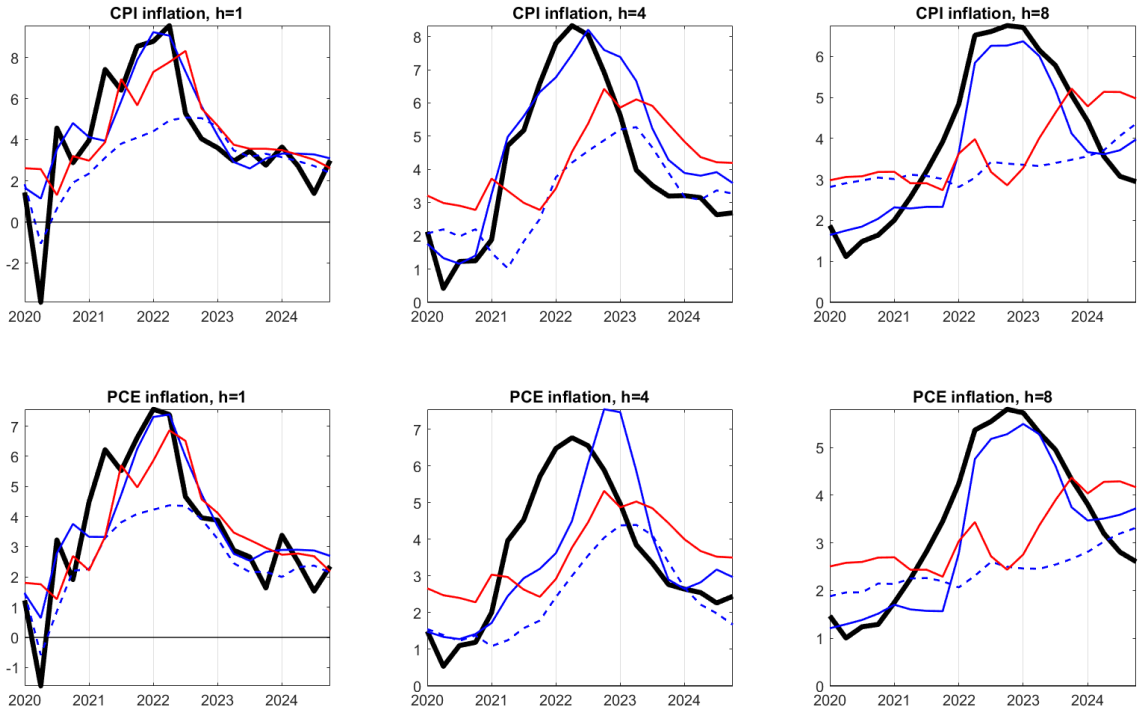

Figure 9: The role of shortage variables in the 2020-21 inflation surge

Notes: This figure plots realized inflation (black lines), together with forecasts from the optimized SOC method including shortage predictors (blue solid lines) and excluding them (blue dashed lines), and with the best-performing time-series (TS) model (red lines). Upper panels: CPI inflation. Lower panels: PCE inflation. Columns correspond to forecast horizons of one quarter (h=1), one year (h=4), and two years (h=8). Sample period: 2020:Q1-2024:Q4.

		D_1	D_2	D_3	D_4	D_5	D_6
h=1	D_1	1.25					
	D_2	1.17	1.27				
	D_3	1.16	1.19	1.30			
	D_4	1.17	1.19	1.24	1.32		
	D_5	1.21	1.23	1.26	1.29	1.36	
	D_6	0.66**	0.69**	0.75**	0.79***	0.87*	0.90
h=4	D_1	1.45					
	D_2	1.41	1.42				
	D_3	1.37	1.35	1.39			
	D_4	1.36	1.34	1.33	1.38		
	D_5	1.40	1.37	1.34	1.34	1.42	
	D_6	0.63***	0.57***	0.48***	0.49***	0.63***	0.68***
h=8	D_1	1.70					
	D_2	1.68	1.68				
	D_3	1.65	1.64	1.66			
	D_4	1.63	1.62	1.61	1.63		
	D_5	1.63	1.62	1.59	1.57	1.64	
	D_6	0.57**	0.53**	0.45**	0.39**	0.51**	0.60**

Table 1: Potential forecasting gains from perfectly predicting frequency components of CPI inflation

Notes: U.S. data, 1978:Q1-2024:Q4. Entries report the root mean squared forecast error (RMSE) of CPI inflation relative to the Atkeson-Ohanian (2001) random walk (AO) benchmark at horizons of one quarter (h=1), one year (h=4), and two years (h=8). Forecasts assume perfect knowledge of the indicated frequency components (D_1-D_6, from highest to lowest frequency). Diagonal entries show results for each component in isolation; off-diagonal entries add components cumulatively. Asterisks denote Diebold-Mariano-West significance at 10 % (*), 5 % (**), and 1 % (***).

	mean	median	standard deviation	minumum	maximum	skewness	kurtosis	AR(1)
Panel A) Target variables								
CPI inflation (h=1)	3.50	3.05	3.06	-9.16	15.8	0.84	6.75	0.70
CPI inflation (h=4)	3.52	2.88	2.63	-1.59	13.7	1.77	6.47	0.96
CPI inflation (h=8)	3.56	2.82	2.46	-0.19	12.2	1.86	6.38	0.98
PCE inflation (h=1)	2.99	2.56	2.44	-6.39	12.0	0.90	5.91	0.78
PCE inflation (h=4)	3.02	2.44	2.20	-1.19	10.7	1.55	5.51	0.97
PCE inflation (h=8)	3.06	2.42	2.10	0.16	10.3	1.61	5.24	0.99
Panel B) Predictors								
DP	-3.75	-3.88	-4.48	-2.83	0.43	0.48	2.24	0.98
TBL	0.04	0.04	0.00	0.15	0.04	0.81	3.43	0.96
LTR	0.02	0.01	-0.11	0.21	0.06	0.76	4.58	-0.06
TMS	0.02	0.02	-0.02	0.04	0.02	-0.56	3.06	0.84
DFR	0.00	0.00	-0.11	0.06	0.03	-0.45	13.2	-0.11
SHR	0.04	0.05	-0.03	0.18	0.04	0.73	3.68	0.96
M2	0.01	0.01	-0.02	0.05	0.01	3.50	30.2	0.48
MSC	3.58	3.07	2.08	9.80	1.60	2.54	9.26	0.95
U	6.05	5.70	3.50	10.6	1.77	0.86	3.60	0.89
ENERGY	4.02	3.89	-50.3	46.8	18.3	-0.89	7.42	0.28
CFNAI	-0.04	0.01	-2.13	1.00	0.52	-1.61	8.24	0.57
JWG	-2.50	-2.15	-8.17	3.18	2.50	-0.25	3.04	0.96
OIL	122	92.1	32.4	321	78.2	0.91	2.71	0.96
SAHM	0.40	0.02	-0.28	5.48	1.01	3.71	20.0	0.79
SH-all	83.5	57.0	37.0	422	74.0	3.28	14.7	0.89
SH-en	19.2	11.1	3.85	201	32.5	6.56	54.0	0.78
SH-food	21.4	17.3	7.55	81.5	13.4	2.40	11.1	0.78
SH-ind	17.3	11.5	4.60	126	21.7	4.40	25.4	0.90
SH-lab	25.6	17.6	9.25	168	27.4	4.21	24.2	0.91
SH-USA	90.9	61.7	41.3	460	82.8	3.22	13.9	0.91

Table 2: Descriptive statistics of inflation measures and predictors

Notes: U.S. data, 1978:Q1-2024:Q4. Panel A reports summary statistics for CPI and PCE inflation at horizons of one quarter (h=1), one year (h=4), and two years (h=8). Panel B reports descriptive statistics for the 20 predictors described in appendix A. Statistics include mean, median, standard deviation, minimum, maximum, skewness, kurtosis, and first-order autocorrelation (AR(1)). All variables are expressed in quarterly frequency; see appendix A for definitions and transformations.

	D_1	D_2	D_3	D_4	D_5	D_6
CPI inflation (h=1)	15	12	10	11	11	**41**
CPI inflation (h=4)	2	5	10	14	13	**56**
CPI inflation (h=8)	1	2	6	13	14	**64**
PCE inflation (h=1)	11	9	9	11	11	**50**
PCE inflation (h=4)	2	4	8	12	12	**63**
PCE inflation (h=8)	1	2	5	10	13	**70**
DP	1	1	2	4	7	**84**
TBL	2	2	4	8	8	**75**
LTR	**53**	25	12	5	2	3
TMS	8	9	13	21	19	**30**
DFR	**55**	23	14	5	1	1
SHR	2	2	4	8	9	**74**
M2	**26**	20	16	17	11	11
MSC	2	3	6	12	16	**59**
U	5	7	11	18	24	**36**
ENERGY	**36**	28	17	10	5	5
CFNAI	21	21	**22**	19	10	6
JWG	2	4	10	18	24	**41**
OIL	2	4	6	8	11	**69**
SAHM	11	19	**27**	23	15	5
SH-all	6	10	14	17	14	**40**
SH-en	11	19	20	13	12	**25**
SH-food	11	11	13	16	12	**37**
SH-ind	5	9	14	19	15	**38**
SH-lab	4	8	10	15	13	**49**
SH-USA	5	9	13	16	14	**44**

Table 3: Variance decomposition of inflation measures and predictors by frequency

Notes: U.S. data, 1978:Q1-2024:Q4. Entries report the percentage of variance of each series explained by frequency components D_1-D_6. D_1 captures fluctuations with a period of 2-4 quarters; D_2-D_5 correspond to cycles of 1-2, 2-4, 4-8, and 8-16 years; D_6 captures fluctuations longer than 16 years. Percentages may not sum to 100 due to rounding. Bold entries highlight the dominant frequency component for each variable.

Inflation measure	CPI			PCE		
Forecasting horizon	h=1	h=4	h=8	h=1	h=4	h=8
Panel A) Benchmark						
AO	2.53	2.07	1.68	1.85	1.66	1.41
Panel B) TS and baseline SOC						
TS vs AO	0.889***	0.848*	0.917	0.896***	0.859	0.918
Baseline SOC vs AO	0.853**	0.789***	0.658**	0.858***	0.770***	0.699*
Baseline SOC vs TS	0.959	0.931	0.717***	0.959	0.897*	0.761**
Panel C) Optimized SOC						
Optimized SOC vs AO	0.769***	0.574**	0.511**	0.762***	0.582***	0.536**
Optimized SOC vs TS	0.864*	0.677**	0.558***	0.851*	0.677***	0.583***

Table 4: Out-of-sample root mean squared forecast errors

Notes: U.S. data. Panel A reports RMSEs of the Atkeson-Ohanian (2001) random walk (AO) benchmark. Panel B reports RMSEs of the aggregate time-series (TS) model and baseline SOC method, expressed relative to AO. Panel C reports results for the optimized SOC method. Forecast horizons are h=1 (2000:Q1-2024:Q4), h=4 (2000:Q4-2024:Q4), and h=8 (2001:Q4-2024:Q4). Asterisks denote Diebold-Mariano-West significance at 10 % (*), 5 % (**), and 1 % (***).

CPI							
	TS	D_1	D_2	D_3	D_4	D_5	D_6
h=1	C-DMSPE 0.25	PLS-1	C-DMSPE 0.5	C-DMSPE 0.5	C-DMSPE 1	DP	C-DMSPE 1
h=4	C-MEAN	M2	SH-USA	DP	DP	PC3	C-DMSPE 1
h=8	C-MEAN	DFR	DFR	SH-en	DP	PC3	PCA
PCE							
	TS	D_1	D_2	D_3	D_4	D_5	D_6
h=1	C-DMSPE 0.25	PLS-1	C-DMSPE 0.5	PLS-2	C-DMSPE 0.75	MSC	C-DMSPE 1
h=4	C-MEAN	PLS-2	LTR	DP	DP	PC3	C-DMSPE 1
h=8	C-MEAN	DFR	DP	DP	DP	PC3	SH-all

Table 5: Best forecasting model for each frequency component of inflation

Notes: The table reports, for each forecast horizon, the model with the lowest RMSE for the aggregate series and for each frequency component (D_1-D_6) of CPI and PCE inflation. Frequency bands ordered from high (D_1) to low (D_6).

CPI				
	M1	M2	M3	M4
h=1	DFR^{D_2}	SH-ind^{D_3}	LTR^{D_4}	PCA^{D_6}
h=4	SAHM^{D_2}	SH-lab^{D_3}	LTR^{D_4}	SH-ind^{D_6}
h=8	TMS^{D_3}	TMS^{D_4}	SAHM^{D_5}	SH-ind^{D_6}
PCE				
	M1	M2	M3	M4
h=1	DFR^{D_2}	SH-ind^{D_3}	LTR^{D_4}	PCA^{D_6}
h=4	DP^{D_2}	SH-lab^{D_3}	LTR^{D_4}	PCA^{D_6}
h=8	RIDGE^{D_2}	TMS^{D_4}	SAHM^{D_5}	PCA^{D_6}

Table 6: Individual component forecasts included in the optimized SOC method

Notes: For each inflation measure (CPI, PCE) and forecast horizon (h=1,4,8), the table reports the frequency-component forecasts selected for inclusion in the optimized SOC method. See section 5.2 for details of the optimization procedure.

Inflation measure	CPI			PCE		
Forecasting horizon	h=1	h=4	h=8	h=1	h=4	h=8
M4	0.870**	0.667**	0.580*	0.882*	0.653***	0.618**
M4 + M3	0.829**	0.621**	0.549*	0.825**	0.619***	0.591**
M4 + M2	0.832**	0.613**	0.612*	0.839**	0.605***	0.616**
M4 + M1	0.849**	0.645**	0.566*	0.864*	0.645***	0.610**
M4 + M3 + M2	0.805***	0.589**	0.517*	0.794***	0.590***	0.543**
M4 + M3 + M1	0.807**	0.603**	0.524*	0.805**	0.612***	0.580**
M4 + M2 + M1	0.798***	0.592**	0.616*	0.809**	0.596***	0.612**
M4 + M3 + M2 + M1	0.769***	0.574**	0.511**	0.762***	0.582***	0.536**

Table 7: Contribution of component forecasts included in the optimized SOC method

Notes: For each inflation measure (CPI, PCE) and forecast horizon (h=1,4,8), the table reports the individual frequency-component forecasts included in the optimized SOC method aggregation (as reported in Table 6) and their incremental contribution to forecast accuracy relative to the Atkeson-Ohanian (2001) random-walk (AO) benchmark. The final row shows the performance of the full optimized SOC method. Asterisks denote Diebold-Mariano-West significance at 10 % (*), 5 % (**), and 1 % (***).

Appendix A - Description of predictors

This appendix describes the predictors used in the empirical analyses.

Monetary and financial indicators. These variables capture monetary policy stance, asset pricing conditions, and term/risk premia that may feed into inflation.

- Real **M2** growth rate: real M2 money stock, deflated by CPI (source: FRED).
- Treasury bill rate (**TBL**): 3-month Treasury bill secondary market rate (source: Goyal, Welch and Zafirov, 2024).
- Shadow rate (**SHR**): Wu-Xia shadow federal funds rate, capturing monetary policy stance at the zero lower bound (source: Wu and Xia, 2016).
- Term spread (**TMS**): difference between the 10-year Treasury constant maturity yield and the 3-month Treasury bill rate (source: Goyal et al., 2024).
- Long-term bond return (**LTR**): return on long-term U.S. government bonds (source: Goyal et al., 2024).
- Default return spread (**DFR**): difference between returns on long-term corporate bonds and long-term government bonds returns (source: Goyal et al., 2024).
- Dividend-price ratio (**DP**): difference between the log of dividends (12-month moving sums of dividends paid on the S&P 500) and the log of prices (the S&P 500 index) (source: Goyal et al., 2024).

Expectations. These indicators reflect households' beliefs about future inflation.

- **MSC**: median one-year-ahead inflation expectations from the University of Michigan Survey of Consumers (source: FRED).

Real activity and slack. These measures proxy cyclical conditions and labor-market tightness.

- Unemployment (**U**): civilian unemployment rate, seasonally adjusted (source: FRED).
- Jobs-workers gap (**JWG**): difference between the total number of jobs (i.e. employment plus job openings) and the number of workers (i.e. the labor force) (source: Hatzius, 2024).
- Sahm rule (**SAHM**): difference between the three-month moving average of the unemployment rate and its minimum over the previous 12 months (source: Sahm, 2019).
- **CFNAI**: Chicago Fed National Activity Index, weighted average of 85 indicators of real activity and inflationary pressure (source: Chicago Fed website).

Commodity and energy indicators. These capture volatile but important cost-push forces.

- CPI energy (**ENERGY**): energy subcomponent of CPI, seasonally adjusted (source: FRED).
- **OIL**: log difference of West Texas Intermediate (WTI) crude oil price (source: FRED).

Shortage indices. These indices are constructed from automated text analysis of major U.S. newspapers, capturing real-time signals of bottlenecks and supply disruptions (source: Caldara et al., 2025).

- **SH-all**: aggregate news-based shortage index.
- **SH-ind**: industry-specific shortage index.
- **SH-en**: energy shortage index.
- **SH-food**: food shortage index.
- **SH-lab**: labor shortage index.
- **SH-USA**: shortage index for the U.S. economy.

Bank of Finland Research Discussion Papers 2026

ISSN 1456-6184, online

1/2026 Fabio Verona

Forecasting inflation: The sum of the cycles outperforms the whole