

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Measuring Inflation Shock Momentum

Kevin J. Lansing and Adam Hale Shapiro
Federal Reserve Bank of San Francisco

April 2026

Working Paper 2026-10

<https://doi.org/10.24148/wp2026-10>

Suggested citation:

Lansing, Kevin J. and Adam Hale Shapiro. 2026. “Measuring Inflation Shock Momentum.”
Federal Reserve Bank of San Francisco Working Paper 2026-10
<https://doi.org/10.24148/wp2026-10>

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the European Central Bank, the Federal Reserve Bank of San Francisco, or the Federal Reserve System.

Measuring Inflation Shock Momentum*

Kevin J. Lansing[†] Adam Hale Shapiro[‡]

April 28, 2026

Abstract

We develop a non-parametric filter that identifies sustained directional runs in shocks to monthly inflation—a concept we define as “inflation shock momentum.” By assessing the shocks to over 100 disaggregated Personal Consumption Expenditures (PCE) inflation categories, we isolate the share of categories experiencing positive or negative inflation shock momentum in a given month. We define the “Inflation Shock Momentum” (ISM) index as the *net* positive momentum share of expenditure-weighted categories (positive minus negative) in a given month. We show that the ISM index helps to forecast aggregate PCE inflation at horizons of 1 to 3 years, even after controlling for a variety of other inflation predictor variables. The ISM index is particularly useful in capturing emerging disinflationary pressure and can be used to help forecast future inflation movements in real time.

JEL classification: E31, E37, E52, C14, C53

Keywords: PCE Inflation, Non-parametric filter, Forecasting

*The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco, or the Board of Governors of the Federal Reserve System. We thank

[†]Federal Reserve Bank of San Francisco, kevin.j.lansing@sf.frb.org

[‡]Federal Reserve Bank of San Francisco, adam.shapiro@sf.frb.org

1 Introduction

In the wake of the pandemic recession, 12-month PCE inflation rose from a trough of 0.46 percent in April 2020 to a peak of 7.2 percent in June 2022. Inflation began to exceed the Federal Reserve’s 2 percent goal in March 2021, but the federal funds rate remained near zero until March 2022.¹ During this one-year interval, five consecutive Federal Open Market Committee (FOMC) statements from April 2021 through November 2021 characterized inflation developments as “largely reflecting transitory factors” or “factors that are expected to be transitory.” But in testimony before the U.S. Senate Committee on Banking, Housing, and Urban Affairs on November 30, 2021, Fed Chair Jerome Powell stated “We tend to use [the word transitory] to mean that it won’t leave a permanent mark in the form of higher inflation. I think it’s probably a good time to retire that word...Clearly, the risk of more persistent inflation has risen.” But over four years later in February 2026, just prior to the start of the Middle East conflict, the 12-month PCE inflation rate remained above the Federal Reserve’s goal at 2.8 percent.

The pandemic-era inflation episode highlights the difficulty faced by policymakers in distinguishing in real time between inflation movements driven by transitory versus persistent shocks (Lansing (2022)). Delayed recognition of persistent inflationary pressures increases the risk that inflation expectations will become de-anchored, requiring a stronger and more prolonged policy response to restore price stability. In contrast, premature policy tightening in response to transitory shocks poses the risk of unnecessary output or employment losses. Early and reliable indicators of shifts in the underlying inflation environment can potentially help improve monetary policy decisions. Our aim is to construct a quantitative measure of persistent inflationary or disinflationary pressure—which we term inflation shock momentum (ISM).

We propose a simple, non-parametric filter that identifies sustained directional runs in innovations to monthly inflation. The filter is based on the premise that as long as the assumed data generating process (DGP) for inflation remains valid, then monthly deviations from the DGP (i.e., the innovations) should exhibit no sustained pattern of positive or negative runs. However, if some aspect of the DGP has recently changed, as in a shift to a regime of higher (lower) average inflation, then one would expect to observe a sustained run of positive (negative) deviations from the original DGP.

We implement our filter at the category-level of inflation, employing the more than 100 Personal Consumption Expenditures (PCE) categories that form the basis of the FOMC’s

¹The median policy path from the FOMC’s March 17, 2021 Summary of Economic Projections indicated that the federal funds rate would remain near zero through the end of 2023.

preferred measure of price inflation. Category-level data increases the density of information available in any given month, allowing us to extract a shock momentum signal, including its direction, using only a narrow window of recent data. For each expenditure category, we estimate a rolling 10-year auto-regressive model to describe a benchmark DGP for category inflation. We then identify sequences of three consecutive months in which category inflation deviates positively or negatively from the benchmark DGP. By focusing on consecutive monthly deviations, the momentum index seeks to capture recent persistent inflationary or disinflationary pressure. Our ISM index is the difference between the expenditure-weighted shares of categories exhibiting consecutive positive versus consecutive negative deviations.

We take a non-parametric approach for two reasons. First, typical parametric estimates of persistence are inherently “sign agnostic.” Specifically, a parametric estimate of persistence does not distinguish between a regime of persistent positive shocks (that would push inflation higher) versus a regime of persistent negative shocks (that would push inflation lower). By contrast, our measure is sign-sensitive—it distinguishes between persistent positive shocks that place upward pressure on inflation and persistent negative shocks that signal disinflationary forces. Second, parametric estimates of persistence represent an *average* value over the entire rolling-window sample period. Given that inflation shock persistence likely varies over time, persistence readings computed using data many months prior to the current month provide less information about current shock persistence. With over 100 category-level observations per month, we can identify coordinated directional pressure using only a short window of recent data.

We find that the ISM index helps to forecast aggregate 12-month PCE inflation at horizons of one to three years, even after controlling for a variety of macroeconomic and financial predictor variables. These include: 1-year ahead household inflation expectations, the vacancy-unemployment ratio, the slope of the Treasury yield curve, oil prices, equity prices, the growth rate of real personal disposable income, and recession indicators. For in-sample forecasts of 12-month PCE inflation with controls, the ISM index is highly statistically significant and improves the adjusted R-squared statistic at forecast horizons of 12 and 24 months. For out-of-sample forecasts with controls, inclusion of the ISM index serves to reduce the root mean squared forecast error (RMSFE) by statistically significant amounts at forecast horizons of 12, 24, and 36 months. Taken together, our results indicate that disaggregated inflation data contains timely information about future aggregate inflation that is not fully captured by standard predictor variables.

Finally, we run an exercise to validate the empirical relevance of the ISM index. We do so by assessing the dynamic response of the index to two well-identified exogenous disturbances: (1) monetary policy attempted disinflations identified by Romer and Romer (2024) and (2)

news shocks about future oil supply identified by Känzig (2021). This exercise serves as a diagnostic test for whether the ISM index captures the impacts of demand or supply shocks on inflation in a manner consistent with economic theory. We provide evidence that the ISM index exhibits a significant and persistent decline following an attempted disinflation. By contrast, negative news about future oil supply yields a transitory and significant rise in the ISM index.

Our study contributes to the growing literature suggesting that aggregate inflation measures often mask important information contained within the cross-sectional distribution of price changes. Category-level data can reflect heterogeneous cost structures and markups across firms or industries (Chiavari et al. (2025)), differences in inflation persistence and volatility across goods or services (Giacomini and Levin (2023)), the relative influence of demand versus supply forces (Shapiro (2024), Eickmeier and Hofmann (2022)), differences in exposure to business cycles (Mahedy and Shapiro (2017), Zaman (2019), Stock and Watson (2020)), differences in exposure to trade policy (Hobijn and Nechio (2025)), or other category-specific factors. These differences may provide early signals of broad-based inflationary or disinflationary pressures before such pressures are fully reflected in aggregate data.

This analysis relates to the extensive applied econometrics literature modeling inflation dynamics. Studies such as Stock and Watson (2007) focus on decomposing inflation into transitory and persistent components, where the persistent component captures the underlying trend. Later studies, including one by the same authors (Stock and Watson (2016)), use the underlying category data to help measure the trend component. Other studies using disaggregated data include a recent paper by Bonaparte et al. (2025) who develop a measure of transitory inflation that captures persistent deviations from trend inflation.² Our measure differs from these types of transitory-trend decompositions in two important ways: (1) our focus on very recent innovations to inflation and (2) differentiating between upward and downward inflation pressure.

Finally, our study builds upon the literature identifying structural breaks in time series forecasting (for example, Giacomini and Rossi (2009)). Recent research, such as Cavallo and Zavaleta (2025), use disaggregated price indexes to detect breaks in the underlying inflation trend. The ISM index’s ability to forecast up to 36 months ahead indicates that it captures persistent structural breaks in the inflation process rather than just temporary fluctuations. This long-horizon predictive power suggests that the ISM index identifies fundamental shifts

²Other related studies using disaggregated data include Quineche and Zapata (2026) who employ disaggregated PCE inflation data to construct directional inflation pressure measures and Verona (2026) who decomposes quarterly inflation into six separate cyclical frequency components and then constructs forecasting models for each component.

in the inflation process which traditional models often fail to detect.

2 Constructing the Inflation Momentum Index

2.1 Methodology

We develop an Inflation Shock Momentum (ISM) index designed to identify coordinated directional pressure across the distribution of category-level inflation rates. Our approach focuses on the realized sign of inflation shocks, where “shocks” are measured as residuals from a simple first-order autoregressive time series model, AR(1), estimated on monthly inflation.

To motivate our approach, we postulate that monthly inflation π_t is governed by the following time series model that allows for a distinction between intrinsic inflation inertia and the persistence of inflation shocks.

$$\pi_t = \mu + \rho\pi_{t-1} + \varepsilon_t \tag{1}$$

$$\varepsilon_t = \alpha\varepsilon_{t-1} + u_t. \tag{2}$$

where $\rho\pi_{t-1}$ represents intrinsic inflation inertia (analogous to the mass of a large ship), ε_t represents a shock to inflation in month t (analogous to thrust imparted to turn the ship), and α measures the persistence of the shock (analogous to the duration of the thrust). If $\alpha = 0$, shocks tend to die out immediately. If $\alpha > 0$, then inflation shocks are self reinforcing—a positive (negative) shock to inflation in month t tends to be followed by another positive (negative) shock in month $t + 1$, $t + 2$, etc. It is in this sense that α can be viewed as a measure of inflation shock momentum.³

We do not directly estimate α for several reasons.⁴ First, non-linear time-series models are not well-identified; parameter estimates are highly sensitive to starting values and estimates often do not converge.⁵ Second, any estimate of α is sign-agnostic, i.e., the estimated value

³Time series momentum is often defined as the autocorrelation of the first difference (or growth rate) of a random variable (Huang et al. (2020)). If the data generating process for inflation is given by equations (1) and (2), then the time series momentum of inflation is given by $Corr(\Delta\pi_t, \Delta\pi_{t-1}) = (\rho + \alpha + \rho\alpha - 1)/2$, which is increasing in both ρ and α .

⁴There are at least two ways to extract an estimate of α . One could jointly estimate ρ and α using non-linear least squares (NLS) using an AR(2) model for monthly inflation—combining equations (1) and (2) implies a restricted AR(2) of the form $\pi_t = \delta + (\rho + \alpha)\pi_{t-1} - (\rho\alpha)\pi_{t-2} + u_t$ where $\delta = \mu(1 - \alpha)$. Alternatively, one could employ a two-step estimation procedure using OLS: perform an OLS regression on (1), and then using those residuals perform an OLS regression on (2). Appendix Figure A1 plots the 10-year rolling window estimated values $\hat{\alpha}$ and $\hat{\rho}$ obtained from this two-step procedure. Appendix Table A6 shows in-sample inflation forecasts using the corresponding $\hat{\alpha}$ measure.

⁵See, for example, Smith Jr (1993)

measures persistence of the inflation shock but does not distinguish between a sustained sequence of *positive* shocks (that would push inflation higher) versus a sustained sequence of *negative* shocks (that would push inflation lower). Third, any estimate of α represents an average value over the *entire* rolling window sample period. While the rolling window estimation would capture some measure of time-varying shock persistence, the pattern of shocks in the most recent months provides a more informative signal about recent changes in the underlying DGP for inflation. Estimating α using a relatively short rolling window, say 3 to 5 years, would not resolve this problem because these windows would still include many “stale” inflation observations.

Based on the above considerations, we measure inflation shock momentum using a non-parametric approach that: (1) distinguishes between positive versus negative shock momentum, and (2) restricts attention to recent shocks. In so doing, we exploit the information in category-level inflation data that underlie the aggregate PCE price index. By shifting to category-level data, we significantly increase the density of information available to gauge inflation behavior in any given month. This approach allows us to extract a shock momentum signal—including its direction—using only a narrow window of recent data. While the aggregate inflation series provides only one new observation each month, the category-level data provides over 100 new observations each month.⁶ The focus on recent shocks allows us to identify the early stages of an underlying shift in the DGP for inflation that would otherwise plausibly take longer to detect using a parametric estimate on aggregate inflation.

To construct the ISM index, we first estimate an AR(1) on each of the 129 separate categories, denoted by i , in the PCE index using 120-month rolling-window regressions:⁷ The rolling-window regressions take the form

$$\pi_{i,t} = \mu_i + \rho_i \pi_{i,t-1} + \varepsilon_{i,t} \tag{3}$$

where the coefficient ρ_i is allowed to vary over time, albeit at a low-frequency (120 months) to ensure estimation efficiency.⁸ The first window begins in February 1959 and ends in January 1969. The estimated AR(1) model generates residuals starting in January 1969. We then advance the window one month and repeat the estimation. We iterate this process

⁶Appendix Figure A5 compares the non-parametric ISM index using just the aggregate data—which can take on values of $\{-1.0, 1\}$ —and the category-level ISM index.

⁷The data on the underlying detail of quantity, price, and expenditure components of the PCE index are available in Tables 2.4.3U, 2.4.4U and 2.4.5U from the “Underlying Detail” page of the BEA’s website. The BEA constructs different levels of aggregation depending on the product category. We use the fourth level of disaggregation. For example, in the case of services, the disaggregation proceeds as follows: (1) services \rightarrow (2) transportation services \rightarrow (3) public transportation \rightarrow (4) air transportation. This breakdown results in 129 categories in the PCE price index going back to 1959.

⁸Appendix Figure A1 shows the distribution of rolling-window estimated values $\hat{\rho}_i$.

for each month until we reach the last window of data. For each rolling-window regression, the reduced-form residuals $\varepsilon_{i,t}$ are collected and then used to identify whether each category i exhibits positive or negative inflation shock momentum near the end date of the rolling window in month t .

We focus on the *recent* sequence of signed residuals for each inflation category. A run of same-signed shocks for a given category is a signal of increased shock persistence, i.e., an increase in α_i . Indeed, as shown in Appendix B and Appendix Figure A3, the probability of observing k consecutive same-signed residuals is monotonically increasing in α_i . By focusing on recent runs, our filter places greater weight on the most recent observations. This methodology enables us to identify shifts in recent shock patterns that influence the behavior of category-level inflation $\pi_{i,t}$, and by extension, aggregate inflation π_t . For our baseline measure, a category is said to exhibit “shock momentum” if the most recent AR(1) residuals share the same sign for at least $k = 3$ consecutive months. Our results are robust to changing the definition of shock momentum to represent $k = 2$ or $k = 4$ consecutive months of same-signed residuals.⁹

Importantly, the non-parametric filter allows us to classify shock momentum as either positive or negative depending on whether recent shocks put upward or downward pressure on inflation. We classify shock momentum as *positive* if the residuals are consecutively greater than zero and *negative* if they are consecutively less than zero. That is, for each category i , we construct the following inflation shock momentum signals:

$$M_{i,t}^+ = \prod_{k=0}^2 \mathbb{1}(\hat{\varepsilon}_{i,t-k} > 0) \quad (4)$$

$$M_{i,t}^- = \prod_{k=0}^2 \mathbb{1}(\hat{\varepsilon}_{i,t-k} < 0) \quad (5)$$

where $M_{i,t}^+$ is equal to 1 if category i experienced $k = 3$ consecutive positive residuals up to period t and zero otherwise. Analogously, $M_{i,t}^-$ is equal to 1 if category i experienced $k = 3$ consecutive negative residuals up to period t and zero otherwise. We examine alternative definitions of (4) and (5) which incorporate the *size* of the consecutive shocks, but this simply adds noise to the measures.¹⁰

Just under 20 percent of the total monthly inflation observations across 129 categories

⁹See Appendix Table A1 and Appendix Figure A6.

¹⁰See Appendix Table A3. The alternative definitions imply $M_{i,t}^+ = \prod_{k=0}^2 \mathbb{1}(\hat{\varepsilon}_{i,t-k} > 0) * |\hat{\varepsilon}_{i,t-k}|$ and $M_{i,t}^- = \prod_{k=0}^2 \mathbb{1}(\hat{\varepsilon}_{i,t-k} < 0) * |\hat{\varepsilon}_{i,t-k}|$, where $|\hat{\varepsilon}_{i,t-k}|$ is the absolute value of the residual. The in-sample fit for predicting 12-month ahead PCE inflation declines slightly, indicating that most of the important variation takes place on the extensive margin involving the *sign*, rather than the *size*, of the shock.

exhibit $k = 3$ consecutive positive residuals ($M_{i,t}^+ = 1$), while approximately 15 percent of total monthly inflation observations exhibit $k = 3$ consecutive negative residuals ($M_{i,t}^- = 1$).¹¹ Taking into account the category-level expenditure weights in the PCE price index, it follows that the shares of the PCE basket experiencing either positive or negative shock momentum in month t are:

$$S_t^+ = \sum_{i=1}^N \omega_{i,t} M_{i,t}^+ \quad (6)$$

$$S_t^- = \sum_{i=1}^N \omega_{i,t} M_{i,t}^- \quad (7)$$

where $\omega_{i,t}$ is the expenditure weight of category i in month t . Note that S_t^+ and S_t^- do not necessarily sum to one since many PCE categories do not experience $k = 3$ consecutive same-signed shocks in a given month. Use of the PCE basket shares allows for a continuous measure of the degree to which aggregate PCE inflation is experiencing positive or negative shock momentum. For example, the value $S_t^+ = 1$ would indicate that the entire PCE basket is experiencing positive shock momentum in month t . Such an outcome would be a strong signal that a persistent positive shock to aggregate inflation was underway. In contrast, if both $S_t^+ = 0.3$ and $S_t^- = 0.3$, then the PCE basket would be experiencing just as many expenditure-weighted categories with positive shock momentum as negative shock momentum, signaling no directional pressure on aggregate inflation. These ideas motivate the construction of the following diffusion index that measures the *net* share of the PCE basket that is experiencing positive versus negative shock momentum:

$$ISM_t = S_t^+ - S_t^- \quad (8)$$

2.2 The Inflation Shock Momentum Index

The upper panel of Figure 1 plots ISM_t , together with 12-month PCE inflation from January 1969 onward. A value $ISM_t > 0$ indicates that the PCE basket is experiencing a higher share of spending with positive shock momentum than negative shock momentum. The opposite is true when $ISM_t < 0$. Figure 1 shows that ISM_t is persistently above zero during the Great Inflation era of the late 1970s and early 80s. The index then declines sharply into

¹¹Appendix Figure A2 shows that these observation shares decrease to 12 percent and 7 percent, respectively, if we employ a stricter definition of shock momentum based on $k = 4$ consecutive months of same-signed residuals. The observation shares increase to 30 percent and 25 percent if shock momentum is defined as $k = 2$ consecutive months of same-signed residuals. For any given value of the shock persistence parameter α , Appendix Figure A3 shows that run probability declines as k increases.

negative territory during the disinflation episode that began under Fed Chair Paul Volcker (Goodfriend and King (2005)). ISM_t remains negative as the disinflation episode continued into the 1990s. PCE inflation declined from around 5 percent in 1990 to 2 percent in 1999. Indeed, the bottom panel of Figure 1 shows that the negative shock component of ISM_t was elevated throughout the decade. The late 1990s coincides with the time when the anchoring process for long-run expected inflation appears to have been completed.¹²

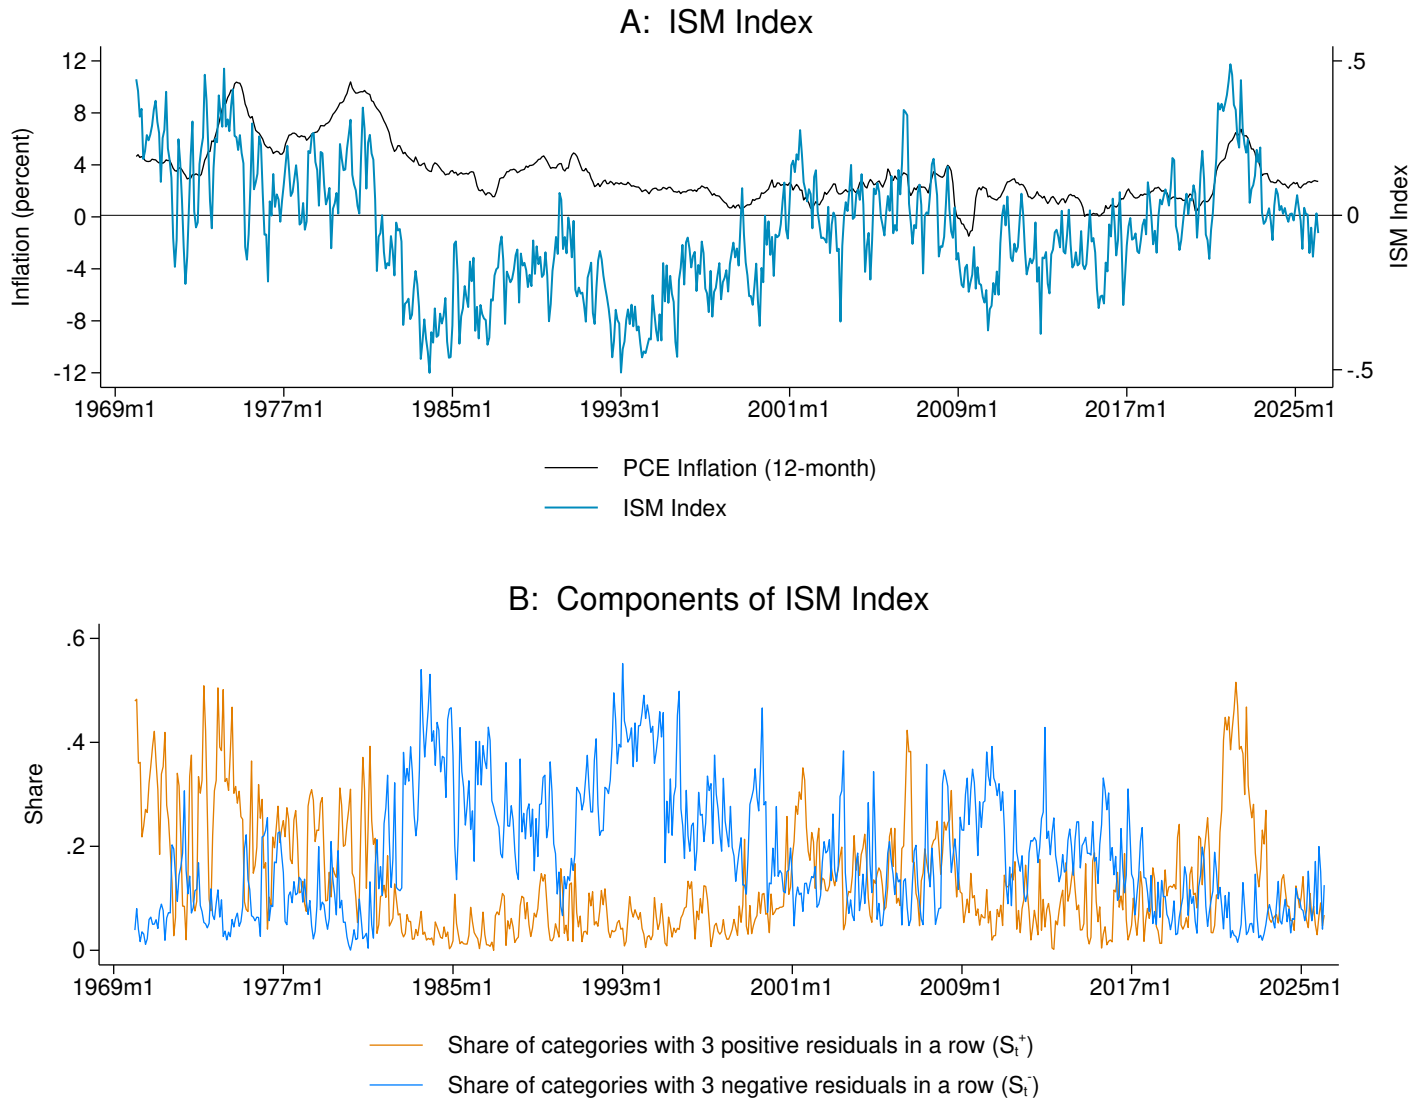
The mid-2000s exhibit low readings for both positive and negative components of ISM_t , which translates to a total index value that fluctuates around zero. This pattern coincides with a period when 12-month PCE inflation remained close to 2 percent. ISM_t declines sharply with the onset of the Great Recession in late 2007 and then remains mostly in negative territory for nearly a decade after the Great Recession ended in June 2009. During the period from June 2009 until the start of the pandemic recession in February 2020, 12-month PCE inflation was below the Fed’s 2 percent goal for 101 out of 129 months, or 78 percent of the time. ISM_t initially declines at the start of the pandemic recession in February 2020, but then reverses course and spikes upward during the period when PCE inflation surged to a 40-year high of 7.2 percent in June 2022. This trajectory is followed by a sharp decline in ISM_t as PCE inflation recedes after the June 2022 peak. At the end of our data sample in February 2026, ISM_t has once again returned to negative territory while PCE inflation remains above the Fed’s goal at 2.83 percent.¹³

To demonstrate the comparative advantage of using the category-level data to construct the ISM index, we implement a version of our non-parametric algorithm using only the aggregate PCE price index. That is, we estimate 120-month rolling-window regressions of (3) using only monthly aggregate PCE inflation. In this case, we observe only one set of inflation-shock momentum signals ($M_{PCE,t}^+$ and $M_{PCE,t}^-$) per month. It follows that this aggregate version the ISM index (the difference between $M_{PCE,t}^+$ and $M_{PCE,t}^-$) can only take on values of $\{-1, 0, 1\}$. Appendix Figure A5 plots this aggregate version of the ISM index against the index constructed using the category-level data. The correlation of the two indexes is 0.51, while the variance of the aggregate measure is noticeably larger due to its ternary structure. We show in Section 4 that the aggregate measure also performs worse in terms of forecasting future inflation than the category-level measure.

¹²See Bernanke (2007), Carvalho et al. (2023), and Jørgensen and Lansing (2025).

¹³Appendix Figure A4 reports ISM indices constructed for broad PCE basket categories (goods and non-housing services) and indices that employ alternative time series models, namely AR(3) or AR (12), for the benchmark rolling window regressions. The resulting ISM indices display somewhat different time-series patterns across goods and services—the correlation between the goods ISM index and the services ISM index is 0.48. However, they display similar patterns across benchmark regression models—the ISM index constructed using the AR(1) model has a correlation of 0.91 with the ISM index constructed using the AR(3) model, and a correlation of 0.75 with the ISM index constructed using the AR(12) model.

Figure 1: Inflation Shock Momentum Index and its Components



Notes: Panel A plots ISM_t together with the 12-month PCE inflation rate. Panel B plots the two components of $ISM_t = S_t^+ - S_t^-$: the PCE basket share of categories with $k = 3$ consecutive positive residuals S_t^+ (orange) and the PCE basket share of categories with $k = 3$ consecutive negative residuals S_t^- (blue).

3 Testing the Inflation Shock Momentum Index

In the next two sections, we examine the ability of the ISM index to improve forecasts of future inflation. As a first exercise, we test whether a shock to the ISM index has a significant impact on the future PCE price level, thereby impacting cumulative future inflation. To do so, we employ the local projection method of Jordà (2005). Specifically, we gauge how a surprise to the ISM index impacts the aggregate PCE price level over specific future horizons. The ISM surprise is constructed as the component of the ISM index that is orthogonal to the current and 12 lagged values of standard economic and financial predictor variables, as well as 12 lagged values of the ISM index itself. We run a distinct regression of the cumulative price change between periods $t - 1$ and $t + h$ for each forecast horizon h :

$$\ln(p_{t+h}) - \ln(p_{t-1}) = c + \beta_0^h ISM_t + \sum_{i=1}^{12} \beta_i^h ISM_{t-i} + \sum_{i=0}^{12} \mathbf{A}_i^h \mathbf{X}_{t-i} + \varepsilon_{t+h} \quad (9)$$


where the matrix \mathbf{X} includes 12-month PCE inflation, household inflation expectations, the vacancy-unemployment ratio V/U , the monthly average spot price per barrel of West Texas Intermediate crude oil, the monthly average value of the S&P 500 stock price index, the 12-month growth rate of real personal disposable income from the Bureau of Economic Analysis (BEA), the monthly average yield curve spread between 10-year Treasury bonds and the federal funds rate, and a monthly dummy variable indicator of NBER recessions. Inflation expectations are the median 1-year ahead household forecast from the University of Michigan Survey.¹⁴ The V/U ratio is measured using data from Barnichon (2010)¹⁵

The left panel of Figure 2 shows the estimates of β_0^h for forecast horizons h that range from 0 to 60 months after the ISM index perturbation. Specifically, β_0^h measures the impact of a 1 standard deviation (sd) increase—about 0.2 percentage points (pp)—in the ISM index on the cumulative increase in the PCE price level between period $t - 1$ and $t + h$. A 1 sd increase in the ISM index causes the PCE price level to increase by around 50 basis points (bp) after 30 months and around 100 bp after 60 months. The center and right panels of Figure 2 repeat the estimation of equation (9) but assess the impact of a shock to each separate component of the ISM index—that is, S_t^+ and S_t^- . A 1 sd increase in S_t^+ —about 0.1 pp—causes close to a 1 pp increase in the ISM index after 60 months. In contrast, the impact of a 1 sd increase in S_t^- —about 0.1 pp—is significantly larger in absolute value and

¹⁴The Michigan survey measure is available from January 1978 onward. Prior to this date, we proxy household inflation expectations using the 12-month trailing PCE inflation rate. The correlation coefficient between the Michigan survey measure and our proxy measure is 0.89 from January 1978 to February 2026.

¹⁵This series is constructed by splicing the Conference Board print newspaper help-wanted index with the Conference Board online help-wanted index.

Figure 2: Impulse Responses to ISM Index Surprises



Notes: The left panel shows the cumulative impact of a one standard deviation (sd) surprise of the ISM index on the log PCE price level over 60 months based on estimating equation (9) at each forecast horizon h . The middle panel shows the cumulative impact of a 1 sd shock of the positive ISM component, S_t^+ . The right panel shows the cumulative impact of a 1 sd shock of the negative ISM component, S_t^- . 1 sd and 90 percent confidence bands shown.

much less noisy than the impact of S_t^+ . The impulse response estimates confirm the patterns observed in Figure 1 which link disinflationary episodes to increased importance of negative shock momentum.

Robustness

Appendix Figures A7, A6, and A8 show that the impulse response function patterns in Figure 2 are robust to: (1) the number of lags in the benchmark autoregressive model used to compute the monthly inflation residuals (AR(1) versus AR(3) or AR(12)), (2) the number of consecutive same-signed residuals used to define the presence of momentum (two, three or four), and (3) the exclusion of data after December 2019 or before January 1990.

4 Forecasting Performance

In this section, we assess the ability of the ISM index to improve forecasts of future inflation, both in-sample and out-of-sample. The in-sample exercise has the advantage of being transparent and provides straightforward information about the additional explanatory power coming from the ISM index and its two components. The out-of-sample exercise

is more computationally intensive but provides a more realistic approximation of real-time forecasting performance.

4.1 In-sample forecasts

For the in-sample exercises, we estimate a parsimonious “baseline” model for forecasting inflation and then assess whether the inclusion of the ISM index in the baseline model improves forecast performance. We consider two baseline models. The first baseline model includes current 12-month PCE inflation and a constant. The second baseline model augments the first model with the current values of the control variables employed in previous section to construct the impulse response functions.¹⁶ The first six columns of Table 1 show the in-sample results when forecasting the 12-month-ahead value of 12-month PCE inflation. Columns 1 to 3 show results without the additional control variables (but still including current 12-month PCE inflation and a constant), while columns 4 to 6 show results with the additional control variables. Columns 7 through 12 show the in-sample results when forecasting the 24-month ahead value of 12-month PCE inflation.

Table 1 shows that including the ISM index serves to increase the adjusted R^2 statistic relative to the baseline model in all cases. Specifically, the adjusted R^2 increases from 0.588 to 0.595 between the simple baseline model (column 1) and the model augmented with the ISM index (column 3). For the baseline model with additional controls, including the ISM index increases the adjusted R^2 from 0.718 to 0.724 (column 4 versus column 6). While the improvements in fit are small, we consider any improvement to be value added given how difficult it is to forecast future inflation (Stock and Watson (2007), Faust and Wright (2013)). In this regard, the increases in the adjusted R^2 statistics when including the ISM index are larger when forecasting 24-month ahead inflation. For instance, the adjusted R^2 increases from 0.319 to 0.333 (column 7 versus column 9) when augmenting the simple baseline model with the ISM index.

Table 1 also shows that the negative component of the ISM index has somewhat stronger explanatory power for future inflation than the positive component. This can be seen from the results in columns 2, 5, 8, and 11 which show that the negative component of the ISM index is statistically significant in all four models, whereas the positive component is statistically significant in only 1 of 4 models. These results conform with the evidence from the impulse response functions in Figure 2 showing that the negative component of the ISM index has a more statistically significant impact on the future price level. The stronger significance of the negative component derives partially from the Volcker disinflation episode

¹⁶We also include 3-month PCE inflation as an additional control.

Table 1: In-sample Forecast Regressions

	12-month ahead forecast						24-month ahead forecast					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
PCE Inflation (y/y)	0.767*** (0.0287)	0.721*** (0.0312)	0.722*** (0.0301)	0.216*** (0.0626)	0.231*** (0.0651)	0.231*** (0.0650)	0.565*** (0.0337)	0.504*** (0.0379)	0.503*** (0.0368)	0.0155 (0.0662)	0.0372 (0.0709)	0.0391 (0.0707)
S_t^+		1.155 (0.813)			1.174* (0.692)			1.208 (0.954)			1.270 (0.858)	
S_t^-		-1.034** (0.504)			-1.219*** (0.459)			-1.763*** (0.642)			-3.678*** (0.651)	
ISM Index			1.090*** (0.284)			1.198*** (0.309)			1.504*** (0.349)			2.523*** (0.390)
N	674	674	674	674	674	674	662	662	662	662	662	662
R ²	0.588	0.596	0.596	0.722	0.728	0.728	0.320	0.335	0.335	0.522	0.553	0.551
Adjusted R ²	0.588	0.595	0.595	0.718	0.723	0.724	0.319	0.332	0.333	0.515	0.545	0.543
Controls	No	No	No	Yes	Yes	Yes	No	No	No	Yes	Yes	Yes

Notes: The dependent variable is the 12-month-ahead value of 12-month PCE inflation in columns 1 through 6 and the 24-month ahead value of 12-month inflation in columns 7 through 12. Controls are 3-month PCE inflation, 1-year ahead household inflation expectations, the V/U ratio, the price of oil, the S&P 500 stock price index, the 12-month growth rate of real personal disposable income, the yield curve spread between 10-year Treasury bonds and the federal funds rate, and a dummy variable indicator of recessions. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

when the negative component was persistently elevated as inflation declined substantially over a decade. Nevertheless, Appendix Figure A8 shows that the general pattern of the impulse response functions is robust to excluding the Volcker disinflation by starting the data sample in 1990. The results are also robust to excluding the pandemic-era inflation surge by ending the data sample in 2019.

Robustness

We perform a series of robustness tests for the in-sample results. Appendix Table A1 provides evidence of robustness with respect to the number of lags included in the benchmark autoregressive model and the number of consecutive same-signed residuals used to define the presence of inflation shock momentum. Results are similar across all specifications, showing that the predictive power of the ISM index is robust to the number of consecutive same sign-residuals used in the definition—either two, three, or four—and to the specification of the benchmark data generating process—either an AR(1), AR(3), or AR(12).

We next test alternative momentum signal measures. Our baseline definitions, equations (4) and (5), treat the “strength” of signals from all categories the same. But some inflation shock momentum signals may be more important than others in predicting future inflation trends. We incorporate two separate measures of “strength” into the momentum signals $M_{i,t}^+$

and $M_{i,t}^+$. The first measure is the size of the shock itself, intended to capture the possibility that larger shocks may have stronger impacts on future inflation:

$$M_{i,t}^+ = \prod_{k=0}^2 \mathbb{1}(\hat{\varepsilon}_{i,t-k} > 0) * \left| \sum_{k=0}^2 \hat{\varepsilon}_{i,t-k} \right| \quad (10)$$

$$M_{i,t}^- = \prod_{k=0}^2 \mathbb{1}(\hat{\varepsilon}_{i,t-k} < 0) * \left| \sum_{k=0}^2 \hat{\varepsilon}_{i,t-k} \right| \quad (11)$$

where $\left| \sum_{k=0}^2 \hat{\varepsilon}_{i,t-k} \right|$ represents the absolute value of the sum of the last three residuals. The second measure is the degree of “inflation stickiness” of category i , intended to capture the possibility that momentum from categories with more persistent inflation will have a stronger impact on future inflation:

$$M_{i,t}^+ = \prod_{k=0}^2 \mathbb{1}(\hat{\varepsilon}_{i,t-k} > 0) * \left(\frac{1}{1 - \hat{\rho}_{i,\tau}} \right) \quad (12)$$

$$M_{i,t}^- = \prod_{k=0}^2 \mathbb{1}(\hat{\varepsilon}_{i,t-k} < 0) * \left(\frac{1}{1 - \hat{\rho}_{i,\tau}} \right) \quad (13)$$

where $\hat{\rho}_{i,\tau}$ is the estimated inflation persistence parameter observed over the rolling-sample period $\tau \in (t - 120, t)$.¹⁷ In this weighting scheme, categories with more persistent inflation (i.e., higher values of $\hat{\rho}_{i,\tau}$) will receive a larger weight. Results of these two exercises are shown in the Appendix Tables A3 and A4. Comparing these results with the baseline estimates in Table 1 show slightly lower statistical significance of the coefficient on the ISM index and lower R^2 , especially at the 24-month ahead horizon. These results indicate that the important component of the ISM index (in terms of a correlation with future inflation) is the extensive margin (i.e., the indicator functions $\mathbb{1}(\cdot)$).

Next, to assess the importance of using category-level data, we report in-sample forecasting results using two alternative measures of the ISM index that employ *only* aggregate data. The first alternative is the aggregate non-parametric ISM index discussed at the end of Section 2.2 and shown in Appendix Figure A5. The in-sample forecast results using this aggregate data based ISM index are shown in Appendix Table A5. The resulting R^2 statistics are lower than the forecasts using the category-level ISM index, as is the statistical significance on the ISM index coefficient.¹⁸ In fact, the coefficient on the aggregate data based ISM index is not statistically significant in the 24-month ahead regressions.

¹⁷We impose a maximum ρ value of 0.9 so that the series remains bounded.

¹⁸For example, the t-statistic falls from 3.8 to 2.0 in the 12-month ahead specification with full controls (column 6).

The second alternative aggregate data based measure is a parametric estimate of $\hat{\alpha}$. To construct $\hat{\alpha}$ we take a two-step approach: For each 120-month rolling window, we run an AR(1) regression on monthly PCE inflation (i.e., equation (1)), and then run an AR(1) regression on the collected residuals (i.e., equation (2)). We report results in Appendix Table A6 using the parametric estimate of $\hat{\alpha}$ in place of our baseline ISM index. Results show a perverse sign using this alternative measure of inflation-shock momentum—that is, $\hat{\alpha}$ is *negatively* correlated with future inflation. Appendix Figure A1 verifies this finding, showing that $\hat{\alpha}$ is below zero during the high inflation period of the 1970s (while estimates of inflation inertia, $\hat{\rho}$, were elevated) and then rises above zero during the low-inflation period of 2010s. This exercise shows the difficulty of using only aggregate data to identify inflation shock momentum and to distinguish it from intrinsic inflation inertia

4.2 Out-of-sample forecasts

To assess the ability of the ISM index to improve out-of-sample inflation forecasts, we employ an adaptive LASSO (least absolute shrinkage and selection operator) estimator. The adaptive LASSO model, which is a weighted variation of the general LASSO model, is defined as:

$$\hat{\beta} = \arg \min \sum_{i=1}^n \left(\pi_{i+h} - \sum_{j=1}^p x_{i,j} b_j \right)^2 + \lambda \sum_{j=1}^p w_j |b_j| \quad (14)$$

where π_{t+h} is forecasted 12-month inflation at horizon h , $x_{i,j}$ represents the set of predictors, and w_j is a known weights vector such that $\hat{w} = 1/|\hat{\beta}|$.

The optimal penalty parameter λ^* is chosen based on a cross-validation procedure. We choose λ^* such that it yields the minimum out-of-sample root mean squared forecast error (RMSFE) across a series of 10-year rolling window regressions. Specifically, we perform a grid search over possible λ values and for each value we estimate the adaptive LASSO model for each 120-month rolling window (starting from the initial month of January 1969) and then collect the h -month ahead forecast error. For each possible λ , we then calculate the RMSFE across the collection of 10-year rolling window regressions. The optimal value λ^* is the λ yielding the minimum RMSFE.

Table A2a reports the results of the out-of-sample forecasting exercise for three different forecast horizons: $h = 12$, $h = 24$, and $h = 36$. We show results using the same two baseline models employed for the in-sample forecasting exercise, namely current 12-month PCE inflation alone (top panel), and current 12-month PCE inflation plus additional controls (bottom panel). All regressions include a constant term. We show the RMSFE computed

across the collection of 10-year rolling window regressions using $\lambda = \lambda^*$ and the Giacomini and White (2006) F-test statistic, which tests whether the two models being compared have equal conditional predictive ability. We also include standardized estimates of the regression coefficients assigned to the included predictors when estimated over the entire sample period using the optimal λ^* . The magnitude of the standardized regression coefficients provides information about the relative importance of the predictors.

Across all three forecast horizons and both baseline models, the RMSFE is lower if the model in question includes the ISM index. The Giacomini-White (GW) test statistics are all statistically significant, indicating that differences in out-of-sample forecast performance, as measured by RMSFE, are not attributable to noise. The standardized regression coefficients provide interesting information. First, the relative importance of the ISM index increases as the forecast horizon become longer, indicating that the index is helpful for picking up changes in longer-term inflation behavior. Second, the ISM index adds more predictive power over the full sample than some other well-known predictor variables, such as real disposable income growth, the interest-rate spread, the S&P 500 stock index, and the V/U ratio. For example, the only variables with more predictive content (in terms of the magnitude of the standardized coefficient) in the $h = 24$ and $h = 36$ forecasts are 1-year ahead expected inflation from households and the oil price. Appendix Table A2 show very similar results when the ISM index is constructed using an AR(3) model.

5 Validation of ISM Index with Externally Identified Shocks

To verify the empirical relevance of the ISM index, we examine its dynamic response to two well-identified exogenous disturbances: (1) monetary policy attempted disinflations identified by Romer and Romer (2024) and (2) news shocks to oil supply identified by Känzig (2021). This exercise serves as a diagnostic test for whether the ISM index captures the impact of demand or supply shocks on inflation in a manner consistent with economic theory.

We estimate the dynamic response of the ISM index using the local projection. For each horizon $h = 0, 1, \dots, 36$ months, we estimate the following specification:

$$ISM_{t+h} - ISM_{t-1} = c + \beta^h \text{shock}_t + \sum_{i=1}^{12} \gamma_i^h ISM_{t-i} + \sum_{i=1}^{12} \delta_i^h \text{Recession}_{t-i} + \epsilon_{t+h}, \quad (15)$$

where ISM_{t+h} is the inflation shock momentum index at horizon h , and shock_t represents

Table 2: Out-of-Sample Forecasting Performance, LASSO Model

(a) Baseline Model

	Baseline			Baseline + ISM		
	h=12	h=24	h=36	h=12	h=24	h=36
ISM Index				0.22	0.25	0.27
Headline PCE (y/y)	1.69	1.21	0.82	1.57	1.11	0.95
RMSFE Ratio				0.898	0.882	0.907
GW F-test				206.7	98.4	120.7
GW p-value				0.00	0.00	0.00

(b) Baseline Model with Controls

	Baseline			Baseline + ISM		
	h=12	h=24	h=36	h=12	h=24	h=36
ISM Index				0.24	0.52	0.75
Headline PCE (y/y)	0.47		0.32	0.50		0.43
Headline PCE (3m)	0.57	0.33		0.48	0.15	-0.25
1-yr Inf. Exp.	0.80	1.16	0.97	0.75	1.11	0.84
V/U ratio	-0.26		0.34	-0.34	-0.20	
Interest Rate Spread (10yr - FFR)	-0.13	0.29	0.63	-0.12	0.30	0.70
Oil price (level)	-0.84	-1.08	-0.97	-0.84	-1.08	-1.04
S&P 500 Index (level)	0.53	0.38		0.49	0.29	
Real Income (y/y)	0.23	0.26	0.10	0.24	0.29	0.20
NBER Recession	-0.19	-0.13		-0.23	-0.22	-0.10
RMSFE Ratio				0.989	0.985	0.969
GW F-test				28.6	39.4	59.1
GW p-value				0.00	0.00	0.00

Notes: All results use the ISM index defined on 3 consecutive same-sign residuals from the AR(1) model. Panel (a) reports results where the baseline model includes 12-month PCE inflation as the only predictor. Panel (b) reports results where the baseline model includes additional predictor variables as controls. Three forecast horizons are estimated, $h = 12$ months, $h = 24$ months, and $h = 36$ months. “RMSFE Ratio” represents the out-of-sample rolling 10-year window RMSFE of the “Baseline + ISM” model divided by the “Baseline” model. Standardized regression coefficients estimated on the *full sample* and optimal λ^* are reported for each predictor. Coefficients on each of the 10-year rolling window regressions used to construct the RMSFE are not reported. All regressions include a constant term.

either the Romer-Romer attempted disinflation shock or the Känzig oil supply news shock. We include 12 lags of the dependent variable to account for the ISM index’s inherent persistence. We also include a dummy variable for NBER recessions to ensure that we are assessing the non-recessionary dynamics of the ISM index in response to an external shock. The results presented below are qualitatively robust to including additional controls, but the estimates are somewhat noisier. The coefficient of interest, β^h , maps the impulse response of the ISM index over a three-year window.

5.1 Response of ISM index to attempted disinflation shocks

Romer and Romer (2024) identify “attempted disinflation dates” as episodes when the FOMC explicitly shifted monetary policy to reduce inflation. Based on narrative evidence from FOMC meeting transcripts, the authors categorize disinflation attempts using three specific criteria: (1) an official declaration that inflation is unacceptable, (2) the implementation of concrete contractionary actions, and (3) an expressed willingness to endure output consequences. Each identified episode is assigned a commitment rating of low, medium, or high, determined by the magnitude of accepted output costs and the rigidity of the inflation target.

Based on the above conceptual framework, Romer and Romer (2024) identify five attempted disinflation episodes since 1969. These include two low-commitment attempts in April 1974 and August 1978. These episodes were followed by two medium-to-high commitment attempts in October 1979 and May 1981 under Fed Chair Paul Volcker. The fifth medium-commitment episode took place in December 1988 under Fed Chair Alan Greenspan. To these five episodes, we add a high-commitment episode in September 2022 under Fed Chair Jerome Powell, but results are similar if this episode is excluded.¹⁹ Najjar and Shapiro (2025) find that successful disinflations (identified as medium to high commitment episodes from Romer and Romer (2024)) are associated with a decline in the demand-driven component of inflation, consistent with the theoretical response to a monetary policy tightening.

Panel A of Figure 3 shows the response of the ISM index to a medium- or high-commitment disinflation shock. The ISM index exhibits a significant and persistent decline following the impact date of the shock. The ISM index drops on impact and then continues to decline for the next three years, albeit somewhat unevenly. The statistically significant

¹⁹Following Powell’s August 2022 Jackson Hole statement that disinflation would require “below-trend growth,” the FOMC raised the target range of the federal funds rate by 0.75 percentage point at its September 2022 meeting. Although Romer and Romer (2024) initially excluded this episode due to data limitations, the meeting’s median year-end projection for inflation from the Summary of Economic Projections (SEP) rose by a full percentage point relative to the previous median projection from the June SEP. These actions signaled a substantial tightening of monetary policy in response to elevated inflation.

downward trajectory confirms that the ISM index captures movements in inflation driven by a credible tightening of monetary policy.

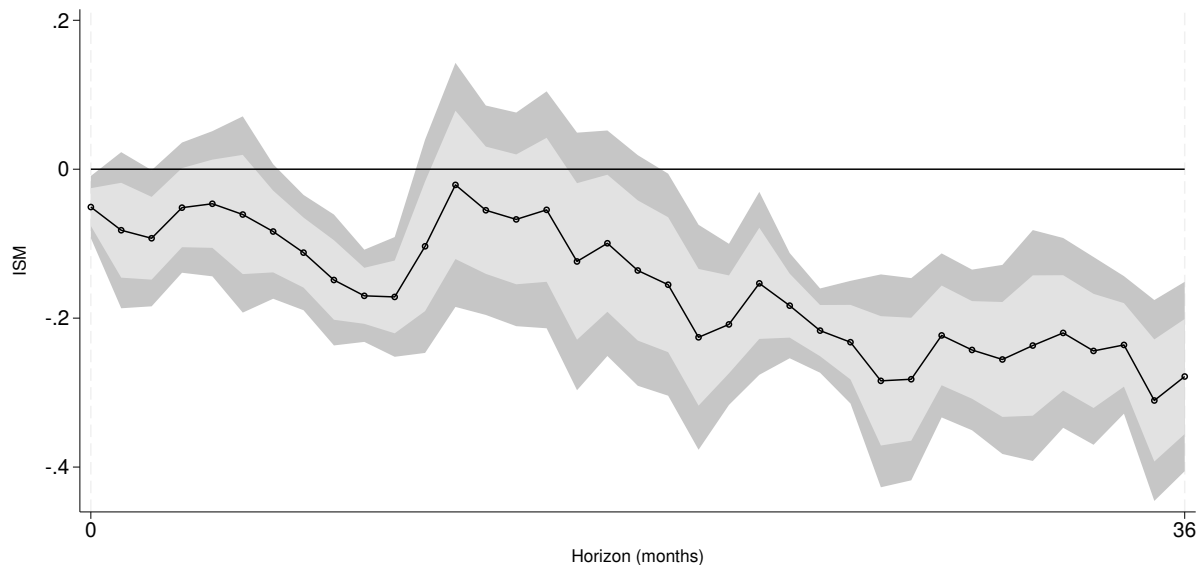
5.2 Response of ISM Index to Oil Supply News Shocks

As a complement to the analysis of demand-side shocks, we use news shocks about future oil supply identified by Känzig (2021) to validate the response of the ISM index. The news shocks capture changes in market expectations about future oil supply following announcements made by the Organization of Petroleum Exporting Countries (OPEC). Each of these announcements represents a high-frequency, exogenous measure of a shock to the prevailing oil supply.

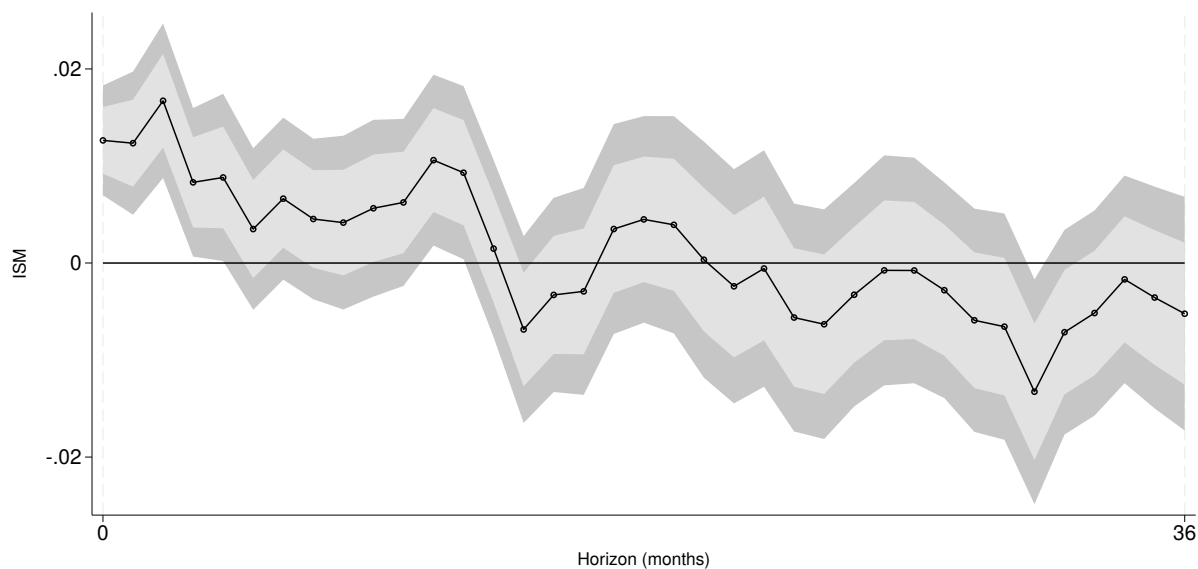
Panel B of Figure 3 shows the response of the ISM index to a one standard deviation move in the shock, representing negative news about the future oil supply. The ISM index increases on impact and remains elevated for approximately one year. The ISM response to the oil supply news shock is much less persistent than the corresponding response to the disinflation shock shown in panel (a). These divergent responses—a persistent ISM index decline following a credible monetary policy tightening and a more transitory ISM index rise following negative news about the future oil supply—provide strong evidence that the ISM captures emerging shifts in the underlying data generating process for inflation.

Figure 3: Impulse Response of ISM index to Externally Identified Shocks

(a) Response of ISM Index to Romer and Romer (2024) Disinflation Shock



(b) Response of ISM Index to Känzig (2021) Oil Supply News Shock



Notes: Panel (a) shows the cumulative impact of the ISM index to a medium-to-high commitment Romer and Romer (2024) attempted disinflation shock over 36 months, based on estimating equation (15) at each forecast horizon h . Panel (b) shows the cumulative impact of the ISM to a 1 sd shock of the Känzig (2021) oil supply news shock. 1 sd and 90 percent confidence bands shown.

6 Conclusion

This paper develops a simple, non-parametric measure of inflation shock momentum that exploits the cross-sectional structure of disaggregated PCE price data. By identifying sustained directional runs in category-level inflation residuals from a rolling autoregressive benchmark data generating process, we construct an Inflation Shock Momentum (ISM) index that tracks the net share of expenditure-weighted categories experiencing positive versus negative shock momentum. The ISM is sign-sensitive and places weight on recent shock sequences, allowing it to distinguish between emerging inflationary and disinflationary pressures in real time. We show that the ISM index contains statistically and economically significant predictive value for aggregate 12-month PCE inflation at horizons of one to three years. The index improves both in-sample fit and out-of-sample forecast accuracy relative to baseline models that include standard macroeconomic and financial predictors. The negative component of the ISM index is particularly informative, consistent with the historical importance of sustained disinflation episodes in shaping medium-term inflation dynamics.

Our findings contribute to a growing literature that leverages disaggregated price data to better understand and forecast aggregate inflation. While much of the traditional forecasting literature focuses on aggregate time series dynamics, our results demonstrate that the cross-sectional distribution of sectoral shocks contains early information about shifts in the underlying inflation environment. By using over 100 simultaneous category-level observations each month, the ISM index identifies coordinated directional pressure that would take substantially longer to detect using aggregate data alone. Moreover, the behavior of the index in response to externally identified shocks aligns with economic theory: The ISM index declines persistently following credible monetary policy disinflations and rises temporarily in response to adverse oil supply news. These validation exercises support the interpretation of the ISM index as capturing meaningful changes in the data generating process for inflation, rather than merely reflecting short-run noise.

From a policy perspective, the ISM index provides a transparent and timely indicator of broad-based inflation momentum. Distinguishing between transitory versus more persistent inflation shocks hitting the economy remains a central challenge for monetary policymakers. Delayed recognition of sustained inflationary pressure risks the de-anchoring of expectations, thus requiring a more forceful subsequent policy response, while overreaction to truly transitory shocks can impose unnecessary costs on output and employment. By summarizing the breadth and direction of recent inflation shocks across categories, the ISM index offers policymakers and practitioners an additional tool for assessing directional inflation pressures in real time.

References

- Barnichon, R. (2010). Building a composite help-wanted index. *Economics Letters* 109(3), 175–178.
- Bernanke, B. (2007). Inflation expectations and inflation forecasting. Technical report, Board of Governors of the Federal Reserve System (US).
- Bonaparte, Y., F. J. Fabozzi, and M. Peron (2025). Measuring transitory inflation: Implications for monetary policy and stock market volatility. *Journal of International Money and Finance* 153, 103284.
- Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2023). Anchored inflation expectations. *American Economic Journal: Macroeconomics* 15(1), 1–47.
- Cavallo, A. and G. G. Zavaleta (2025). Turning points in inflation: A structural breaks approach with micro data. Technical report, National Bureau of Economic Research.
- Chiavari, A., M. Morazzoni, and D. Smirnov (2025). Heterogeneous markups cyclical and monetary policy.
- Eickmeier, S. and B. Hofmann (2022). What drives inflation? disentangling demand and supply factors.
- Faust, J. and J. H. Wright (2013). Forecasting inflation. In *Handbook of economic forecasting*, Volume 2, pp. 2–56. Elsevier.
- Giacomini, R. and Y. Levin (2023). Microforecasting inflation. *Economic Perspectives* (3).
- Giacomini, R. and B. Rossi (2009). Detecting and predicting forecast breakdowns. *The Review of Economic Studies* 76(2), 669–705.
- Giacomini, R. and H. White (2006). Tests of conditional predictive ability. *Econometrica* 74(6), 1545–1578.
- Goodfriend, M. and R. G. King (2005). The incredible volcker disinflation. *Journal of Monetary Economics* 52(5), 981–1015.
- Hobijn, B. and F. Nechio (2025). The effects of tariffs on inflation and production costs. *FRBSF Economic Letter* 2025(12), 1–6.
- Huang, D., J. Li, L. Wang, and G. Zhou (2020). Time series momentum: is it there? *Journal of financial economics* 135(3), 774–794.

- Jorda, O. (2005, March). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review* 95(1), 161–182.
- Jørgensen, P. L. and K. J. Lansing (2025). Anchored inflation expectations and the slope of the phillips curve. *European Economic Review*, 105073.
- Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from opec announcements. *American Economic Review* 111(4), 1092–1125.
- Lansing, K. J. (2022). Untangling persistent versus transitory shocks to inflation. *FRBSF Economic Letter* 13.
- Mahedy, T. and A. H. Shapiro (2017). What’s down with inflation? *FRBSF Economic Letter* 35.
- Najjar, R. and A. H. Shapiro (2025). Does monetary policy tightening reduce inflation? *FRBSF Economic Letter* 3, 1–6.
- Quineche, R. and J. Zapata (2026). The two sides of inflation: How price increases and price cuts shape inflation. *Economic Modelling* 158, 107533.
- Romer, C. D. and D. H. Romer (2024). Lessons from history for successful disinflation. *Journal of Monetary Economics* 148, 103654.
- Shapiro, A. H. (2024). Decomposing supply-and demand-driven inflation. *Journal of Money, Credit and Banking*.
- Sheppard, W. F. (1899). On the application of the theory of error to cases of normal distribution and normal correlation. *Philosophical Transactions of the Royal Society of London, Series A: Containing Papers of a Mathematical or Physical Character* (192), 101–167.
- Smith Jr, A. A. (1993). Estimating nonlinear time-series models using simulated vector autoregressions. *Journal of Applied Econometrics* 8(S1), S63–S84.
- Stock, J. H. and M. W. Watson (2007). Why has us inflation become harder to forecast? *Journal of Money, Credit and banking* 39, 3–33.
- Stock, J. H. and M. W. Watson (2016). Core inflation and trend inflation. *Review of Economics and Statistics* 98(4), 770–784.

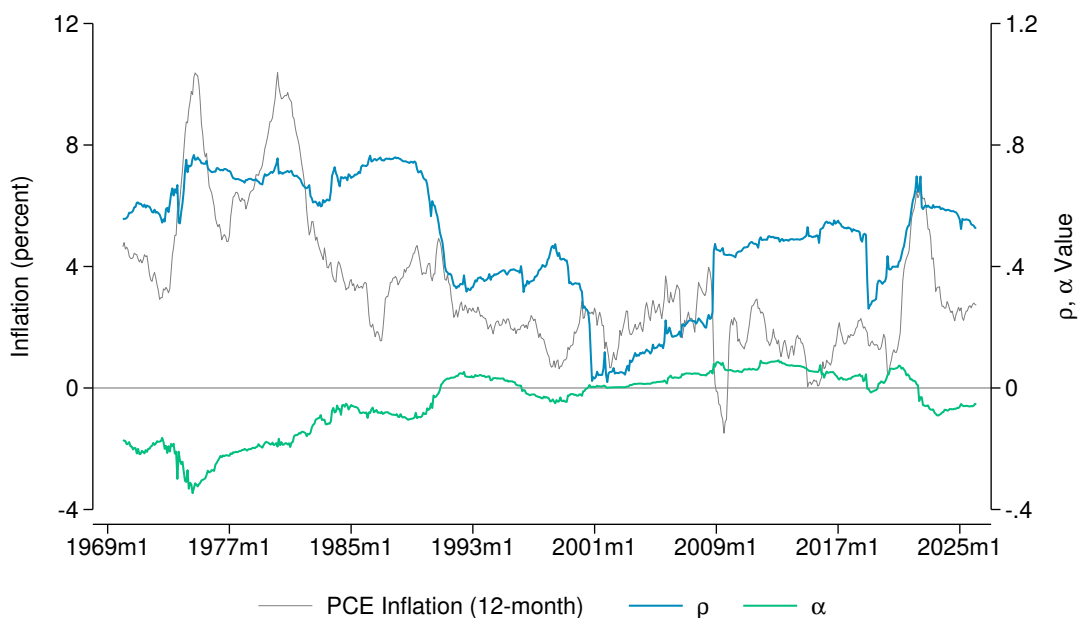
- Stock, J. H. and M. W. Watson (2020). Slack and cyclically sensitive inflation. *Journal of Money, Credit and Banking* 52(S2), 393–428.
- Timmermann, A. (2000). Moments of markov switching models. *Journal of econometrics* 96(1), 75–111.
- Van Vleck, J. H. and D. Middleton (1966). The spectrum of clipped noise. *Proceedings of the IEEE* 54(1), 2–19.
- Verona, F. (2026). Forecasting inflation: The sum of the cycles outperforms the whole. Technical report, Bank of Finland Research Discussion Papers.
- Zaman, S. (2019). Cyclical versus acyclical inflation: A deeper dive. *Economic Commentary* (2019-13).

Online Appendix

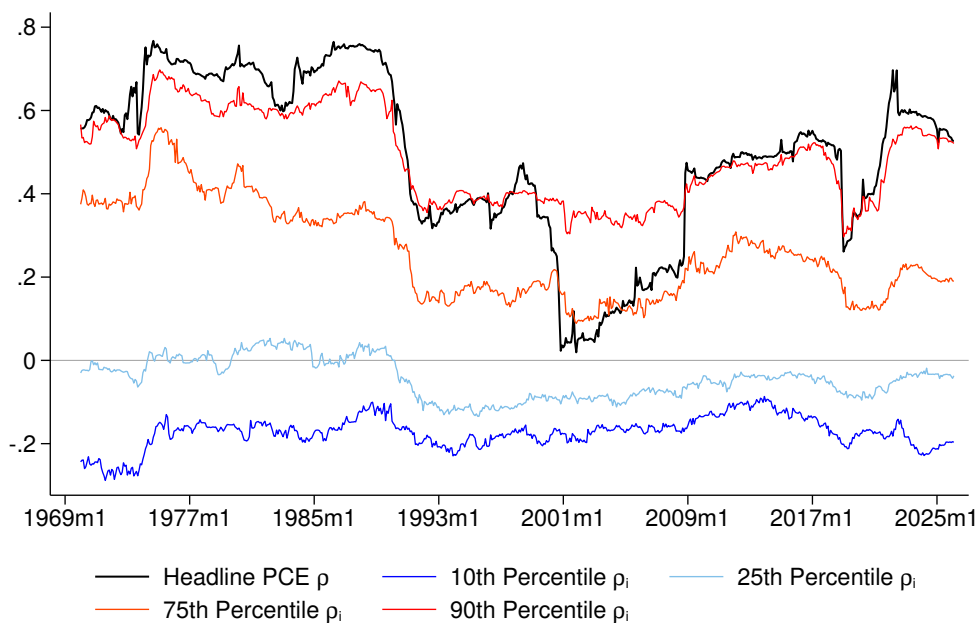
A Tables and Figures

Figure A1: Rolling-Window Estimates of ρ and α

(a) Rolling Window Estimates of $\hat{\rho}$ and $\hat{\alpha}$ for Headline PCE Inflation

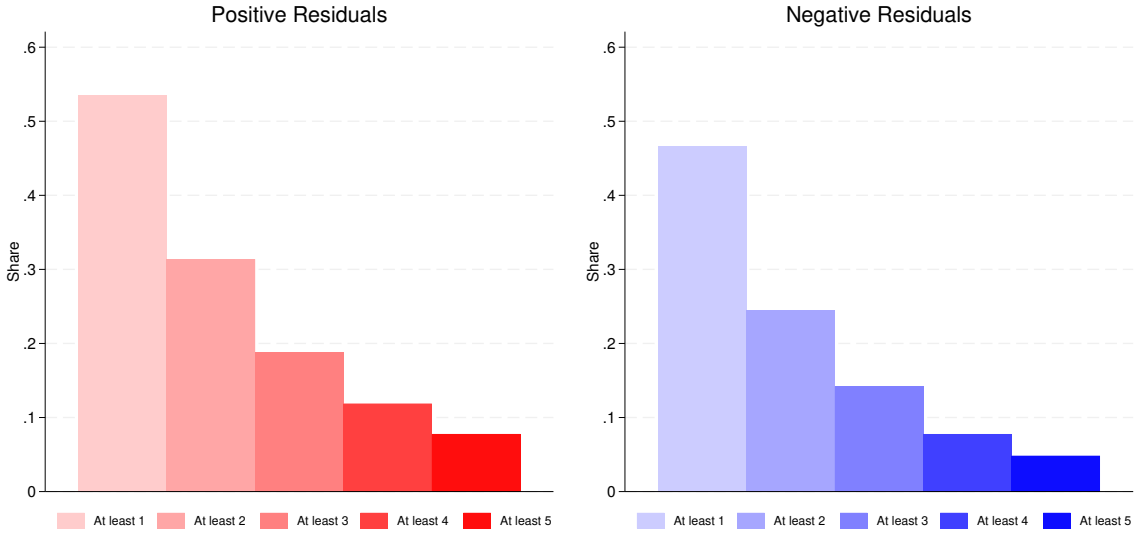


(b) Distribution of Estimated $\hat{\rho}_i$ Values Across Categories




Note: Panel (a) shows the 10-year rolling window estimates of $\hat{\rho}$ and $\hat{\alpha}$ from estimating equations on (1) and (2) using monthly headline (aggregate) PCE inflation. Panel (b) shows the distribution of 10-year rolling window estimates of $\hat{\rho}_i$ from estimating equation (3) using monthly category-level inflation, arranged by percentile at the end date of each rolling window. The black line in panel (b) shows 10-year rolling window estimates of the persistence parameter ρ from estimating a simple AR(1) model on monthly headline PCE inflation.

Figure A2: Distribution of Consecutive Residuals Across Observations



Note: Each bar represents the share of total monthly observations across categories with at least k consecutive positive or negative residuals.

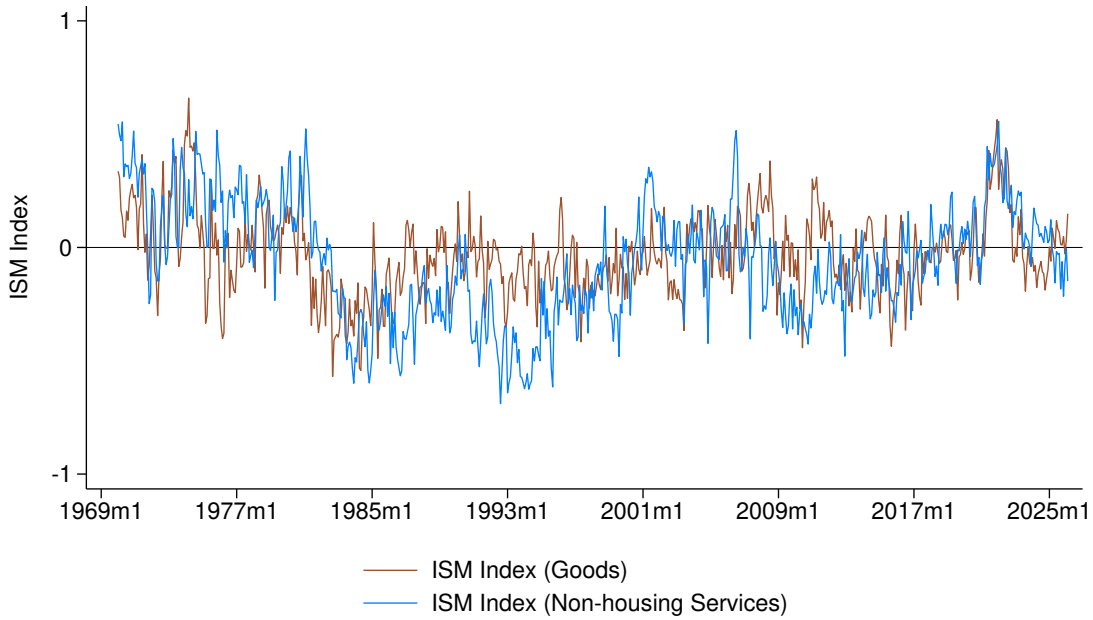
Figure A3: Probabilities of Consecutive Same-Signed Residuals



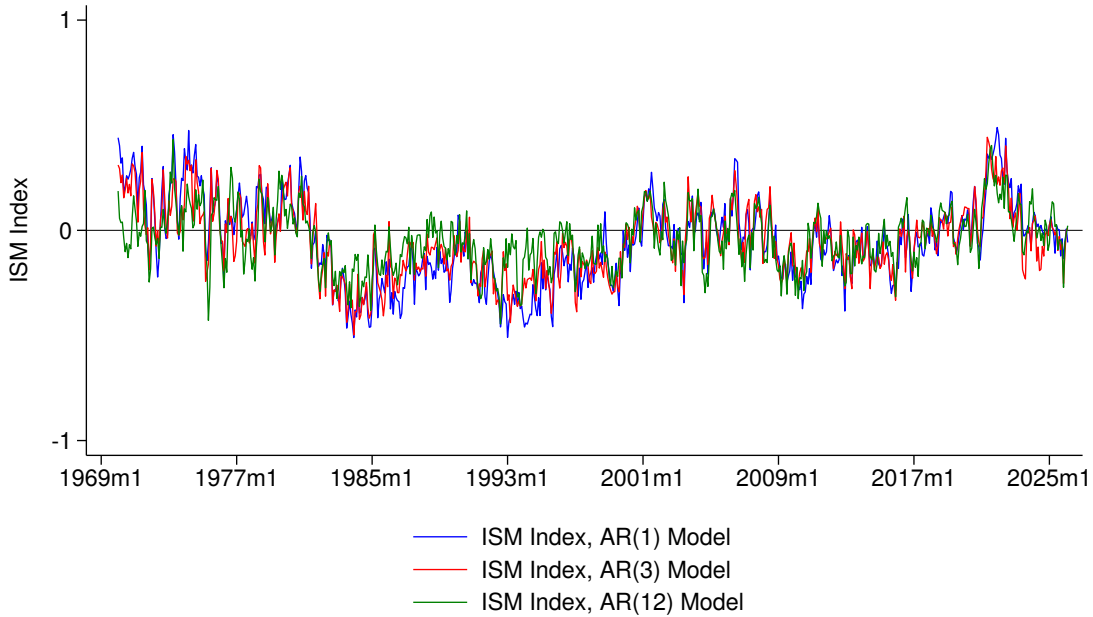
Notes: Given equation (2), the figure plots the probability of observing at least k consecutive positive residuals as a function of the persistence parameter α . The probability of consecutive negative residuals is identical. The probability of observing at least k consecutive residuals of the *same sign* (i.e., consecutive positive *or* consecutive negative) would be the same chart with the y-axis scaled by a factor of two. Details are shown in Appendix B.

Figure A4: ISM Index by Broad Category and by AR(\cdot) Model

(a) ISM Index by Goods and Services (non-housing)

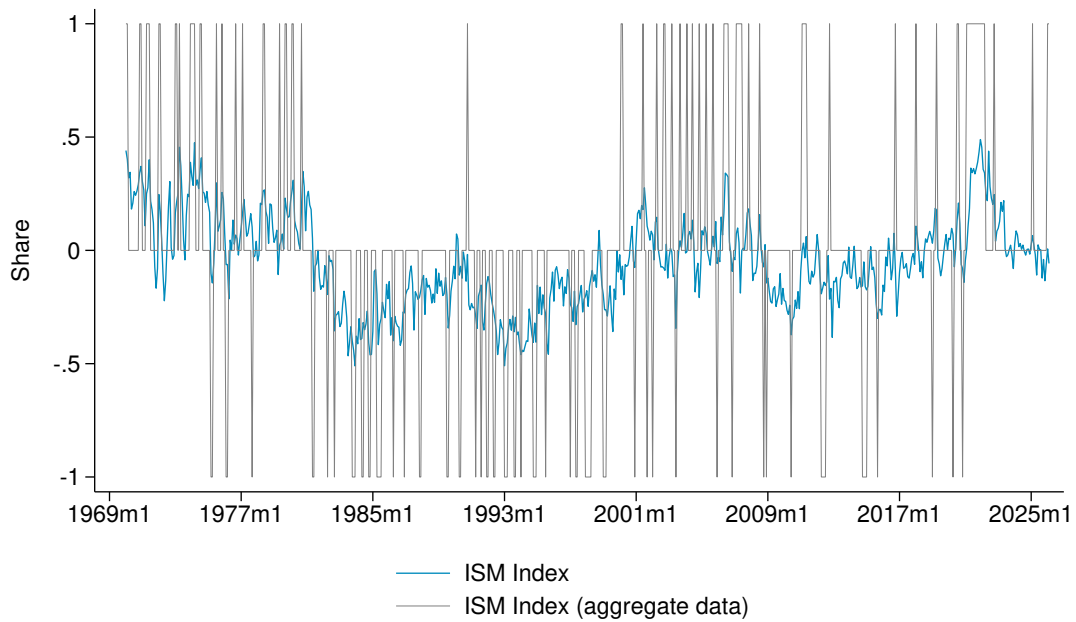


(b) ISM Index by AR(\cdot) Model



Note: Panel (a) shows the ISM index by broad category: goods and non-housing services. Housing services are not included in panel (a). Panel (b) shows the ISM index under alternative assumptions regarding the benchmark data generating process—an AR(3) or an AR(12). The AR(1) and AR(3) model have a correlation of 0.91, while the AR(1) and AR(12) model have a correlation of 0.75.


Figure A5: ISM Index Constructed with Aggregate vs Category-level Data




Note: ISM Index (aggregate) is constructed by running 120-month rolling-window regressions on aggregate monthly inflation: $\pi_t = \mu + \rho\pi_{t-1} + \varepsilon_t$. The index is equal to 1 in period t if 3 consecutive residuals are greater than zero, equal to -1 if 3 consecutive residuals are less than zero, and equal to 0 otherwise. The correlation of the two series is 0.51.

Figure A6: Impulse Response to Momentum Surprises: By number of consecutive same-sign residuals, AR(1) momentum model


(a) Two consecutive same-sign residuals



(b) Three consecutive same-sign residuals

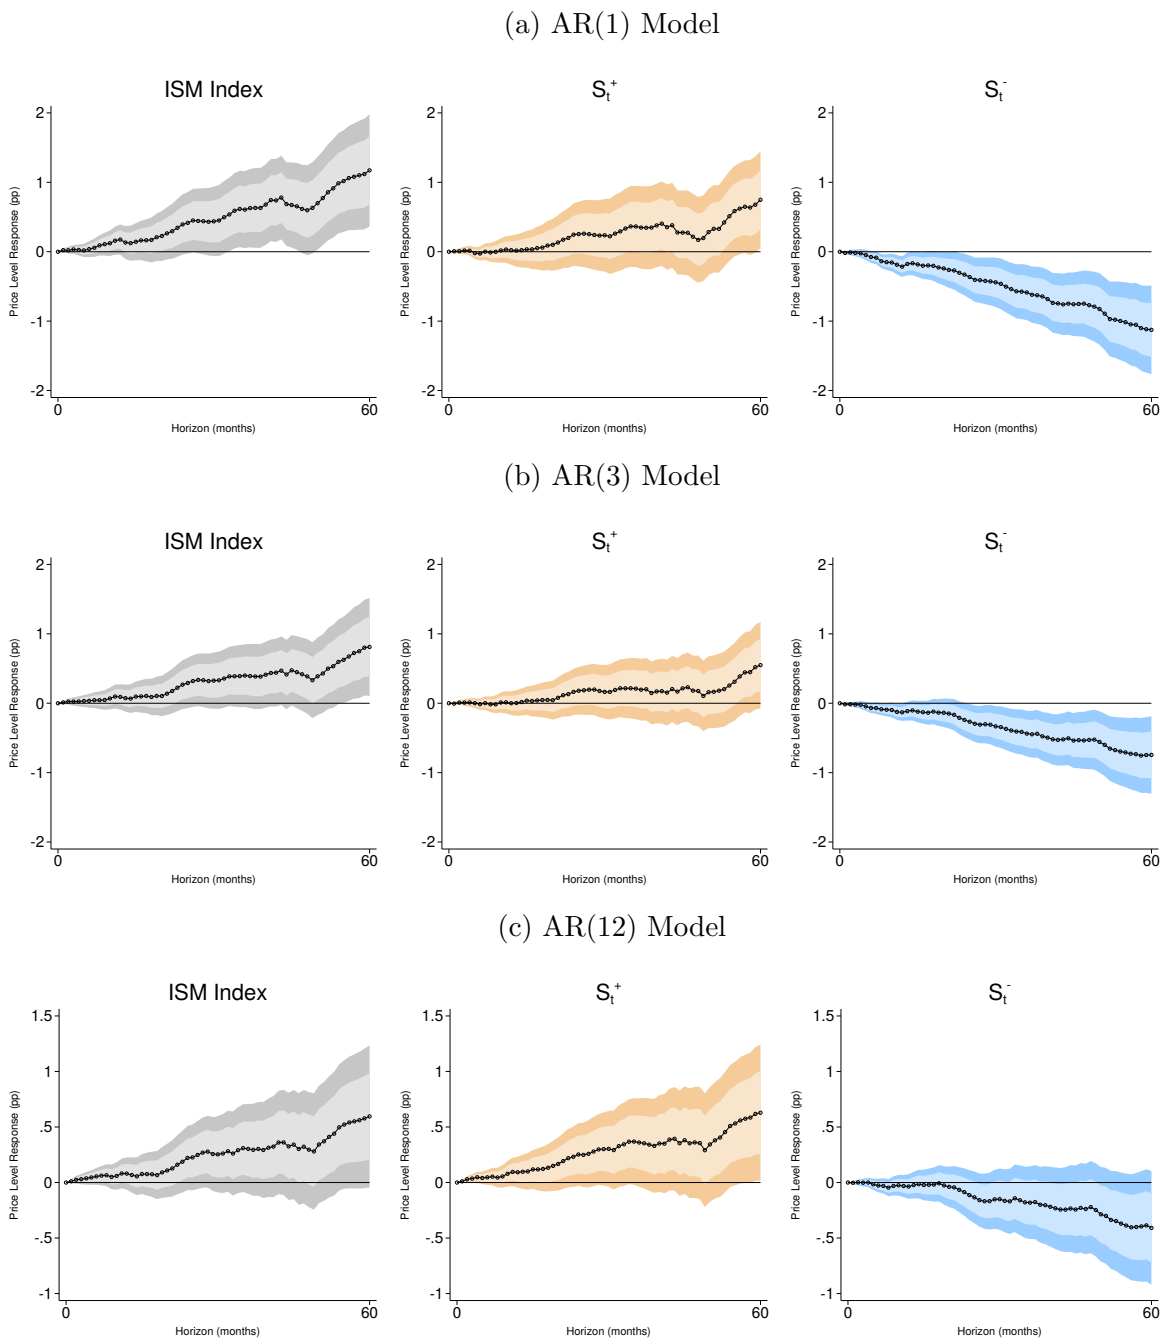


(c) Four consecutive same-sign residuals



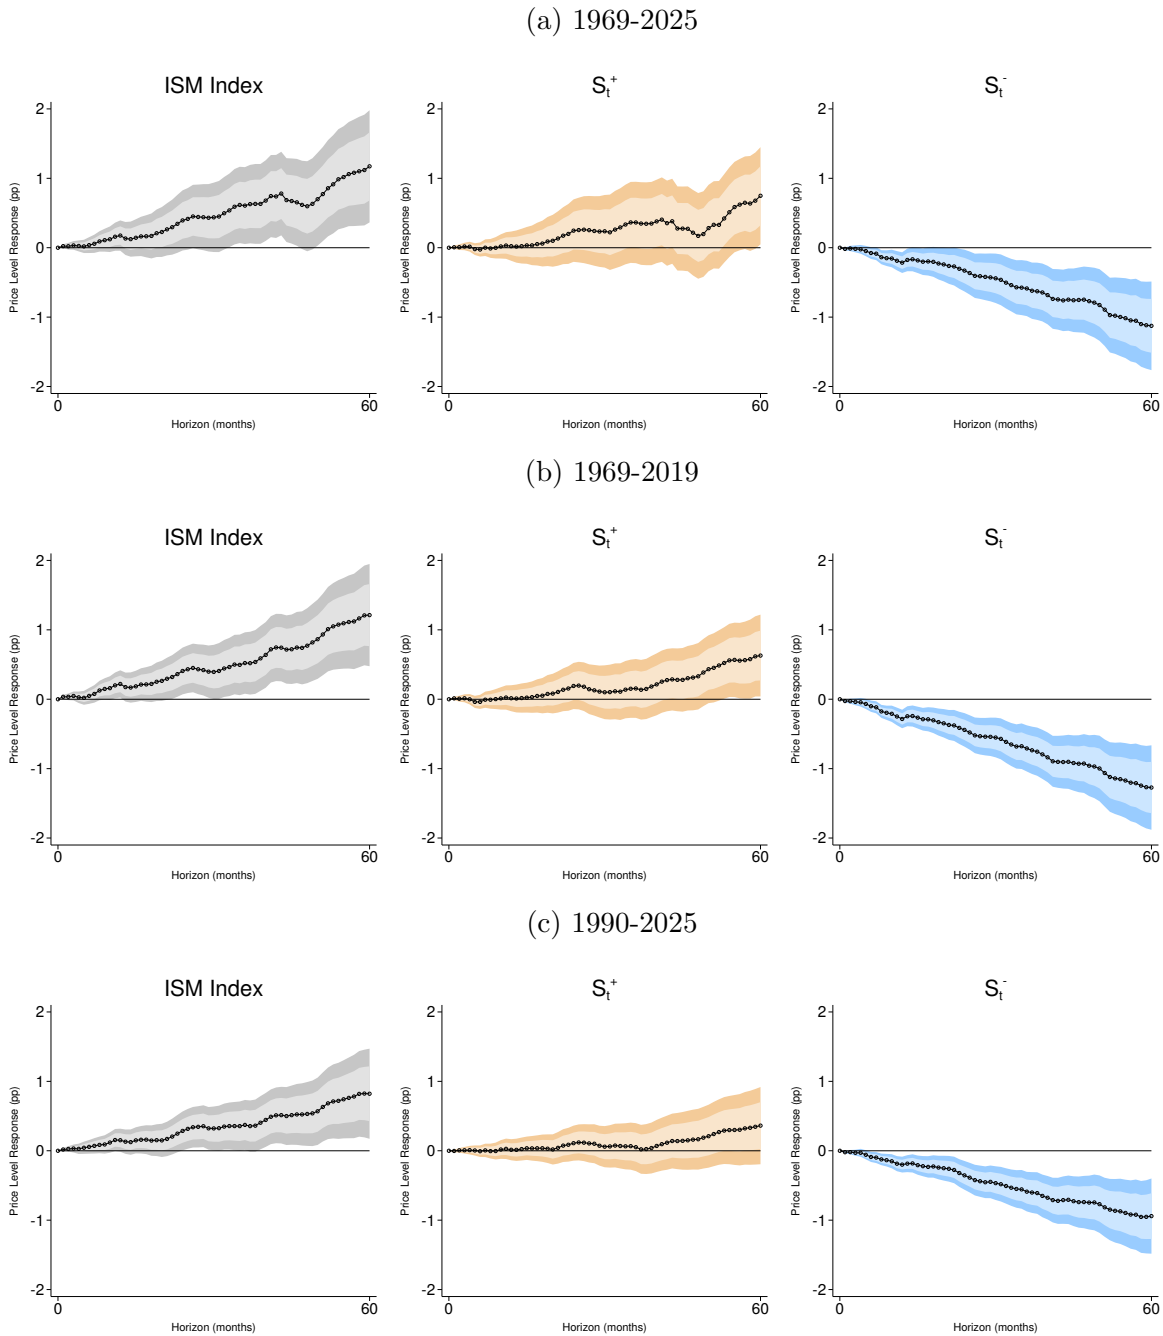
Notes: The left panels show the cumulative impact of a one standard deviation (sd) surprise of the ISM index on the log PCE price level over 60 months based on estimating equation (9) at each forecast horizon h . The middle panels show the cumulative impact of a 1 sd shock of the positive ISM component, S_t^+ . The right panels show the cumulative impact of a 1 sd shock of the negative ISM component, S_t^- . 1 sd and 90 percent confidence bands shown. Panels (a), (b), and (c) show results when the ISM index is defined based on two, three, and four consecutive same-signed residuals, respectively.

Figure A7: Impulse Response to Momentum Surprises: By AR(\cdot) model, three consecutive same-sign residuals



Notes: The left panels show the cumulative impact of a one standard deviation (sd) surprise of the ISM index on the log PCE price level over 60 months based on estimating equation (9) at each forecast horizon h . The middle panels show the cumulative impact of a 1 sd shock of the positive ISM component, S_t^+ . The right panels shows the cumulative impact of a 1 sd shock of the negative ISM component, S_t^- . 1 sd and 90 percent confidence bands shown. Panels (a), (b), and (c) show results when the ISM is defined based on a data-generating process of inflation (that is, the functional form of equation (1)) of an AR(1), AR(3), or an AR(12) model, respectively.

Figure A8: Impulse Response to Momentum Surprises: By subsample, AR(1) model, three consecutive same-sign residuals



Notes: The left panels show the cumulative impact of a one standard deviation (sd) surprise of the ISM index on the log PCE price level over 60 months based on estimating equation (9) at each forecast horizon h . The middle panels show the cumulative impact of a 1 sd shock of the positive ISM component, S_t^+ . The right panels shows the cumulative impact of a 1 sd shock of the negative ISM component, S_t^- . 1 sd and 90 percent confidence bands shown. Panels (a), (b), and (c) show results under alternative subsamples: the full sample (1969-2025), pre-Covid (1969-2019), and post-Volcker disinflation (1990-2025).

Table A1: In-sample Forecast Fit: Robustness across models, with controls

(a) Definition of Momentum: Two consecutive same-sign residuals

	12-month ahead forecast			24-month ahead forecast			36-month ahead forecast		
	(1) AR-1	(2) AR-3	(3) AR-12	(4) AR-1	(5) AR-3	(6) AR-12	(7) AR-1	(8) AR-3	(9) AR-12
PCE Inflation (y/y)	0.234*** (0.0652)	0.242*** (0.0666)	0.235*** (0.0661)	0.0404 (0.0712)	0.0533 (0.0721)	0.0449 (0.0702)	0.225*** (0.0644)	0.241*** (0.0648)	0.234*** (0.0667)
ISM Index	1.080*** (0.244)	0.959*** (0.256)	0.504 (0.313)	1.898*** (0.307)	1.722*** (0.342)	0.860** (0.379)	2.810*** (0.354)	2.390*** (0.370)	1.043** (0.414)
N	674	674	674	662	662	662	650	650	650
R ²	0.731	0.728	0.723	0.549	0.541	0.526	0.514	0.493	0.462
Adjusted R ²	0.726	0.723	0.719	0.541	0.534	0.518	0.506	0.484	0.453
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

(b) Definition of Momentum: Three consecutive same-sign residuals

	12-month ahead forecast			24-month ahead forecast			36-month ahead forecast		
	(1) AR-1	(2) AR-3	(3) AR-12	(4) AR-1	(5) AR-3	(6) AR-12	(7) AR-1	(8) AR-3	(9) AR-12
PCE Inflation (y/y)	0.231*** (0.0650)	0.240*** (0.0663)	0.239*** (0.0659)	0.0391 (0.0707)	0.0578 (0.0715)	0.0531 (0.0702)	0.224*** (0.0626)	0.244*** (0.0632)	0.232*** (0.0662)
ISM Index	1.198*** (0.309)	1.154*** (0.333)	0.927** (0.421)	2.523*** (0.390)	2.471*** (0.448)	1.729*** (0.516)	3.696*** (0.457)	3.409*** (0.491)	1.485** (0.589)
N	674	674	674	662	662	662	650	650	650
R ²	0.728	0.727	0.724	0.551	0.545	0.530	0.517	0.499	0.463
Adjusted R ²	0.724	0.722	0.720	0.543	0.537	0.522	0.509	0.490	0.453
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

(c) Definition of Momentum: Four consecutive same-sign residuals

	12-month ahead forecast			24-month ahead forecast			36-month ahead forecast		
	(1) AR-1	(2) AR-3	(3) AR-12	(4) AR-1	(5) AR-3	(6) AR-12	(7) AR-1	(8) AR-3	(9) AR-12
PCE Inflation (y/y)	0.221*** (0.0644)	0.232*** (0.0656)	0.235*** (0.0653)	0.0195 (0.0699)	0.0444 (0.0704)	0.0430 (0.0698)	0.194*** (0.0619)	0.228*** (0.0626)	0.218*** (0.0654)
ISM Index	1.444*** (0.370)	1.320*** (0.412)	1.469*** (0.516)	3.078*** (0.464)	3.068*** (0.576)	2.504*** (0.648)	4.583*** (0.563)	4.429*** (0.666)	1.642** (0.765)
N	674	674	674	662	662	662	650	650	650
R ²	0.728	0.726	0.725	0.550	0.542	0.531	0.517	0.496	0.461
Adjusted R ²	0.723	0.721	0.721	0.542	0.534	0.523	0.509	0.488	0.452
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The dependent variable is 12-month-ahead 12-month inflation in columns 1 through 3, 24-month ahead 12-month inflation in columns 4 through 6, and 36-month-ahead 12 month inflation in columns 7 through 9. Columns 1, 4, and 7 depicts results where residuals are estimated using the AR(1) model, columns 2, 5, and 8 using the AR(3) model, and columns 3, 6, and 9 using the AR(12) model. Panel (a) depicts results where the ISM index is defined based on two consecutive same-sign residuals, panel (b) depicts results based on three consecutive same-sign residuals (our baseline), and panel (c) depicts results based on four consecutive residuals. Controls include 3-month PCE inflation, 1-year ahead household inflation expectations, the V/U ratio, the price of oil, the S&P 500 stock price index, real personal disposable income growth (year-over-year), the interest rate spread (10 year minus the federal funds rate), and a recession dummy variable indicator. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A2: Out-of-Sample Forecasting Performance: LASSO Model using AR(3) momentum index

(a) Baseline Model

	Baseline			Baseline + ISM		
	h=12	h=24	h=36	h=12	h=24	h=36
ISM Index				0.20	0.22	
Headline PCE (y/y)	1.69	1.21	0.82	1.59	1.14	1.06
RMSFE Ratio				0.897	0.903	0.941
GW F-test				131.1	122.1	178.7
GW p-value				0.00	0.00	0.00

(b) Baseline Model with Controls

	Baseline			Baseline + ISM		
	h=12	h=24	h=36	h=12	h=24	h=36
ISM Index				0.20	0.44	0.59
Headline PCE (y/y)	0.47		0.32	0.52	0.04	0.30
Headline PCE (3m)	0.57	0.33		0.47	0.13	-0.03
1-yr Inf. Exp.	0.80	1.16	0.97	0.77	1.15	0.83
V/U ratio	-0.26		0.34	-0.33	-0.18	
Interest Rate Spread (10yr - FFR)	-0.13	0.29	0.63	-0.13	0.28	0.65
Oil price (level)	-0.84	-1.08	-0.97	-0.85	-1.11	-1.03
S&P 500 Index (level)	0.53	0.38		0.51	0.34	
Real Income (y/y)	0.23	0.26	0.10	0.23	0.28	0.20
NBER Recession	-0.19	-0.13		-0.22	-0.20	-0.01
RMSFE Ratio				0.987	0.988	0.993
GW F-test				5.4	26.6	34.8
GW p-value				0.00	0.00	0.00

Notes: All results use the ISM index defined on 3 consecutive same-sign residuals from the AR(3) model. Panel (a) reports results where the baseline model includes 12-month PCE inflation as the only predictor. Panel (b) reports results where the baseline model includes additional predictor variables as controls. Three forecast horizons are estimated, $h = 12$ months, $h = 24$ months, and $h = 36$ months. “RMSFE Ratio” represents the out-of-sample rolling 10-year window RMSFE of the “Baseline + Momentum” model divided by the “Baseline” model. Standardized regression coefficients estimated on the *full sample* and optimal λ^* are reported for each predictor. Coefficients on each of the 10-year rolling window regressions that are used to construct the RMSFE values are not reported. All regressions include a constant term.

Table A3: In-sample Forecast Regressions: Weighting by shock size ($\hat{\varepsilon}$)

	12-month ahead forecast						24-month ahead forecast					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
PCE Inflation (y/y)	0.767*** (0.0287)	0.721*** (0.0326)	0.722*** (0.0322)	0.216*** (0.0626)	0.230*** (0.0651)	0.229*** (0.0659)	0.565*** (0.0337)	0.519*** (0.0401)	0.518*** (0.0396)	0.0155 (0.0662)	0.0316 (0.0743)	0.0306 (0.0720)
S_t^+		0.508** (0.254)			0.667*** (0.217)			0.197 (0.289)			0.388 (0.274)	
S_t^-		-0.148 (0.209)			0.0654 (0.207)			-0.727** (0.290)			-1.335*** (0.287)	
ISM Index			0.373*** (0.131)			0.408*** (0.140)			0.395*** (0.148)			0.727*** (0.171)
N	674	674	674	674	674	674	662	662	662	662	662	662
R ²	0.588	0.595	0.594	0.722	0.728	0.726	0.320	0.328	0.327	0.522	0.540	0.536
Adjusted R ²	0.588	0.593	0.593	0.718	0.724	0.722	0.319	0.325	0.325	0.515	0.531	0.529
Controls	No	No	No	Yes	Yes	Yes	No	No	No	Yes	Yes	Yes

Notes: Here we incorporate the *size* of the shock by assuming that $M_{i,t}^+ = \prod_{k=0}^2 \mathbb{1}(\hat{\varepsilon}_{i,t-k} > 0) * |\sum_{k=0}^2 \hat{\varepsilon}_{i,t-k}|$ and $M_{i,t}^- = \prod_{k=0}^2 \mathbb{1}(\hat{\varepsilon}_{i,t-k} < 0) * |\sum_{k=0}^2 \hat{\varepsilon}_{i,t-k}|$, where $|\sum_{k=0}^2 \hat{\varepsilon}_{i,t-k}|$ is the absolute value of the sum of the residuals. The dependent variable is the 12-month-ahead value of 12-month PCE inflation in columns 1 through 6 and the 24-month ahead value of 12-month inflation in columns 7 through 12. Controls are 3-month PCE inflation, 1-year ahead household inflation expectations, the V/U ratio, the price of oil, the S&P 500 stock price index, real personal disposable income growth (year-over-year), the yield curve spread between 10-year Treasury bonds and the federal funds rate, and a dummy variable indicator of recessions. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: In-sample Forecast Regressions: Weighting by $(\frac{1}{1-\hat{\rho}_{i,\tau}})$

	12-month ahead forecast						24-month ahead forecast					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
PCE Inflation (y/y)	0.767*** (0.0287)	0.735*** (0.0309)	0.732*** (0.0291)	0.216*** (0.0626)	0.233*** (0.0652)	0.235*** (0.0652)	0.565*** (0.0337)	0.541*** (0.0369)	0.533*** (0.0353)	0.0155 (0.0662)	0.0438 (0.0716)	0.0373 (0.0707)
S_t^+		0.440 (0.442)			0.842** (0.367)			0.150 (0.477)				0.734* (0.383)
S_t^-		-0.679** (0.267)			-0.619** (0.259)			-0.884*** (0.340)				-1.529*** (0.345)
ISM Index			0.563*** (0.192)			0.729*** (0.195)			0.527** (0.213)			1.141*** (0.223)
N	674	674	674	674	674	674	662	662	662	662	662	662
R ²	0.588	0.594	0.594	0.722	0.728	0.728	0.320	0.326	0.325	0.522	0.538	0.537
Adjusted R ²	0.588	0.592	0.593	0.718	0.723	0.724	0.319	0.323	0.323	0.515	0.530	0.529
Controls	No	No	No	Yes	Yes	Yes	No	No	No	Yes	Yes	Yes

Notes: Here we incorporate the observed inflation persistence of the category by assuming that $M_{i,t}^+ = \prod_{k=0}^2 \mathbb{1}(\hat{\varepsilon}_{i,t-k} > 0) * \frac{1}{1-\hat{\rho}_{i,\tau}}$ and $M_{i,t}^- = \prod_{k=0}^2 \mathbb{1}(\hat{\varepsilon}_{i,t-k} < 0) * \frac{1}{1-\hat{\rho}_{i,\tau}}$, where $\hat{\rho}_{i,\tau}$ is the estimated inflation persistence parameter observed over the rolling-sample period $\tau \in (t-120, t)$. The dependent variable is the 12-month-ahead value of 12-month PCE inflation in columns 1 through 6 and the 24-month ahead value of 12-month inflation in columns 7 through 12. Controls are 3-month PCE inflation, 1-year ahead household inflation expectations, the V/U ratio, the price of oil, the S&P 500 stock price index, real personal disposable income growth (year-over-year), the yield curve spread between 10-year Treasury bonds and the federal funds rate, and a dummy variable indicator of recessions. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A5: In-sample Forecast Regressions: Non-parametric measure using aggregate PCE

	12-month ahead forecast						24-month ahead forecast					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
PCE Inflation (y/y)	0.767*** (0.0287)	0.752*** (0.0290)	0.754*** (0.0287)	0.216*** (0.0626)	0.238*** (0.0653)	0.238*** (0.0651)	0.565*** (0.0337)	0.561*** (0.0350)	0.560*** (0.0347)	0.0155 (0.0662)	0.0349 (0.0679)	0.0322 (0.0682)
S_t^+		0.475** (0.212)			0.305 (0.186)			0.0297 (0.261)			-0.00559 (0.224)	
S_t^-		-0.172 (0.133)			-0.166 (0.147)			-0.220 (0.156)			-0.420** (0.168)	
ISM Index			0.329*** (0.120)			0.239* (0.123)			0.121 (0.144)			0.200 (0.142)
N	674	674	674	674	674	674	662	662	662	662	662	662
R ²	0.588	0.594	0.593	0.722	0.724	0.724	0.320	0.321	0.321	0.522	0.525	0.523
Adjusted R ²	0.588	0.592	0.592	0.718	0.719	0.719	0.319	0.318	0.319	0.515	0.516	0.515
Controls	No	No	No	Yes	Yes	Yes	No	No	No	Yes	Yes	Yes

Notes: This ISM index is constructed by running 120-month rolling-window regressions on aggregate monthly inflation: $\pi_t = \mu_i + \rho\pi_{t-1} + \varepsilon_t$. The index is equal to 1 if 3 consecutive residuals are greater than zero in period t , equal to -1 if 3 consecutive residuals are less than zero, and equal to 0 otherwise. See Figure A5. The dependent variable is the 12-month-ahead value of 12-month PCE inflation in columns 1 through 6 and the 24-month ahead value of 12-month inflation in columns 7 through 12. Controls are 3-month PCE inflation, 1-year ahead household inflation expectations, the V/U ratio, the price of oil, the S&P 500 stock price index, real personal disposable income growth (year-over-year), the yield curve spread between 10-year Treasury bonds and the federal funds rate, and a dummy variable indicator of recessions. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A6: In-sample Forecast Regressions: Parametric measure ($\hat{\alpha}$) using aggregate PCE

	12-month ahead forecast				24-month ahead forecast			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
PCE Inflation (y/y)	0.767*** (0.0287)	0.480*** (0.0449)	0.216*** (0.0626)	0.160** (0.0679)	0.565*** (0.0337)	-0.0342 (0.0444)	0.0155 (0.0662)	-0.303*** (0.0628)
Aggregate $\hat{\alpha}$		-7.941*** (1.024)		-2.292** (1.053)		-16.52*** (1.056)		-14.58*** (1.280)
N	674	674	674	674	662	662	662	662
R ²	0.588	0.640	0.722	0.724	0.320	0.544	0.522	0.608
Adjusted R ²	0.588	0.639	0.718	0.720	0.319	0.542	0.515	0.601
Controls	No	No	Yes	Yes	No	No	Yes	Yes

Notes: To construct $\hat{\alpha}$ we run 10-year rolling-window regressions of the residuals from an AR(1) on monthly PCE inflation. See Panel (a) of Figure A1. The dependent variable is the 12-month-ahead value of 12-month PCE inflation in columns 1 through 6 and the 24-month ahead value of 12-month inflation in columns 7 through 12. Controls are 3-month PCE inflation, 1-year ahead household inflation expectations, the V/U ratio, the price of oil, the S&P 500 stock price index, real personal disposable income growth (year-over-year), the yield curve spread between 10-year Treasury bonds and the federal funds rate, and a dummy variable indicator of recessions. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

B Probabilities of consecutive same-signed residuals

Assume that ε_t follows the stationary Gaussian AR(1) process

$$\varepsilon_t = \alpha\varepsilon_{t-1} + u_t, \quad |\alpha| < 1, \quad (16)$$

where u_t is white noise. We derive an approximation for the probability of observing k consecutive positive realizations of ε_t . The same approximation applies for k consecutive negative realizations of ε_t .

Since $(\varepsilon_t, \varepsilon_{t-1})$ is jointly normal with correlation coefficient α , the arcsine law for correlated Gaussian variables implies that the unconditional probability that two consecutive residuals will have the same sign is

$$P(\varepsilon_t > 0, \varepsilon_{t-1} > 0) = \frac{1}{4} + \frac{1}{2\pi} \arcsin(\alpha), \quad (17)$$

which is often referred to as Sheppard's formula (Sheppard (1899), Van Vleck and Middleton (1966)). Because the stationary Gaussian distribution is symmetric,

$$P(\varepsilon_{t-1} > 0) = \frac{1}{2}. \quad (18)$$

Hence, the conditional probability that an already realized positive residual is followed by another positive residual is twice the unconditional probability given by Sheppard's formula, as shown below.

$$p = P(\varepsilon_t > 0 \mid \varepsilon_{t-1} > 0) = \frac{1}{2} + \frac{1}{\pi} \arcsin(\alpha). \quad (19)$$

The sign process of a Gaussian AR(1) model is not identical to a first-order Markov chain. This is because the AR(1) model has a continuous state space, whereas the Markov chain has a discrete state space.²⁰ Nevertheless, a convenient approximation is obtained by representing the sign process as a symmetric two-state Markov chain with transition probability p , which matches the exact one-step sign persistence probability implied by the AR(1) model. Under this approximation,

$$P(\varepsilon_t > 0, \varepsilon_{t+1} > 0, \dots, \varepsilon_{t+k-1} > 0) \approx P(\varepsilon_t > 0) p^{k-1}. \quad (20)$$

Since $P(\varepsilon_t > 0) = 1/2$, the probability of k consecutive positive realizations is approxi-

²⁰Timmermann (2000) derives the analytical moments of a 2-state Markov switching model that can be used to match the corresponding moments of an AR(1) model.

mately

$$P(\varepsilon_t > 0, \varepsilon_{t+1} > 0, \dots, \varepsilon_{t+k-1} > 0) \approx \frac{1}{2} \left(\frac{1}{2} + \frac{1}{\pi} \arcsin(\alpha) \right)^{k-1}. \quad (21)$$