Permanent and Transitory Shocks to the Expected Inflation Term Structure

Fabio Gómez-Rodríguez*

February 2026

Abstract

What (de-)anchors inflation expectations? This paper studies the long-run dynamics of inflation expectations by modeling the entire expected inflation term structure as a single curve-valued (functional) time series. We document persistent low-frequency movements across maturities that are inconsistent with stationarity at the functional level. Motivated by this evidence, we model the term structure as a functional autoregressive process with a unit root and decompose innovations into permanent and transitory components using long-run restrictions in the spirit of Blanchard and Quah, adapted to a functional setting. Empirically, we find that long-horizon inflation expectations are driven exclusively by permanent shocks: transitory disturbances have no effect at the far end of the curve. Permanent shocks therefore represent belief re-anchoring episodes that shift the long-run inflation anchor, while transitory shocks generate short-lived deviations. Although both shocks have similar short-run effects on inflation and real activity when normalized by their impact on inflation, monetary policy responds more strongly and persistently to permanent shifts in expectations. Large permanent shocks coincide with episodes widely interpreted as shifts in the inflation regime or in monetary policy credibility. Moreover, permanent shocks are positively correlated with monetary policy and oil supply surprises, reinforcing their interpretation as belief re-anchoring disturbances. Taken together, the evidence suggests that inflation expectations are anchored by monetary policy credibility and de-anchored by shocks that alter perceptions of the inflation regime, rather than by conventional cyclical demand fluctuations.

JEL classification: C32, C14, E31, E52, E58

Keywords: inflation expectations, expected inflation term structure, functional time series, permanent and transitory shocks, long-run identification, monetary policy responses

*Lehigh University and Central Bank of Costa Rica. Email: fabio.gomez-rodriguez@lehigh.edu. The views expressed in this paper are those of the author and do not necessarily reflect those of any institution.

1 Introduction

What (de-)anchors inflation expectations? Which shocks permanently re-anchor long-run inflation beliefs, and what are the macroeconomic and policy consequences of such belief shifts? This paper addresses these questions by studying the long-run dynamics of inflation expectations across the entire term structure. Inflation expectations play a central role in modern macroeconomics and monetary policy, shaping wage-setting behavior, price dynamics, and the transmission of policy actions. Over recent decades, central banks have placed increasing emphasis on anchoring inflation expectations—particularly at medium and long horizons—as a cornerstone of monetary credibility. Yet much of the empirical literature treats inflation expectations as stationary objects fluctuating around a stable long-run mean, or analyzes expectations at a small number of discrete maturities in isolation.

Using the Federal Reserve Bank of Cleveland's model-based estimates of expected inflation—which recover the full term structure from short to very long maturities—we document pronounced persistence and low-frequency movements that are difficult to reconcile with stationarity at the level of the entire curve. These estimates are constructed from nominal Treasury yields, inflation-indexed securities, and inflation swap data within a no-arbitrage term structure framework and therefore reflect model-implied inflation expectations derived from asset prices rather than survey responses at a limited set of maturities. The availability of horizons up to thirty years provides a model-based window into the behavior of long-run inflation beliefs. Across maturities as long as ten to thirty years, expectations display gradual but persistent level shifts, suggesting that inflation expectations behave as a slow-moving macroeconomic state that can be re-anchored over time rather than as stationary forecasts subject only to transitory noise.

A central empirical result emerges from this perspective: long-horizon inflation expectations are driven exclusively by permanent shocks. Transitory disturbances have no effect at the far end of the curve. Movements in long-run expectations therefore reflect belief re-anchoring episodes rather than conventional cyclical fluctuations. This finding provides a direct answer to the anchoring question: inflation expectations are anchored or de-anchored by shocks that alter long-run beliefs about the inflation regime.

These features of the data have important implications for modeling and identification. If inflation expectations are shaped by persistent belief revisions, then expectation shocks differ fundamentally depending on whether they alter long-run beliefs or generate only short-lived deviations. Identification strategies based solely on short-run dynamics are therefore insufficient. Instead, the presence of non-stationarity provides a natural basis for distinguishing between shocks with permanent effects on the expectations curve and shocks whose effects dissipate over time, following the logic of long-run identification in the spirit of Blanchard and Quah (1993).

To motivate this approach, Appendix C reports complementary maturity-by-maturity evidence based on standard univariate time-series tools. Across short, medium, and long maturities, we document systematic rejections of stationarity and highly persistent level movements. At the same time, unit-root test outcomes vary across horizons, deterministic specifications, and sample windows. Rather than undermining the presence of persistence, this heterogeneity highlights the limitations of a purely scalar perspective and points to the presence of common low-frequency components operating at the level of the term structure.

These limitations motivate modeling the term structure of inflation expectations as a single curve-valued time series evolving jointly over time. We adopt a functional time-series framework in which inflation expectations are treated as random functions over maturity. The empirical evidence supports a parsimonious representation in which the term structure is driven by a small number of low-frequency components. In the baseline specification, we find a single common stochastic trend. Economically, this corresponds to a dominant long-run belief about inflation that can be permanently re-anchored by certain shocks, alongside transitory forces that generate short-lived deviations around that trend.

Related literature and contribution. This paper makes three main contributions.

First, we document and formally characterize non-stationarity in the entire term structure of inflation expectations. Rather than assessing persistence horizon by horizon, we show that expectations are driven by a small number of common low-frequency components operating at the level of the curve.

Second, we extend the logic of permanent–transitory decompositions and long-run identification to a functional setting. By applying long-run restrictions to the term structure viewed as an infinite-dimensional object, we identify belief re-anchoring shocks that permanently shift the long-run inflation anchor and distinguish them from transitory expectation disturbances.

Third, we provide new evidence on the economic content and macroeconomic consequences of expectation shocks identified through their long-run effects. Permanent expectation shocks align with episodes widely interpreted as shifts in the inflation regime or in monetary policy credibility and are positively correlated with monetary policy and oil supply surprises. In contrast, transitory shocks have no effect at long maturities. While both types of shocks generate similar short-run responses of inflation and real activity when normalized by their impact on inflation, monetary policy responds more strongly and persistently to permanent shifts in expectations. Taken together, these findings suggest that inflation expectations are anchored by monetary policy credibility and de-anchored by shocks that alter perceptions of the inflation regime, rather than by conventional cyclical demand fluctuations.

The remainder of the paper develops and applies this framework in four steps. Section 2 documents persistence and non-stationarity in the term structure of inflation expectations. Section 3 introduces the econometric methodology and formalizes the identification of permanent and transitory expectation shocks within a functional time-series framework. Section 5 interprets the identified shocks by relating them to external macroeconomic and policy news measures. Section 6 examines the macroeconomic implications of these shocks, focusing on inflation, labor market outcomes, and monetary policy responses. The final section concludes.

2 Persistence and Non-Stationarity in the Term Structure of Inflation Expectations

This section documents a basic empirical fact that motivates the rest of the analysis: inflation expectations at short, medium, and long horizons display persistent level movements that are inconsistent with mean reversion. Visual inspection reveals long episodes of gradual re-anchoring associated with major macroeconomic regimes, including the disinflation of the early 1980s, the Great Moderation, and the post–Global Financial Crisis period. Even at very long horizons, expectations respond to shocks in a manner that appears small in magnitude but persistent over time.

2.1 Visual Evidence

Figure 1 plots market-based inflation expectations[1] at horizons of one, three, ten, and thirty years from 1982 to 2025. Across all maturities, inflation expectations display slow-moving and highly persistent level shifts, with little evidence of reversion toward a fixed long-run mean. The strong co-movement across horizons indicates the presence of common low-frequency forces shaping inflation beliefs over time. Importantly, even expectations at very long horizons exhibit small but durable changes, consistent with gradual belief re-anchoring rather than purely transitory fluctuations.

[1] Data from the Federal Reserve Bank of Cleveland Federal Reserve Bank of Cleveland (2025).

Figure 1: Inflation Expectations Across Horizons

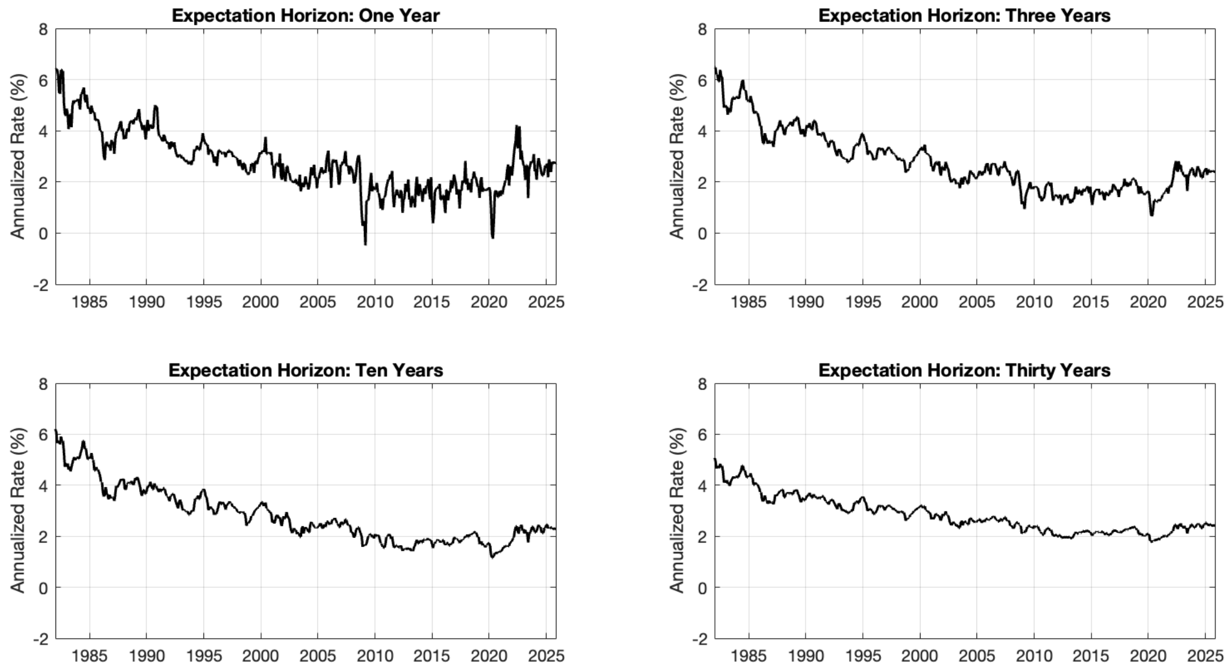

Note: The figure plots annualized market-based inflation expectations at horizons of one, three, ten, and thirty years from January 1982 to November 2025. Expectations are constructed using a no-arbitrage term structure model based on nominal Treasury yields and inflation-indexed securities, as reported by the Federal Reserve Bank of Cleveland. The series exhibit persistent, low-frequency movements across all horizons.

2.2 Formal Statistical Evidence

The visual patterns documented above are consistent with a view of inflation expectations as highly persistent objects shaped by slow-moving belief revisions. As a first, transparent benchmark, Appendix C reports maturity-by-maturity evidence based on standard univariate stationarity and unit-root tests. Taken together, these results provide suggestive support for the idea that inflation expectations exhibit persistent level shifts that are difficult to reconcile with rapid mean reversion. While the outcomes of scalar unit-root tests vary across horizons, sample lengths, and deterministic specifications, the overall pattern is easy to reconcile with the presence of common low-frequency forces operating across maturities rather than with purely horizon-specific dynamics.

This heterogeneity across individual maturities is not a drawback of the evidence, but rather highlights the limitations of a purely scalar perspective. Univariate tests are not designed to detect stochastic trends that are shared across horizons, nor do they require uniform unit-root behavior at each maturity when expectations are jointly determined. Instead, the maturity-by-maturity results naturally motivate a framework in which persistence is assessed at the level of the entire term structure.

Accordingly, the core statistical evidence in this paper relies on functional time-series methods that formally test for non-stationarity in the term structure of inflation expectations as a whole. Within this framework, we directly assess the dimension of the non-stationary space and test for the presence of common stochastic trends driving expectations across horizons. As shown in Section D, these functional unit-root tests confirm the presence of persistent shifts in inflation expectations that operate at the level of the entire curve, providing a coherent statistical foundation for distinguishing between permanent and transitory expectation shocks.

Importantly, the presence of persistent stochastic trends does not imply that inflation expectations are unanchored or unstable. Rather, it implies that anchoring operates through the magnitude of permanent belief revisions rather than through rapid mean reversion. Inflation expectations can therefore remain tightly anchored in economic terms while still exhibiting non-stationary dynamics, with small but durable shifts reflecting gradual changes in beliefs about the inflation regime.

2.3 Economic Interpretation

Persistent shifts in inflation expectations have implications that extend beyond statistical characterization. They suggest that expectation anchoring operates through the scale of permanent belief revisions rather than through rapid mean reversion, that some policy actions re-anchor beliefs about the inflation regime itself, and that expectation shocks may act as slow-moving state variables with long-lived macroeconomic effects. These considerations motivate an empirical framework that distinguishes permanent from transitory expectation shocks and studies their economic consequences at the level of the entire term structure.

3 Econometric Methodology

Scope of the methodological exposition. The econometric framework employed in this paper builds directly on the theory of non-stationary functional autoregressive processes developed in Chang et al. (2023). That work provides a comprehensive treatment of functional unit roots, the associated decomposition into permanent and transitory subspaces, and the asymptotic properties of the relevant estimators.

The purpose of the present section is not to reproduce or extend those theoretical results, but to summarize only the essential concepts and assumptions required to implement the methodology in the context of inflation expectations and to interpret the empirical findings. Readers interested in formal proofs, technical conditions, and the broader theoretical foundations are referred to Chang et al. (2023).

This section presents the econometric framework used to model the dynamics of the term structure of inflation expectations. The central empirical feature motivating our approach is the presence of persistent, low-frequency movements that are shared across maturities. To capture these joint dynamics, we model inflation expectations as a non-stationary functional time series driven by a common stochastic trend.

The exposition proceeds in three steps. First, we formalize inflation expectations as random functions evolving in a separable Hilbert space. Second, we introduce a functional autoregressive (FAR) model with unit roots and discuss its implications for decomposing the term structure into permanent and transitory components. Third, we describe the finite-dimensional projection used for estimation, interpretation, and empirical implementation.

3.1 Inflation Expectations as Random Functions

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, and let $H = L^2(I)$ denote the separable Hilbert space of square-integrable functions defined on the horizon interval $I = [1, 30]$, measured in years.[2] At each date t, the term structure of inflation expectations is represented by a random function

$$e_t(\tau) \in H,$$

where τ indexes the forecast horizon. The sequence $\{e_t\}_{t\in\mathbb{Z}}$ thus forms a functional time series.

The Hilbert space structure of H is endowed with the inner product

$$\langle f, g\rangle = \int_1^{30} f(\tau)g(\tau)\, d\tau,$$

and the associated norm $\|f\| = \langle f, f\rangle^{1/2}$. This framework allows us to analyze inflation expectations jointly across horizons, rather than treating individual maturities in isolation.

Because H is separable, there exists a countable orthonormal basis $\{v_j\}_{j=1}^{\infty}$ such that any $f \in H$ admits the expansion

$$f = \sum_{j=1}^{\infty} \alpha_j v_j.$$

This representation provides the foundation for dimension reduction and estimation while preserving the infinite-dimensional nature of the object of interest.

3.2 A Functional Autoregressive Model with Unit Roots

We model the dynamics of the expected inflation term structure using a functional autoregressive process of order one.[3] The restriction to a first-order process is adopted for convenience of exposition; higher-order functional autoregressions can be written in companion form without loss of generality Chang et al. (2023).

$$e_t = Ae_{t-1} + \varepsilon_t, \tag{1}$$

where $A : H \to H$ is a bounded linear operator and $\{\varepsilon_t\}$ is an H-valued white noise process with mean zero and covariance operator Σ.

[2] For the inflation expectations constructed by Aruoba (2020), the available horizon range is $I = [0.25, 10]$.

[3] For intuition, it is helpful to think of (1) as the infinite-dimensional analogue of a vector autoregression. In this interpretation, the operator A plays the role of a very large (in fact, infinite-dimensional) coefficient matrix that maps the entire past term structure into its current shape, while preserving the joint dynamics across horizons.

The operator A governs the persistence and propagation of inflation expectations across horizons, while ε_t captures new information arriving at date t.

Our empirical analysis provides strong evidence that inflation expectations are non-stationary at the functional level and are driven by a small number of common stochastic trends. Accordingly, we impose the following conditions.

Assumption 1.

(a) The operator A is compact.

(b) The eigenvalue 1 belongs to the spectrum of A, with finite multiplicity.

(c) The innovations $\{\varepsilon_t\}$ are i.i.d. in H with $\mathbb{E}\|\varepsilon_t\|^4 < \infty$ and covariance operator Σ.

Compactness of A ensures that its spectrum consists of isolated eigenvalues with zero as the only accumulation point, a property that plays a central role in the identification of permanent and transitory components. The presence of the unit eigenvalue implies that the process admits a Beveridge–Nelson–type decomposition at the functional level.

Using the testing procedure developed by Chang et al. (2016), which exploits the asymptotic behavior of the sample variance operator, we test sequentially for the number of unit roots in the functional process. The results reject the presence of more than one unit root while failing to reject the null of exactly one unit root. We therefore model the expected inflation term structure as a functional process driven by a single common stochastic trend.

3.3 Permanent and Transitory Subspaces

Under Assumption 1, the space H admits a decomposition into permanent and transitory subspaces,

$$H = H_P \oplus H_T,$$

where H_P is the finite-dimensional eigenspace associated with the unit eigenvalue of A, and H_T collects all remaining components.

Let Π_P and Π_T denote the corresponding Riesz projection operators,

$$\Pi_P = \frac{1}{2\pi i} \oint_{\Gamma_P} (\lambda I - A)^{-1} d\lambda, \qquad \Pi_T = \frac{1}{2\pi i} \oint_{\Gamma_T} (\lambda I - A)^{-1} d\lambda,$$

where Γ_P encloses the eigenvalue $\lambda = 1$ and Γ_T encloses the remainder of the spectrum. These projections are mutually orthogonal, invariant under A, and yield a unique decomposition

$$e_t = e_t^{(P)} + e_t^{(T)},$$

with $e_t^{(P)} = \Pi_P e_t$ capturing permanent belief shifts and $e_t^{(T)} = \Pi_T e_t$ capturing transitory fluctuations.

Economically, the permanent component represents slow-moving re-anchoring of long-run inflation beliefs that affects the entire term structure, while the transitory component reflects short-lived deviations driven by temporary shocks.

3.4 Finite-Dimensional Approximation and Estimation

Although the process e_t evolves in an infinite-dimensional space, empirical implementation requires a finite-dimensional approximation. We therefore project the data onto the leading eigenfunctions of the covariance operator Σ of the innovations.

Let $\{(\lambda_j, v_j)\}_{j\geq 1}$ denote the eigenvalue–eigenfunction pairs of Σ, ordered so that $\lambda_1 \geq \lambda_2 \geq \cdots$. In the data, the first few eigenfunctions explain the vast majority of total variation in the term structure of inflation expectations. In particular, the first three components account for over 99% of total variance, indicating a strong empirical concentration of variability in low-dimensional functional modes.

To implement the model, we approximate e_t using its first m functional principal components and estimate the resulting m-dimensional score process. The truncation dimension m and VAR lag order p are selected jointly using a rolling one-step-ahead forecast evaluation procedure. Specifically, for each candidate pair (m, p), we estimate a VAR(p) model on the first m score series and compute expanding-window forecast errors over the test sample. The baseline specification corresponds to the (m, p) pair minimizing the average root mean squared forecast error across components.

This predictive selection approach follows the methodology developed in Chang et al. (2026), where similar functional autoregressive models are estimated using forecast-based truncation criteria.

4 Modeling Inflation Expectations as a Functional Time Series

This section develops the econometric framework used to study the dynamics of inflation expectations across the entire term structure. Motivated by the evidence in Section 2, which documents persistent and highly co-moving behavior across maturities, we model inflation expectations as a functional time series: a sequence of curves indexed by time, each describing expected inflation over a continuum of horizons. This approach allows us to formally characterize the long-run behavior of the expectations curve, to assess the presence of common stochastic trends, and to provide a foundation for identifying permanent and transitory shocks to inflation expectations.

4.1 The Functional Representation of Inflation Expectations

We begin by formalizing the term structure of inflation expectations as a functional object. At each point in time, inflation expectations are observed across a range of horizons, forming a smooth curve that maps forecast horizons into expected inflation rates. Treating these observations as discretizations of an underlying function allows us to study the joint dynamics of expectations across maturities, rather than analyzing a finite set of horizons in isolation. We specify the function space in which these curves reside and describe the properties required for subsequent analysis.

4.2 Random Functions and Functional Time Series

Once inflation expectations are represented as functions, their evolution over time can be modeled as a stochastic process taking values in a function space. In this framework, each realization of the inflation expectations curve is a random function, and the sequence of such curves constitutes a functional time series. This perspective generalizes multivariate time-series analysis to infinite-dimensional settings and provides the probabilistic foundation for defining dependence, persistence, and long-run behavior at the level of the entire expectations curve.

4.3 Functional Autoregressive Processes

To capture the dynamic dependence of inflation expectations over time, we model the functional time series using a functional autoregressive (FAR) process. The FAR framework extends linear autoregressive models to function-valued data by allowing the current expectations curve to depend on past curves through linear operators. This representation provides a flexible yet tractable way to describe how shocks propagate across horizons and over time, while preserving the joint structure of the term structure.

4.4 Stationarity, Non-Stationarity, and Common Trends

A central question for the analysis is whether inflation expectations are stationary or exhibit persistent non-stationary behavior. In the functional setting, non-stationarity can arise through the presence of one or more stochastic trends that affect the entire curve. We characterize the stationary and non-stationary components of FAR processes and discuss how these components can be separated using orthogonal projections. Importantly, FAR processes defined with compact operators admit non-stationary behavior only in finite-dimensional subspaces, a property that allows us to interpret persistent movements in inflation expectations as arising from a small number of common stochastic trends.

4.5 Permanent and Transitory Spaces and Shock Identification

Conditional on the estimated functional autoregressive model and the number of common stochastic trends, we decompose the dynamics of inflation expectations into permanent and transitory components. The permanent space captures shocks that have lasting effects on the level of the expectations curve, while the transitory space contains shocks whose effects dissipate over time. This decomposition provides the basis for identifying permanent and transitory inflation-expectations shocks using long-run restrictions in the spirit of Blanchard and Quah (1993). In the context of inflation expectations, permanent shocks are interpreted as belief re-anchoring events, whereas transitory shocks generate short-lived fluctuations around the prevailing belief anchor.

Figure 2: Response of the Expected Inflation Term Structure to Permanent and Transitory Shocks

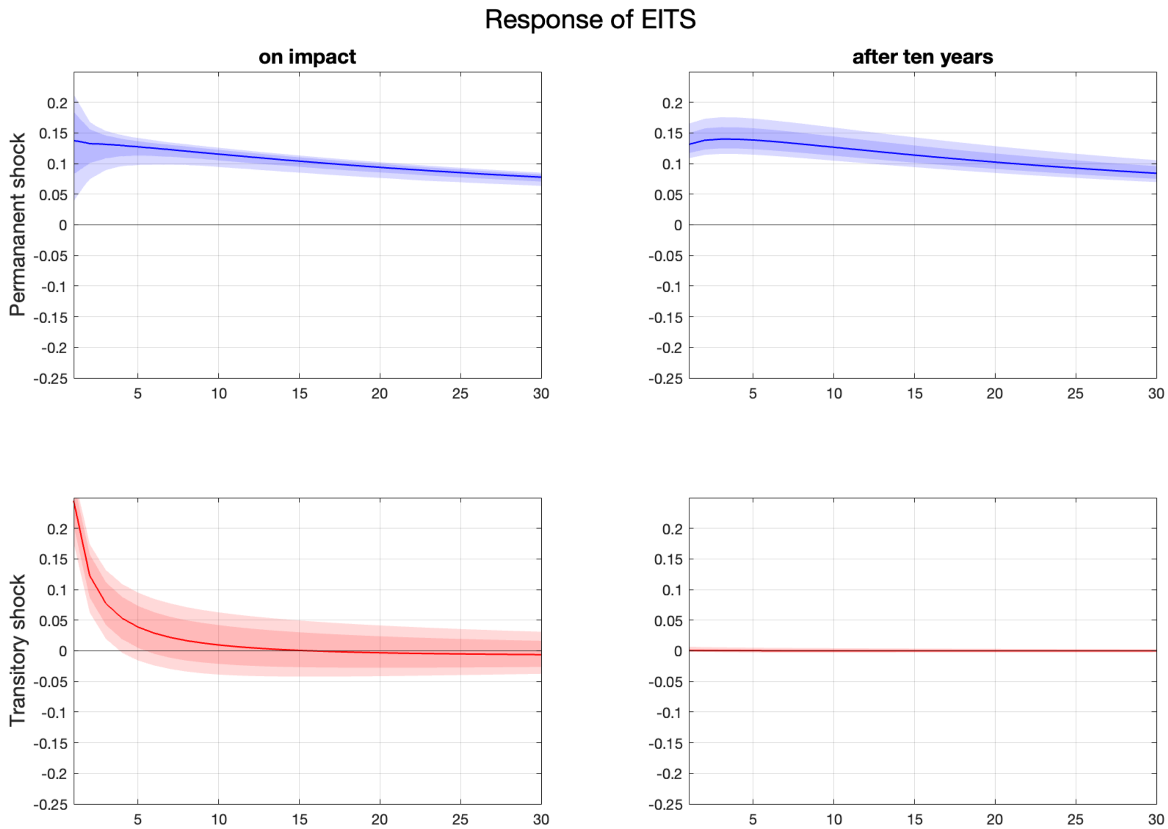

Note: The figure plots the response of the expected inflation term structure (EITS) to a one-standard-deviation permanent shock (top panels) and a transitory shock (bottom panels), evaluated on impact (left column) and after ten years (right column). Responses are shown as functions of the expectation horizon, measured in years. Solid lines denote point estimates. Shaded areas represent pointwise bootstrap confidence bands: darker bands correspond to the 68 percent interval and lighter bands to the 90 percent interval. Permanent shocks induce persistent shifts in expected inflation across all horizons, while transitory shocks generate short-lived responses that vanish at long horizons.

The distinction between these two types of shocks is illustrated by the dynamic responses of the expected inflation term structure. Permanent shocks induce persistent shifts in expected inflation across all horizons, including very long maturities, with effects that remain economically meaningful even a decade after the shock. In contrast, transitory shocks primarily affect short- and medium-term expectations and have negligible long-run effects on the term structure. This horizon-dependent behavior provides direct empirical support for the long-run identification strategy and clarifies the economic interpretation of the permanent and transitory spaces.

5 Interpreting the Shocks: What Moves Them?

This section provides economic interpretation for the permanent and transitory inflation-expectations shocks identified in the previous analysis. We proceed in two complementary steps. First, we examine the historical behavior of these shocks, focusing on episodes in which they attain unusually large realizations. By relating these episodes to well-known macroeconomic and policy developments, we assess whether permanent shocks tend to coincide with events plausibly associated with changes in long-run inflation beliefs—such as shifts in monetary policy credibility, sustained inflationary pressures, or major fiscal and supply-side disturbances—while transitory shocks align more closely with short-lived inflation news or temporary macroeconomic shocks.

Second, we complement this historical narrative with a more systematic analysis based on external macroeconomic and policy-related news measures. These external expectation and policy shocks are taken from the harmonized survey-based dataset developed by Adams and coauthors. For attribution and detailed descriptions of the construction of these measures, see Adams and Barrett (2025) and Adams and Matthes (2026). Rather than treating these external shocks as instruments or proposing an alternative identification strategy, this exercise is intended to characterize which types of information tend to be associated with persistent versus transitory movements in the expected inflation term structure (EITS).

Together, the historical and correlation-based evidence helps clarify the economic content of the permanent–transitory decomposition. It provides interpretive guidance on the nature of belief revisions captured by the model and sheds light on the types of macroeconomic and policy developments that are more likely to generate lasting re-anchoring of inflation expectations rather than short-lived deviations around a stable belief anchor.

5.1 Historical Behavior of Permanent and Transitory Shocks

Figure 3 plots the estimated permanent and transitory inflation-expectations shocks over time, with annotations marking the largest positive and negative realizations in each series. Several features stand out.

First, permanent shocks are relatively infrequent but economically sizable. Large realizations tend to be clustered around periods associated with major and persistent changes in the macroeconomic environment. Notable negative permanent shocks occur in the early 1980s, consistent with the disinflationary regime shift associated with the Volcker period, as well as during the Global Financial Crisis in late 2008 and early 2009. More recent negative realizations around 2020 and 2023–2025 coincide with episodes of extraordinary macroeconomic uncertainty, including the onset of the COVID-19 pandemic and the subsequent tightening phase of U.S. monetary policy. On the positive side, permanent shocks appear during episodes such as the late 1980s, the early 2000s, and the post-pandemic inflation surge, plausibly reflecting upward revisions to long-run inflation beliefs rather than short-lived news.

Second, transitory shocks are both more frequent and more symmetric, with large positive and

Figure 3: Estimated Permanent and Transitory Inflation-Expectations Shocks

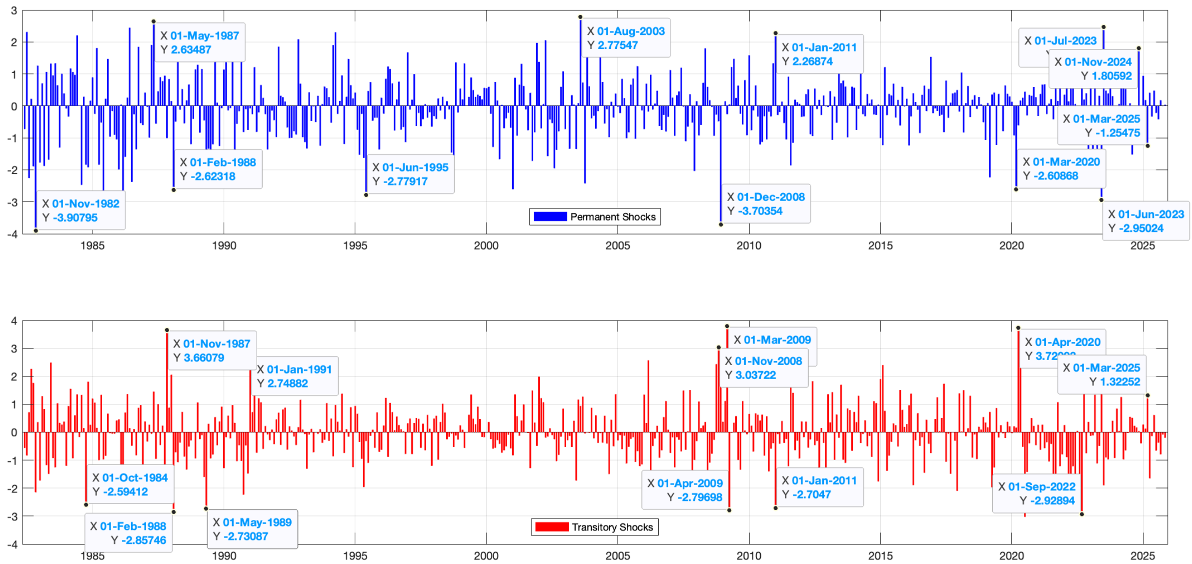

Note: The figure plots the estimated permanent (top panel) and transitory (bottom panel) shocks to the expected inflation term structure (EITS) over time. Shocks are identified from the functional autoregressive model using long-run restrictions that distinguish disturbances with persistent effects on the expectations curve from those whose effects dissipate over time. Selected large positive and negative realizations are annotated to highlight economically relevant episodes. Permanent shocks are interpreted as belief re-anchoring events affecting long-horizon inflation expectations, while transitory shocks capture short-lived fluctuations driven by temporary macroeconomic or policy news.

negative realizations occurring throughout the sample. These shocks often spike during periods of heightened short-run macroeconomic or financial volatility, including inflation surprises, policy announcements, and crisis-related news. For example, large transitory movements are evident around the late 1980s, the Global Financial Crisis, and the COVID-19 episode. Unlike permanent shocks, however, these realizations do not appear to be associated with lasting changes in the level of long-horizon inflation expectations.

Third, the timing and magnitude of the two shocks differ in economically meaningful ways. Episodes that generate large transitory shocks do not necessarily coincide with large permanent shocks, and vice versa. This distinction is particularly evident during periods such as 2009–2011 and the post-2020 recovery, where transitory shocks are sizable and frequent, while permanent shocks are more episodic. This pattern reinforces the interpretation of permanent shocks as belief re-anchoring events—rare but persistent—while transitory shocks capture short-run fluctuations in expectations driven by temporary news.

Overall, the historical behavior of the shocks supports the long-run identification strategy. Permanent shocks align with episodes plausibly associated with shifts in long-run inflation beliefs and policy regimes, whereas transitory shocks reflect short-lived disturbances that leave the long-run expectations curve largely unaffected.

Figure 4 reports correlations between the estimated permanent and transitory EITS shocks and a collection of commonly used macroeconomic news and surprise measures. The top panel focuses on the permanent shock, while the bottom panel reports results for the transitory shock. Shocks are ordered by the magnitude of their correlation, and confidence intervals are shown to illustrate sampling uncertainty.

The permanent EITS shock exhibits positive correlations with a wide range of shocks that plausibly convey information relevant for long-run inflation beliefs. These include oil-price and supply-related shocks, fiscal and tax news with medium- to long-horizon implications, and several monetary-policy and information-type shocks. While individual correlations are moderate in size, the overall pattern suggests that permanent movements in inflation expectations tend to coincide with broad, macroeconomic disturbances that affect the perceived long-run inflation environment rather than short-lived fluctuations. At the same time, shocks commonly interpreted as technology or productivity improvements display negative correlations with the permanent component, consistent with persistent downward revisions in inflation expectations following favorable supply-side developments.

In contrast, the transitory EITS shock displays a markedly different correlation profile. Positive correlations are concentrated among shocks associated with short-run fiscal actions, near-term tax changes, and demand-driven disturbances, while correlations with long-horizon or belief-related shocks are small or negative. Monetary policy and information shocks tend to load negatively on the transitory component, indicating that news which affects long-run inflation beliefs tends to offset, rather than amplify, short-lived movements in expectations. Overall, the transitory shock appears to be associated with temporary reshaping of the expected inflation curve rather than persistent re-anchoring.

Taken together, the correlation evidence reinforces the interpretation of the permanent and transitory shocks as capturing distinct dimensions of expectation dynamics. Permanent shocks align with information that plausibly shifts long-run inflation beliefs, while transitory shocks are more closely related to short-run macroeconomic and policy disturbances. Importantly, no single external shock dominates either component, underscoring that the identified EITS shocks reflect latent combinations of information rather than simple transformations of existing macroeconomic surprise measures.

6 Macroeconomic Implications of Inflation Expectations Shocks

This section studies the macroeconomic and financial effects of permanent and transitory shocks to the expected inflation term structure (EITS). Having characterized the persistence properties and economic content of these shocks, we now examine how they propagate to the real economy, monetary policy, and asset prices. The analysis is conducted using a separate VAR in which the permanent and transitory EITS shocks are used as external instruments, allowing us to trace their dynamic effects without imposing additional structural restrictions on the macroeconomic system.

Figure 4: Correlations between permanent and transitory expected inflation term structure shocks and external macroeconomic shocks

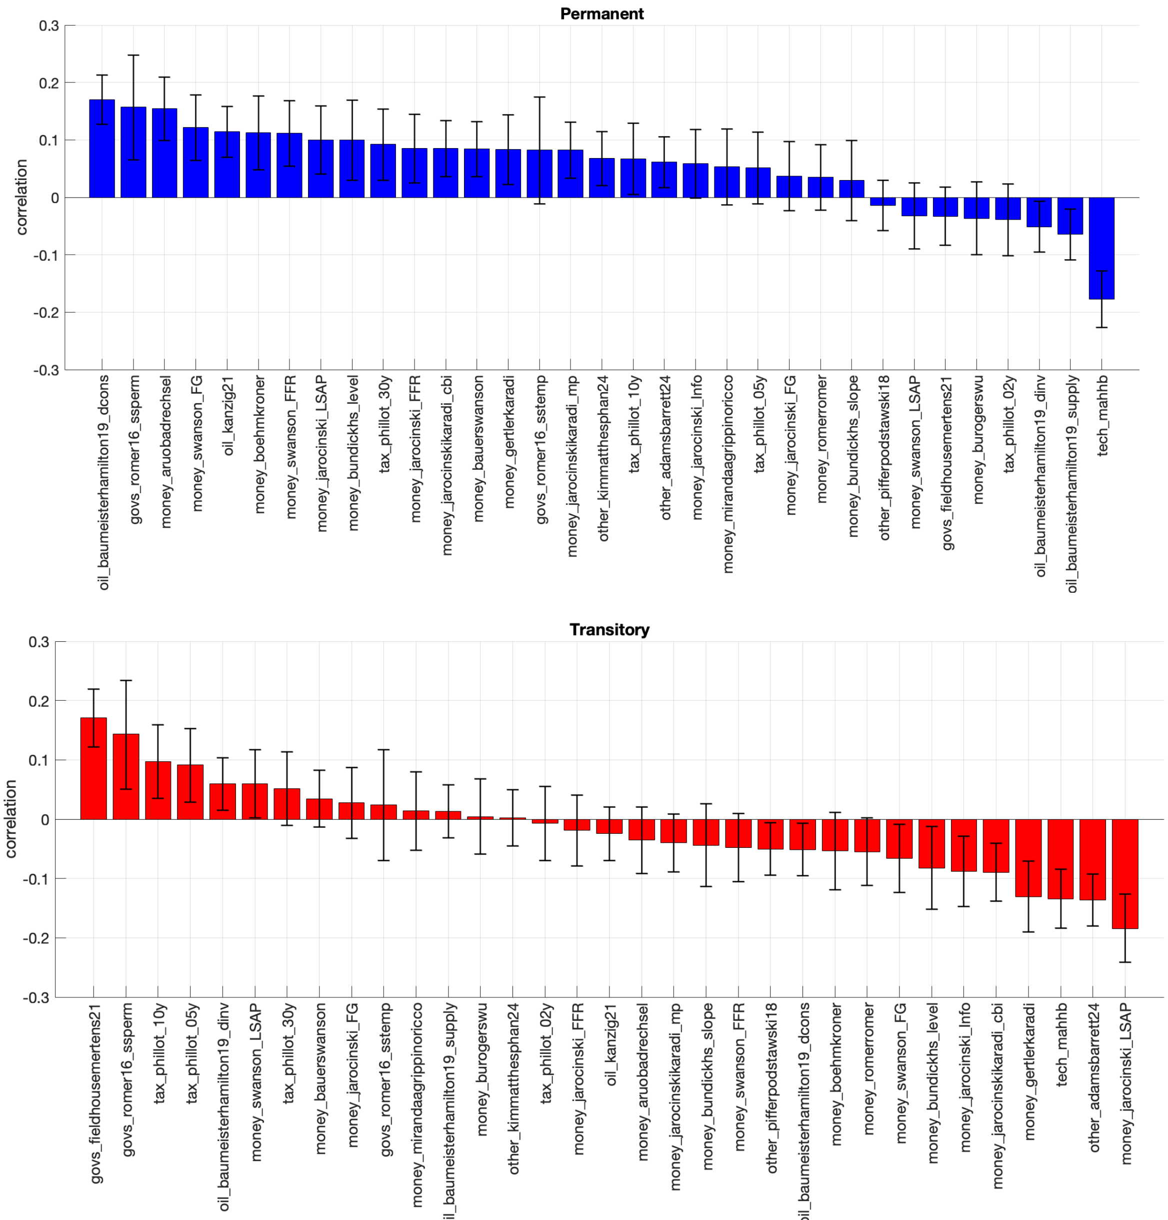

Note: The figure reports pairwise correlations between the estimated permanent and transitory shocks to the expected inflation term structure (EITS) and a set of external macroeconomic, policy, and news-based shocks commonly used in the literature. External shocks (see Adams and Barrett (2025) and Adams and Matthes (2026)). are taken as given and are not used as instruments in the estimation of the EITS shocks. Correlations are computed using standardized series, and confidence intervals are shown to reflect sampling uncertainty. This exercise is intended to provide interpretive evidence on the economic content of the identified shocks rather than to establish causal relationships.

We analyze the responses of inflation, unemployment, and the policy interest rate. Figure 5 reports impulse responses to permanent and transitory expectation shocks, normalized so that inflation increases on impact in both cases. Once scaled in this way, inflation and unemployment display broadly similar short-run dynamics following the two shocks. In both cases, inflation rises on impact and gradually reverts, while unemployment declines and recovers slowly over time. This similarity suggests that, conditional on the initial inflationary impulse, near-term real activity and price dynamics are largely insensitive to whether the underlying expectation shock is permanent or transitory.

In contrast, the response of monetary policy differs sharply across the two shocks. The federal funds rate reacts more strongly and persistently to permanent shifts in inflation expectations than to transitory ones. While transitory expectation shocks generate a smaller and hump-shaped policy response that fades relatively quickly, permanent shocks induce a sustained tightening of policy that remains elevated over the medium run. This divergence indicates that monetary policy responds not only to the level of inflationary pressure but also to the perceived persistence of movements in inflation expectations. In this sense, the decomposition of expectation shocks into permanent and transitory components reveals a dimension of policy behavior that would be difficult to detect using standard scalar measures of expectations.

Taken together, the results highlight the importance of distinguishing between permanent and transitory movements in inflation expectations when assessing macroeconomic dynamics. Although both types of shocks generate similar short-run responses of inflation and unemployment once normalized by their initial impact, they have markedly different implications for monetary policy. Treating the expected inflation term structure as a curve-valued object and decomposing its innovations by persistence therefore provides new insights into how expectations influence policy transmission, even in the absence of large differences in real-side responses.

7 Conclusion

This paper studies the term structure of inflation expectations through the lens of persistence and long-run dynamics. Rather than focusing on individual maturities or summary measures of long-run expectations, we treat the entire expectations curve as the object of interest and analyze its evolution using a functional time-series framework. This perspective allows us to distinguish between shocks that permanently alter the level or shape of the expectations curve and shocks whose effects are inherently transitory.

Our empirical results document pronounced low-frequency movements in inflation expectations that operate at the level of the entire curve. These dynamics are well captured by a small number of nonstationary functional components, while remaining dimensions of the expectations curve display stationary behavior. Importantly, the presence of permanent components does not imply that inflation expectations are unanchored. Instead, our findings suggest that anchoring operates through the magnitude and frequency of permanent belief revisions rather than through rapid mean

Figure 5: Impulse responses of macroeconomic variables to permanent and transitory shocks to the expected inflation term structure

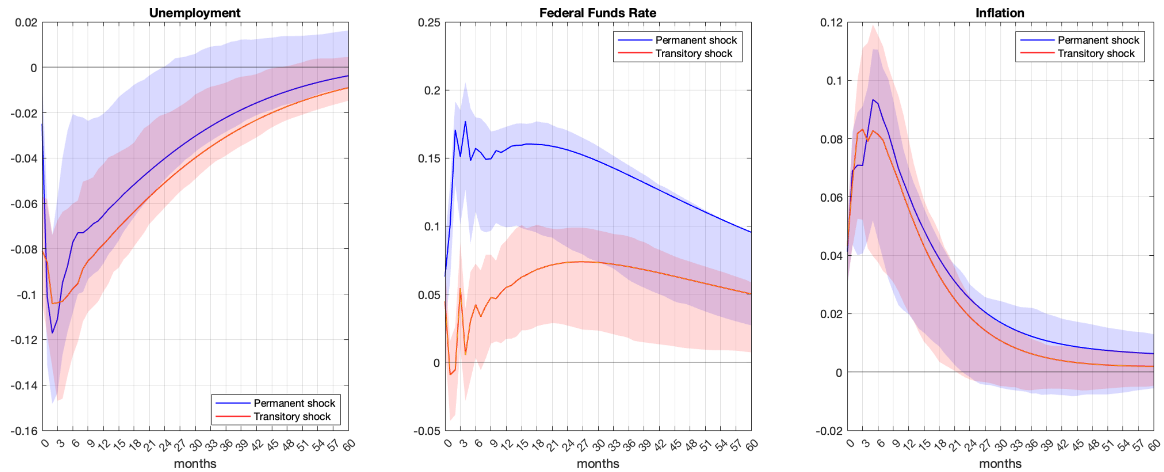

Note: Shaded areas are 95% Monte Carlo confidence intervals for the impulse responses (percentile intervals computed from simulated sample paths and re-estimation of the VAR). Responses correspond to one–standard-deviation (orthogonalized) innovation shocks.

reversion of expectations at each horizon.

By identifying expectation shocks based on their long-run effects on the term structure, we provide a decomposition into permanent and transitory disturbances that is both empirically tractable and economically interpretable. Permanent shocks correspond to persistent shifts in inflation beliefs that affect expectations across maturities, while transitory shocks generate short-lived deviations that primarily reshape the near-term portion of the curve. Correlations with external measures of macroeconomic and policy news further support this interpretation, indicating that the two components are associated with distinct types of information.

The distinction between permanent and transitory expectation shocks proves particularly informative when examining macroeconomic responses. While both types of shocks generate similar short-run movements in inflation and real activity when normalized by their impact on inflation, monetary policy responds more strongly and persistently to shocks that permanently re-anchor expectations. This pattern highlights the importance of persistence, rather than the immediate magnitude of inflationary pressure, in shaping policy responses and provides a new perspective on how central banks react to changes in inflation beliefs.

The framework developed in this paper opens several avenues for future research. First, the functional common-trend approach can be applied to alternative measures of inflation expectations, including survey-based term structures and cross-country data, to study differences in anchoring across institutional and policy regimes. Second, tighter links between the identified expectation shocks and structural macroeconomic disturbances—such as monetary, fiscal, or supply-side shocks—could help clarify the mechanisms underlying persistent belief revisions. Finally, extending the analysis to financial markets would allow for a more comprehensive assessment of how expectation-driven shocks propagate through asset prices and policy transmission channels.

Overall, our results underscore the value of analyzing inflation expectations as a high-dimensional object with rich persistence properties. Viewing the term structure of inflation expectations through a functional and long-run lens provides new insights into how beliefs evolve over time and how central bank credibility is reflected not in the absence of movement, but in the gradual and limited nature of permanent shifts in expectations.

References

Adams, J. and Barrett, P. (2025). What are empirical monetary policy shocks? estimating the term structure of policy news. Technical report, International Monetary Fund.

Adams, J. J. and Matthes, C. (2026). How ricardian are we?

Aruoba, S. B. (2016). Measuring the term structure of inflation expectations. *Review of Economics and Statistics*, 98(4):718–731.

Aruoba, S. B. (2020). Term structures of inflation expectations and real interest rates. *Journal of Business & Economic Statistics*, 38(3):542–553.

Bauer, M. D. and Pflueger, C. E. (2018). International inflation persistence. *Review of Economics and Statistics*, 100(1):1–17.

Beechey, M. J., Johannsen, B. K., and Levin, A. T. (2011). The rise and fall of u.s. inflation persistence. *Journal of Monetary Economics*, 58(2):137–155.

Bernanke, B. S. (2007). Inflation expectations and inflation forecasting. *Speech at the Monetary Economics Workshop of the National Bureau of Economic Research*.

Beveridge, S. and Nelson, C. R. (1981). A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the 'business cycle'. *Journal of Monetary economics*, 7(2):151–174.

Blanchard, O. J. and Quah, D. (1993). The dynamic effects of aggregate demand and supply disturbances: Reply. *The American Economic Review*, 83(3):653–658.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2020). Average inflation targeting and the anchoring of inflation expectations. *Brookings Papers on Economic Activity*.

Chang, Y., Gómez-Rodríguez, F., and Matthes, C. (2026). The influence of fiscal and monetary policies on the shape of the yield curve. *Journal of Economic Dynamics and Control*, 184:105276.

Chang, Y., Hu, B., and Park, J. Y. (2023). Econometric analysis of persistent functional dynamics. Manuscript, submitted. This version: June 2023.

Chang, Y., Kim, C. S., and Park, J. Y. (2016). Nonstationarity in time series of state densities. *Journal of Econometrics*, 192(1):152–167.

Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–2678.

Coibion, O., Gorodnichenko, Y., and Kamdar, R. (2018). Is the phillips curve alive and well after all? inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics*, 10(1):197–232.

Del Negro, M., Giannoni, M. P., and Schorfheide, F. (2015). Inflation in the great recession and new keynesian models. *American Economic Journal: Macroeconomics*, 7(1):168–196.

Federal Reserve Bank of Cleveland (2025). Inflation expectations. https://doi.org/10.26509/frbc-infexp. Accessed December 19, 2025.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle*. Princeton University Press, Princeton, NJ, 2nd edition.

Gonzalo, J. and Granger, C. (1995). Estimation of common long-memory components in cointegrated systems. *Journal of Business & Economic Statistics*, 13(1):27–35.

Gürkaynak, R. S., Sack, B., and Wright, J. H. (2010). The term structure of inflation expectations. *Journal of Monetary Economics*, 57(6):585–595.

Haubrich, J. G., Pennacchi, G., and Ritchken, P. (2012). Inflation expectations, real rates, and risk premia: Evidence from inflation swaps. *Review of Financial Studies*, 25(5):1588–1629.

Manski, C. F. (2004). Measuring expectations. *Econometrica*, 72(5):1329–1376.

Stock, J. H. and Watson, M. W. (2017). Slack and cyclically sensitive inflation. *Journal of Money, Credit and Banking*, 49(S1):293–334.

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, NJ.

A Inflation Expectations: Measurement, Dynamics, and Macroeconomic Role

Inflation expectations play a central role in modern macroeconomic theory and monetary policy. In standard New Keynesian frameworks, expected inflation enters directly into firms' price-setting decisions and households' intertemporal choices, shaping the dynamics of inflation, output, and real interest rates (Woodford, 2003; Galí, 2015). From a policy perspective, the management of inflation expectations—often described in terms of "anchoring"—is widely viewed as a key channel through which central banks influence macroeconomic outcomes and maintain price stability (Bernanke, 2007; **?**). As a result, a large empirical literature seeks to measure inflation expectations, assess their stability over time, and understand how they respond to economic news and policy actions.

Empirical work on inflation expectations relies on two broad classes of measures: survey-based expectations and market-based expectations. Survey measures provide direct information on agents' beliefs but are typically available only at a limited number of horizons and may be affected by reporting biases, rounding behavior, or changes in survey design (Manski, 2004; Coibion and Gorodnichenko, 2015). Market-based measures, by contrast, infer expected inflation from asset prices—most commonly from the spread between nominal Treasury yields and inflation-indexed securities—using no-arbitrage term structure models to separate expected inflation from risk and liquidity premia (Haubrich et al., 2012; Aruoba, 2016). These approaches allow researchers to recover expectations at a wide range of maturities and to study their joint dynamics across horizons. A growing literature uses such measures to analyze the sensitivity of inflation expectations to macroeconomic shocks, monetary policy announcements, and changes in the policy regime (Gürkaynak et al., 2010; Bauer and Pflueger, 2018).

A.1 Persistence, Anchoring, and Long-Run Movements in Expectations

A central question in the literature concerns the persistence of inflation expectations and the degree to which they are anchored over time. Much empirical work assesses anchoring by examining the stability of long-horizon or far-forward expectations, their sensitivity to short-run inflation surprises, or their co-movement with realized inflation (Beechey et al., 2011; Coibion et al., 2018). Evidence of low volatility or weak sensitivity at long horizons is often interpreted as a sign of well-anchored expectations. At the same time, a number of studies document gradual but persistent movements in medium- and long-term expectations, particularly around major macroeconomic events or shifts in the monetary policy environment (Del Negro et al., 2015; Stock and Watson, 2017).

These findings raise important questions about how to interpret persistence in inflation expectations. Persistent movements need not imply unanchored expectations; instead, they may reflect infrequent but economically meaningful revisions to long-run beliefs about inflation (Bordalo et al., 2020). Distinguishing between permanent belief shifts and transitory fluctuations is therefore central to understanding the dynamics of expectations and their macroeconomic consequences. While much of the existing literature approaches this issue using horizon-by-horizon analysis or reducedform regressions, relatively little work explicitly models the joint long-run dynamics of expectations across the entire term structure. This gap motivates approaches that treat inflation expectations as a high-dimensional or functional object and that exploit long-run information to identify shocks with fundamentally different economic interpretations.

B Analysis of the EITS series of Aruoba (2020)

Figure B-1: Inflation Expectations Across Horizons

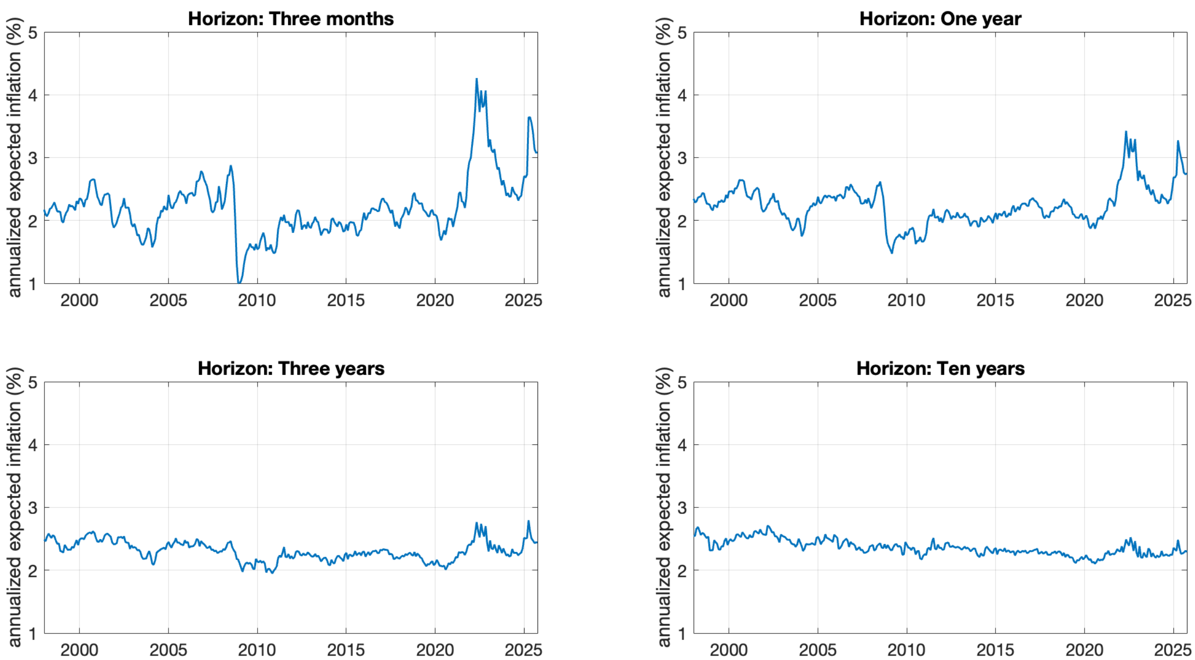

Note: The figure plots annualized market-based inflation expectations at horizons of three months, one year, three years, and ten years from January 1998 to October 2025. The series are taken from the Aruoba–Term Structure of Inflation Expectations (ATSIX), constructed by Aruoba using a no-arbitrage term structure model that jointly exploits information from nominal Treasury yields and Treasury Inflation-Protected Securities (TIPS). The model delivers zero-coupon inflation expectations at multiple horizons in a manner consistent with absence of arbitrage across maturities. Across all horizons, inflation expectations display persistent low-frequency movements, with shorter-horizon expectations exhibiting greater volatility than longer-horizon measures.

C Stationarity and Unit-Root Tests for Individual Inflation Expectations Maturities

This appendix reports a detailed analysis of the time-series properties of inflation expectations at individual maturities. While the main text focuses on the joint dynamics of the term structure using a functional time-series framework, examining persistence at selected horizons using standard scalar methods provides useful and transparent evidence that motivates the functional approach.

C.1 Methodology

For each maturity, we assess persistence using two complementary tests with opposing null hypotheses. First, we apply the Kwiatkowski–Phillips–Schmidt–Shin (KPSS) test, which takes stationarity as the null hypothesis. Rejection of the KPSS null provides direct evidence against stationarity. Second, we apply the augmented Dickey–Fuller (ADF) test, which takes the presence of a unit root as the null hypothesis. Failure to reject the ADF null is consistent with non-stationary, unit-root-type behavior.

Using both tests allows us to assess persistence in a manner that is less sensitive to the maintained null hypothesis and provides a more complete characterization of the scalar time-series properties of expectations at each horizon.

C.2 Deterministic Components

The baseline specification for both tests includes an intercept but excludes a deterministic time trend. This choice is motivated by visual inspection and economic considerations. Inflation expectations display persistent level shifts but do not exhibit systematic linear trends over the sample period. Including a deterministic trend risks attributing low-frequency stochastic movements to deterministic components. Specifications allowing for trend stationarity are reported as robustness checks.

C.3 Lag and Bandwidth Selection

Implementation of the KPSS and ADF tests requires selecting tuning parameters to account for short-run serial correlation. For the KPSS test, the long-run variance bandwidth is chosen using a data-driven autoregressive approximation, with lag order selected by the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC). For the ADF test, the number of lagged differences is selected using the same criteria. This approach avoids ad hoc parameter choices and ensures comparability across maturities.

C.4 Results

Tables 1, 2, and 3 report maturity-by-maturity scalar stationarity and unit-root tests for market-based inflation expectations. Two broad patterns emerge.

First, the KPSS test overwhelmingly rejects the null of level or trend stationarity across maturities. This finding is robust to alternative bandwidth choices and deterministic specifications. Treated as scalar series, inflation expectations do not behave as fluctuations around a fixed mean (or deterministic trend) with purely transitory deviations.

Second, the ADF evidence varies across horizons and sample windows. In shorter subsamples—such as the Aruoba-based sample and the last-20-years Cleveland subsample—the ADF test frequently fails to reject the unit-root null for a nontrivial set of maturities. Such non-rejections are consistent with the well-known low power of unit-root tests against highly persistent stationary alternatives, particularly in shorter samples.

In the full Cleveland sample, the ADF test often rejects the unit-root null at conventional levels. These rejections, however, should not be interpreted as decisive evidence of stationarity. Unit-root tests are sensitive to the inclusion of deterministic components, to near-unit-root behavior, and to structural shifts. Gradual level shifts or regime changes can generate apparent rejections even when the underlying dynamics remain highly persistent and dominated by low-frequency movements. The instability of ADF conclusions across specifications and subsamples therefore cautions against relying exclusively on scalar tests to characterize the persistence of inflation expectations.

Taken together, the KPSS rejections provide systematic evidence against simple scalar stationarity, while the mixed ADF results point to very persistent low-frequency dynamics that are not easily summarized by horizon-by-horizon tests. This motivates the next step of the analysis: rather than deciding the stochastic properties of the term structure one maturity at a time, we model the entire cross-section of horizons as a functional time series. This allows us to assess directly whether the term structure as a curve contains a non-stationary component—and how many—which is the relevant question for identifying permanent versus transitory expectation shocks and studying their macroeconomic implications.

Table 1: Scalar stationarity and unit-root tests across maturities (Cleveland Fed inflation expectations; Full sample)

Horizon	T	KPSS (null: stationary)				ADF (null: unit root)			
		Constant		Trend		Constant (ARD)		Trend (TS)	
(years)		AIC: bw / p	BIC: bw / p	AIC: bw / p	BIC: bw / p	AIC: lag / p	BIC: lag / p	AIC: lag / p	BIC: lag / p
1	527	2 / 0.0100	1 / 0.0100	2 / 0.0100	1 / 0.0100	1 / 0.0035	0 / 0.0010	0 / 0.0010	0 / 0.0010
2	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0178	0 / 0.0178	0 / 0.0106	0 / 0.0106
3	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0236	0 / 0.0236	0 / 0.0378	0 / 0.0378
4	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0229	0 / 0.0229	0 / 0.0495	0 / 0.0495
5	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0218	0 / 0.0218	0 / 0.0503	0 / 0.0503
6	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0210	0 / 0.0210	0 / 0.0475	0 / 0.0475
7	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0206	0 / 0.0206	0 / 0.0440	0 / 0.0440
8	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0204	0 / 0.0204	0 / 0.0408	0 / 0.0408
9	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0203	0 / 0.0203	0 / 0.0381	0 / 0.0381
10	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0203	0 / 0.0203	0 / 0.0357	0 / 0.0357
11	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0204	0 / 0.0204	0 / 0.0338	0 / 0.0338
12	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0204	0 / 0.0204	0 / 0.0322	0 / 0.0322
13	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0205	0 / 0.0205	0 / 0.0307	0 / 0.0307
14	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0207	0 / 0.0207	0 / 0.0295	0 / 0.0295
15	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0208	0 / 0.0208	0 / 0.0285	0 / 0.0285
16	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0210	0 / 0.0210	0 / 0.0276	0 / 0.0276
17	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0211	0 / 0.0211	0 / 0.0268	0 / 0.0268
18	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0213	0 / 0.0213	0 / 0.0260	0 / 0.0260
19	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0215	0 / 0.0215	0 / 0.0253	0 / 0.0253
20	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0217	0 / 0.0217	0 / 0.0247	0 / 0.0247
21	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0219	0 / 0.0219	0 / 0.0243	0 / 0.0243
22	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0221	0 / 0.0221	0 / 0.0238	0 / 0.0238
23	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0223	0 / 0.0223	0 / 0.0234	0 / 0.0234
24	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0226	0 / 0.0226	0 / 0.0230	0 / 0.0230
25	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0228	0 / 0.0228	0 / 0.0226	0 / 0.0226
26	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0231	0 / 0.0231	0 / 0.0223	0 / 0.0223
27	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0233	0 / 0.0233	0 / 0.0219	0 / 0.0219
28	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0236	0 / 0.0236	0 / 0.0216	0 / 0.0216
29	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0238	0 / 0.0238	0 / 0.0213	0 / 0.0213
30	527	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0241	0 / 0.0241	0 / 0.0211	0 / 0.0211

Notes: The table reports p-values for the KPSS test (null: stationarity) and the augmented Dickey–Fuller (ADF) test (null: unit root) applied to Cleveland Fed market-based inflation expectations at horizons from 1 to 30 years (monthly sample, $T = 527$). For KPSS, the long-run variance bandwidth is selected via AIC or BIC using an AR(p) approximation; entries report *bandwidth / p-value*. For ADF, the number of lagged differences is selected via AIC or BIC based on the ADF regression; entries report *lags / p-value*. "Constant" corresponds to a specification with an intercept only; "Trend" allows for a linear time trend. Reported KPSS p-values are truncated at 0.0100 by the software.

Table 2: Scalar stationarity and unit-root tests across maturities (Cleveland Fed inflation expectations; last 20 years)

		KPSS (null: stationary)				ADF (null: unit root)			
Horizon	T	Constant		Trend		Constant (ARD)		Trend (TS)	
(years)		AIC: bw / p	BIC: bw / p	AIC: bw / p	BIC: bw / p	AIC: lag / p	BIC: lag / p	AIC: lag / p	BIC: lag / p
1	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0010	0 / 0.0010	0 / 0.0010	0 / 0.0010
2	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0039	0 / 0.0039	0 / 0.0119	0 / 0.0119
3	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0200	0 / 0.0200	0 / 0.0705	0 / 0.0705
4	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0408	0 / 0.0408	0 / 0.1414	0 / 0.1414
5	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0564	0 / 0.0564	0 / 0.1946	0 / 0.1946
6	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0663	0 / 0.0663	0 / 0.2385	0 / 0.2385
7	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0724	0 / 0.0724	0 / 0.2665	0 / 0.2665
8	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0761	0 / 0.0761	0 / 0.2847	0 / 0.2847
9	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0785	0 / 0.0785	0 / 0.2973	0 / 0.2973
10	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0800	0 / 0.0800	0 / 0.3065	0 / 0.3065
11	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0812	0 / 0.0812	0 / 0.3140	0 / 0.3140
12	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0820	0 / 0.0820	0 / 0.3198	0 / 0.3198
13	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0826	0 / 0.0826	0 / 0.3248	0 / 0.3248
14	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0831	0 / 0.0831	0 / 0.3293	0 / 0.3293
15	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0836	0 / 0.0836	0 / 0.3334	0 / 0.3334
16	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0841	0 / 0.0841	0 / 0.3373	0 / 0.3373
17	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0846	0 / 0.0846	0 / 0.3412	0 / 0.3412
18	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0851	0 / 0.0851	0 / 0.3450	0 / 0.3450
19	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0857	0 / 0.0857	0 / 0.3488	0 / 0.3488
20	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0863	0 / 0.0863	0 / 0.3526	0 / 0.3526
21	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0870	0 / 0.0870	0 / 0.3565	0 / 0.3565
22	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0877	0 / 0.0877	0 / 0.3605	0 / 0.3605
23	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0884	0 / 0.0884	0 / 0.3646	0 / 0.3646
24	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0892	0 / 0.0892	0 / 0.3687	0 / 0.3687
25	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0900	0 / 0.0900	0 / 0.3730	0 / 0.3730
26	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0910	0 / 0.0910	0 / 0.3773	0 / 0.3773
27	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0919	0 / 0.0919	0 / 0.3818	0 / 0.3818
28	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0929	0 / 0.0929	0 / 0.3863	0 / 0.3863
29	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0940	0 / 0.0940	0 / 0.3910	0 / 0.3910
30	240	1 / 0.0100	1 / 0.0100	1 / 0.0100	1 / 0.0100	0 / 0.0951	0 / 0.0951	0 / 0.3957	0 / 0.3957

Notes: The table reports p-values for the KPSS test (null: stationarity) and the augmented Dickey–Fuller (ADF) test (null: unit root) applied to Cleveland Fed market-based inflation expectations using only the last 20 years of monthly observations ($T = 240$). For KPSS, the long-run variance bandwidth is selected via AIC or BIC using an AR(p) approximation; entries report *bandwidth / p-value*. For ADF, the number of lagged differences is selected via AIC or BIC based on the ADF regression; entries report *lag / p-value*. "Constant" corresponds to a specification with an intercept only; "Trend" allows for a linear time trend.

Table 3: Scalar stationarity and unit-root tests at short horizons (Aruoba inflation expectations, Jan. 1998–Oct. 2025)

Horizon	T	KPSS (null: stationary)				ADF (null: unit root)			
		Constant		Trend		Constant (ARD)		Trend (TS)	
(months)		AIC: bw / p	BIC: bw / p	AIC: bw / p	BIC: bw / p	AIC: lag / p	BIC: lag / p	AIC: lag / p	BIC: lag / p
3	334	5 / 0.0100	2 / 0.0100	5 / 0.0100	2 / 0.0100	4 / 0.1174	1 / 0.0274	4 / 0.1750	1 / 0.0505
6	334	5 / 0.0100	5 / 0.0100	5 / 0.0100	5 / 0.0100	4 / 0.0972	4 / 0.0972	4 / 0.1700	4 / 0.1700
9	334	5 / 0.0100	5 / 0.0100	5 / 0.0100	5 / 0.0100	4 / 0.0799	4 / 0.0799	4 / 0.1651	4 / 0.1651
12	334	5 / 0.0102	4 / 0.0100	5 / 0.0100	4 / 0.0100	4 / 0.0660	3 / 0.0078	4 / 0.1609	3 / 0.0248
15	334	5 / 0.0168	4 / 0.0100	5 / 0.0100	4 / 0.0100	4 / 0.0546	3 / 0.0076	4 / 0.1559	3 / 0.0285
18	334	4 / 0.0109	4 / 0.0109	4 / 0.0100	4 / 0.0100	3 / 0.0071	3 / 0.0071	3 / 0.0307	3 / 0.0307
21	334	4 / 0.0118	4 / 0.0118	4 / 0.0100	4 / 0.0100	3 / 0.0065	3 / 0.0065	3 / 0.0318	3 / 0.0318
24	334	4 / 0.0100	4 / 0.0100	4 / 0.0100	4 / 0.0100	3 / 0.0059	3 / 0.0059	3 / 0.0320	3 / 0.0320
36	334	4 / 0.0100	4 / 0.0100	4 / 0.0100	4 / 0.0100	3 / 0.0049	3 / 0.0049	3 / 0.0283	3 / 0.0283
48	334	4 / 0.0100	4 / 0.0100	4 / 0.0100	4 / 0.0100	3 / 0.0053	3 / 0.0053	3 / 0.0206	3 / 0.0206
60	334	4 / 0.0100	4 / 0.0100	4 / 0.0100	4 / 0.0100	3 / 0.0067	3 / 0.0067	3 / 0.0128	3 / 0.0128
72	334	4 / 0.0100	3 / 0.0100	4 / 0.0100	3 / 0.0100	3 / 0.0083	2 / 0.0348	3 / 0.0077	2 / 0.0505
84	334	3 / 0.0100	3 / 0.0100	3 / 0.0100	3 / 0.0100	2 / 0.0331	2 / 0.0331	3 / 0.0046	2 / 0.0255
96	334	3 / 0.0100	3 / 0.0100	3 / 0.0100	3 / 0.0100	2 / 0.0327	2 / 0.0327	2 / 0.0139	2 / 0.0139
108	334	3 / 0.0100	3 / 0.0100	3 / 0.0100	3 / 0.0100	2 / 0.0329	2 / 0.0329	2 / 0.0083	2 / 0.0083
120	334	3 / 0.0100	3 / 0.0100	3 / 0.0100	3 / 0.0100	2 / 0.0334	2 / 0.0334	2 / 0.0050	2 / 0.0050

Notes: The table reports p-values for the KPSS test (null: stationarity) and the augmented Dickey–Fuller (ADF) test (null: unit root) applied to Cleveland Fed market-based inflation expectations at short horizons from 3 to 120 months (monthly sample, $T = 334$). For KPSS, the long-run variance bandwidth is selected via AIC or BIC using an AR(p) approximation; entries report *bandwidth / p-value*. For ADF, the number of lagged differences is selected via AIC or BIC; entries report *lags / p-value*. "Constant" denotes an intercept-only specification; "Trend" allows for a linear time trend. KPSS p-values are truncated at 0.0100 by the software.

C.5 Discussion

The scalar tests reported in this appendix are not intended to provide a definitive characterization of the number or nature of long-run components driving inflation expectations. Instead, they establish a robust empirical regularity: inflation expectations at individual maturities display persistence consistent with the presence of stochastic trends. This evidence motivates the functional time-series framework developed in the main text, which formally assesses common stochastic trends across the entire term structure and provides a basis for identifying permanent and transitory shocks to inflation expectations.

D Unit Root Tests for Functional Time Series

This appendix describes the unit root tests for functional time series used in the paper. Because functional data are infinite-dimensional objects, inference on persistence and stochastic trends necessarily relies on finite-dimensional representations or projections. We employ two complementary approaches.

The Johansen-based approach adapts standard multivariate cointegration methods to the functional setting. The functional time series is first projected onto a finite number of principal components, yielding a low-dimensional vector process. Cointegration tests are then applied to this projected system to infer the number of cointegrating relations, and by implication the dimension of the non-stationary space. This procedure is transparent and closely connected to familiar VAR-based methods. However, inference is conditional on the chosen finite-dimensional approximation, and the results may depend on the number of components retained and on the adequacy of the VAR representation for the projected scores.

By contrast, the approach developed by Chang et al. (2016) directly tests the dimension of the unit-root subspace in the underlying functional process. Rather than inferring non-stationarity indirectly through cointegration among projected components, this method treats the functional object itself as the primitive and assesses how many stochastic trends are present at the level of the entire curve. The resulting test statistics have nuisance-parameter-free limiting distributions and do not rely on specifying or estimating a finite-dimensional VAR. As such, they provide a complementary and robust check on the presence and dimensionality of persistent components, independent of the particular finite-dimensional approximation used for estimation.

D.1 Finite-Dimensional Representation

Let $\{f_t\}_{t=1}^{n}$ denote a sequence of functional observations, observed on a common discretized grid and represented as vectors $f_t \in \mathbb{R}^d$. Define the sample mean

$$\bar{f} = \frac{1}{n}\sum_{t=1}^{n} f_t,$$

and the demeaned process $\tilde{f}_t = f_t - \bar{f}$. Let

$$\widehat{Q} = \frac{1}{n} \sum_{t=1}^{n} \tilde{f}_t \tilde{f}_t'$$

denote the empirical covariance operator. Let $\{v_1, \ldots, v_m\}$ be the eigenvectors associated with the m largest eigenvalues of $\widehat{Q}$, and define the projection operator

$$\Pi_m = \sum_{j=1}^{m} v_j v_j'.$$

The finite-dimensional score process is defined as

$$[f]_t = \Pi_m f_t \in \mathbb{R}^m,$$

which represents the functional process projected onto the leading m principal components.

D.2 Error-Correction Representation

We model the projected process $\{[f]_t\}$ using a VAR(p) in error-correction form,

$$\Delta[f]_t = \Gamma[f]_{t-1} + \sum_{i=1}^{p-1} \Phi_i \Delta[f]_{t-i} + u_t.$$

To isolate long-run dynamics, all variables entering the cointegration analysis are first residualized with respect to the lagged differences $\{\Delta[f]_{t-1}, \ldots, \Delta[f]_{t-(p-1)}\}$.

D.3 Sequential Testing of Unit Roots

Let $\ell_0 \in \{m, m-1, \ldots, 1\}$ denote the number of unit roots under the null hypothesis. Partition the score vector as

$$[f]_t = \begin{pmatrix} [f]_t^N \\ [f]_t^S \end{pmatrix}, \qquad \dim([f]_t^N) = \ell_0, \quad \dim([f]_t^S) = m - \ell_0.$$

Let U_t^N denote the residuals from regressing $\Delta[f]_t^N$ on lagged differences, U_t^S the residuals from regressing $\Delta[f]_t^S$ on lagged differences, and U_t^1 the residuals from regressing $[f]_{t-1}^S$ on lagged differences.

The relevant blocks of the error-correction matrix are estimated as

$$\widehat{\Gamma}_{NS} = \left(\sum_t U_t^N U_t^{1\prime} \right) \left(\sum_t U_t^1 U_t^{1\prime} \right)^{-1}, \qquad \widehat{\Gamma}_{SS} = \left(\sum_t U_t^S U_t^{1\prime} \right) \left(\sum_t U_t^1 U_t^{1\prime} \right)^{-1}.$$

D.4 Permanent–Transitory Decomposition

Following Beveridge and Nelson (1981) and Gonzalo and Granger (1995), we construct a candidate permanent component as

$$[f]_t^P = [f]_t^N - [f]_t^S \widehat{A}', \qquad \widehat{A} = \widehat{\Gamma}_{SS}^{-1} \widehat{\Gamma}_{NS}.$$

By construction, $[f]_t^P$ eliminates error-correction feedback and captures the common stochastic trends implied by the null hypothesis of ℓ_0 unit roots.

D.5 Johansen Max-Eigenvalue Test

We apply a Johansen maximum-eigenvalue test to the extracted permanent component $\{[f]_t^P\}$. Define the moment matrices

$$S_{00} = \sum_t \Delta[f]_t^P \Delta[f]_t^{P\prime}, \qquad S_{11} = \sum_t [f]_{t-1}^P [f]_{t-1}^{P\prime},$$

$$S_{10} = \sum_t [f]_{t-1}^P \Delta[f]_t^{P\prime}, \qquad S_{01} = S_{10}'.$$

Let $\hat{\lambda}_{\max}$ denote the largest generalized eigenvalue solving

$$\det\bigl(S_{10} S_{00}^{-1} S_{01} - \lambda S_{11}\bigr) = 0.$$

The test statistic is

$$\text{LR}_{\max} = -(n-1)\log(1 - \hat{\lambda}_{\max}),$$

which is compared to standard Johansen critical values. A sequential testing procedure is used, stopping at the first non-rejection to obtain an estimate $\hat{\ell}$ of the number of unit roots in the functional process.

D.6 Test for the Dimension of the Unit Root Subspace Chang et al. (2016)

As a complementary approach, we employ the unit root dimension test proposed by Chang et al. (2016), which directly targets the dimension of the unit root subspace of a functional time series and yields nuisance-parameter-free asymptotic distributions.

D.6.1 Hypotheses and Testing Strategy

Let H_N denote the unit root subspace of the functional process $\{f_t\}$. The objective is to determine its dimension $M = \dim(H_N)$. For each candidate value $M \in \{0, 1, \ldots, M_{\max}\}$, we test

$$H_0 : \dim(H_N) = M,$$

against the alternative

$$H_1 : \dim(H_N) \leq M - 1.$$

Rejection of H_0 at level $\alpha = 0.05$ implies that the data support fewer than M unit-root directions. The test is implemented sequentially, starting from the upper bound $M_{\max} = 5$ and proceeding downward. As long as the null is rejected, we decrease M by one. The estimate of the unit-root dimension, $\widehat{M}$, is defined as the largest value of M for which the null hypothesis is not rejected:

$$\widehat{M} = \max\{\, M \leq M_{\max} : \text{do not reject } H_0(M)\,\}.$$

Equivalently, $\widehat{M}$ is the first value encountered (moving downward from $M_{\max}$) at which the test fails to reject.

D.6.2 Construction of the Test Statistic

For a given M, let $\{v_1, \ldots, v_M\}$ be any set of vectors spanning the unit root subspace H_N; the choice of basis is immaterial. Define the projected process

$$z_t = \big(\langle v_1, f_t\rangle, \ldots, \langle v_M, f_t\rangle\big)', \qquad t = 1, \ldots, T.$$

Let $Z_T = (z_1, \ldots, z_T)'$ and define

$$Q_M^T = Z_T' Z_T.$$

Under the null hypothesis, $T^{-2} Q_M^T$ converges weakly to an integrated Brownian motion covariance matrix.

To account for serial correlation, the test is based on generalized eigenvalues of Q_M^T with respect to a consistent estimator of the long-run variance matrix

$$\Omega_M = \sum_{|i| \leq \ell} \varpi_\ell(i) \Gamma_T(i),$$

where $\Gamma_T(i)$ denotes the sample autocovariance of Δz_t, $\varpi_\ell(i)$ is a bounded kernel, and the truncation lag ℓ increases with T.

Let

$$\lambda_1(Q_M^T, \Omega_M^T) \leq \cdots \leq \lambda_M(Q_M^T, \Omega_M^T)$$

denote the generalized eigenvalues of Q_M^T with respect to Ω_M^T.

D.6.3 Test Statistic and Asymptotic Distribution

The test statistic for $H_0 : \dim(H_N) = M$ is

$$\tau_M^T = T^{-2} \lambda_M(Q_M^T, \Omega_M^T),$$

i.e., the smallest generalized eigenvalue.

Under suitable regularity conditions,

$$\tau_M^T \xrightarrow{d} \lambda_M(Q_M^*),$$

where $Q_M^* = \int_0^1 W_M^*(r)W_M^*(r)'dr$ and W_M^* is an M-dimensional standard Brownian motion. The limiting distribution is free of nuisance parameters and depends only on M.

Under the alternative hypothesis, $\tau_M^T \xrightarrow{p} 0$, implying that the test is consistent.

D.6.4 Implementation and Interpretation

In practice, the vectors spanning H_N are replaced by consistent estimators obtained from the leading eigenfunctions of the empirical covariance operator. The test statistic is invariant to the particular choice of basis for H_N.

Critical values for τ_M^T depend only on M and are tabulated in Chang et al. (2016). This test provides a robustness check for the finite-dimensional, Johansen-based approach by delivering inference directly on the unit root structure of the functional process.

E Choice of the number of components m and VAR lag length p

This appendix documents how we choose (i) the number of functional principal components retained, m, and (ii) the lag length, p, of the unrestricted VAR used to forecast the component scores. The goal is to keep the main text focused on the economic interpretation, while making the empirical choice transparent and reproducible.

E.1 Forecast loss and rolling-window summaries

Let $\{\tau_j\}_{j=1}^J$ denote the maturity grid. For each rolling estimation window $r = 1, \ldots, 291$, forecast horizon $h \in \{1, \ldots, 24\}$ months, component dimension m, and VAR lag length p, we compute the maturity-aggregated squared forecast error

$$\text{SFE}_r(h; m, p) = \sum_{j=1}^{J} \left(\widehat{f}_{r+h|r}^{(m,p)}(\tau_j) - f_{r+h}(\tau_j) \right)^2, \qquad (2)$$

where $f_t(\tau)$ denotes the observed term structure object (e.g., expected inflation at maturity τ), and $\widehat{f}_{r+h|r}^{(m,p)}(\tau)$ is the corresponding forecast implied by an m-dimensional score process governed by an unrestricted VAR(p).

Rather than reporting means (which can be sensitive to occasional large forecast misses), we summarise the distribution of $\text{SFE}_r(h; m, p)$ across rolling windows using quantiles. For each (h, m, p) we record

$$Q_q(h; m, p) \equiv q\text{-th quantile of } \{\text{SFE}_r(h; m, p)\}_{r=1}^{291}, \qquad q \in \{0.025, 0.25, 0.50, 0.75, 0.975\}.$$

The spreadsheet accompanying the paper contains these quantiles for horizons $h = 1, \ldots, 24$, component choices $m \in \{1, \ldots, 6\}$, and lag lengths $p \in \{1, \ldots, 12\}$.

E.2 Why we impose $m \geq 3$

A purely predictive criterion can favour very low-dimensional representations (e.g. $m = 1$ or $m = 2$) because they are parsimonious. However, the empirical objective of the paper is not only forecasting accuracy but also an economically interpretable decomposition of movements in the term structure into:

- a *non-stationary direction* (a persistent "level/trend" component capturing slow-moving belief revisions), and
- a *stationary subspace* with *at least two* dynamic degrees of freedom (e.g. slope/curvature-type movements).

This decomposition requires at least three dimensions: one to accommodate the non-stationary direction and two to allow richer stationary dynamics. For that reason, we treat $m = 3$ as the smallest specification that delivers the conceptual split used throughout the analysis, and we evaluate whether $m > 3$ yields meaningful additional benefits.

E.3 A joint selection rule for (m, p)

To discipline the choice of (m, p) while respecting the interpretability constraint above, we use two principles:

(i) Integrated typical performance. Define the horizon-averaged median loss

$$\overline{Q}_{0.50}(m, p) \equiv \frac{1}{24} \sum_{h=1}^{24} Q_{0.50}(h; m, p). \qquad (3)$$

This measures typical performance across horizons.

(ii) Tail-risk/stability. We also track

$$\overline{Q}_{0.75}(m, p) = \frac{1}{24} \sum_{h=1}^{24} Q_{0.75}(h; m, p), \qquad \overline{Q}_{0.975}(m, p) = \frac{1}{24} \sum_{h=1}^{24} Q_{0.975}(h; m, p),$$

to ensure that a specification does not achieve a slightly better median at the cost of very poor performance in a non-trivial fraction of rolling windows.

Finally, we apply *parsimony* as a tie-breaker: increasing m and p quickly increases the number of VAR coefficients (approximately $m^2 p$ slope parameters, plus intercepts), which can amplify estimation noise in rolling samples and lead to unstable long-horizon dynamics.

E.4 Empirical choice and robustness

Lag length. Conditioning on $m = 3$, the specification $p = 1$ minimises $\overline{Q}_{0.50}(3, p)$ and also yields the smallest (or essentially smallest) $\overline{Q}_{0.75}(3, p)$. Horizon-by-horizon, the best-performing p under $m = 3$ is overwhelmingly $p = 1$, with only negligible improvements from $p > 1$ at a small number of horizons. Given the large increase in parameters implied by longer lag lengths, we select $p = 1$.

Number of components. Relative to $m = 3$, increasing the number of components beyond three does not deliver systematic improvements in typical performance, while it can worsen tail outcomes in some cases (a symptom of overfitting/instability in certain rolling windows). Thus, once we enforce the economically motivated requirement $m \geq 3$, $m = 3$ is the natural baseline.

Summary evidence. Table 4 reports horizon-averaged quantile summaries for a small set of representative specifications. The baseline $(m, p) = (3, 1)$ is very close to the best-performing model in terms of median loss, while remaining parsimonious and stable in the tails. By contrast, larger specifications may exhibit undesirable tail behaviour (large $\overline{Q}_{0.975}$), despite only modest differences in the median.

Table 4: Horizon-averaged rolling-window SFE quantiles for selected (m, p) choices

Model	m	p	$\overline{Q}_{0.50}$	$\overline{Q}_{0.75}$	$\overline{Q}_{0.975}$
Overall best by $\overline{Q}_{0.50}$ (unconstrained)	1	6	0.0053	0.0156	0.0572
Baseline (interpretable)	3	1	0.0056	0.0164	0.0571
Closest low-dimensional alternative	2	1	0.0054	0.0162	0.0584
Next dimension (illustrative)	4	1	0.0069	0.0167	0.0655
High-dimensional example (illustrative)	5	5	0.0065	0.0199	2.39×10^6

E.5 Joint selection rule for (m, p)

To choose (m, p) jointly, we adopt a robust summary of predictive performance that avoids undue sensitivity to a few extreme windows. Let $Q_{0.50}(h, m, p)$ denote the median (across rolling windows) of the SFE distribution at horizon h for specification (m, p). We then aggregate across horizons using

$$\overline{Q}_{0.50}(m, p) \;\equiv\; \frac{1}{H} \sum_{h=1}^{H} Q_{0.50}(h, m, p), \qquad H = 24,$$

and select the pair (m, p) that minimises $\overline{Q}_{0.50}(m, p)$ subject to $m \geq 3$.

Table 5 reports $\overline{Q}_{0.50}(m, p)$ for $m \geq 3$ and $p \in \{1, \ldots, 12\}$. Two patterns stand out: (i) for each $m \geq 3$, shorter lag orders dominate in terms of average median SFE; (ii) within the admissible set $m \geq 3$, the overall minimum occurs at $(m, p) = (3, 1)$ (bold in the table). Moreover, for $m = 3$, $p = 1$ is the horizon-by-horizon minimiser for the vast majority of horizons (in our grid, it is best for 21 of the 24 horizons), reinforcing the choice of a parsimonious dynamics specification.

Accordingly, the baseline model in the main text fixes $m = 3$ and $p = 1$. All remaining quantiles (25th, 75th, and tail percentiles) and alternative choices (e.g., $p = 2$ or $m = 4$) are reported in this Appendix to document robustness without burdening the main exposition.

$m \backslash p$	1	2	3	4	5	6	7	8	9	10	11	12
3	0.005574	0.006129	0.006411	0.006485	0.006506	0.006463	0.006684	0.006847	0.007072	0.007110	0.007101	0.007122
4	0.006420	0.007110	0.007486	0.007556	0.007618	0.007640	0.007706	0.007848	0.008157	0.008414	0.008439	0.008437
5	0.007323	0.008211	0.008938	0.009167	0.009234	0.009243	0.009272	0.009286	0.009298	0.009346	0.009364	0.009375
6	0.008753	0.010017	0.010864	0.010922	0.010969	0.011171	0.011245	0.011281	0.011287	0.011294	0.011313	0.011312

Table 5: Average (across horizons) of the rolling-window median sum of squared forecast errors, $\overline{Q}_{0.50}(m,p) \equiv \frac{1}{H}\sum_{h=1}^{H} Q_{0.50}(h,m,p)$ with $H = 24$. Entries are reported for $m \geq 3$ and $p \in \{1,\ldots,12\}$; the bold entry denotes the minimum in this grid.