The Effects of Economic Shocks on the Distribution of U.S. Household Inflation Expectations *

Yoosoon Chang† Fabio Gómez-Rodríguez‡ Gee Hee Hong§

This version: February 15, 2026

Abstract

Using a functional time-series framework, we study how the full distribution of U.S. household inflation expectations—rather than only its average—responds to economic shocks. Contractionary monetary policy shocks lower average inflation expectations and reduce the share of households expecting high inflation. Personal income tax and gasoline price shocks shift the distribution upward, with gasoline price shocks also increasing disagreement in medium-run expectations. In contrast, government spending shocks have little effect on the distribution. We further show that responses depend on households' pre-shock beliefs. Taken together, these findings show that focusing on average expectations misses economically meaningful shifts in dispersion and tail risk.

JEL classification: E52, E61, E62, E63, H11, H30

Keywords: inflation expectations, household survey, functional autoregression, functional impulse responses, economic shocks, policy transmission

*We would like to thank Chaewon Baek, Olivier Coibion, Jaewon Lee, Jaewoo Lee, Emi Nakamura, Joon Y. Park, Woongyong Park, Jae Sim, Michael Weber, Deniz Igan, seminar participants at the International Monetary Fund, Norges Bank, Federal Reserve Bank of Chicago, BI Norwegian Business School, Bank of International Settlements, the KAEA seminar series, the conference of the Society of Economic Measurement and anonymous referees. We are also grateful to Sangmyung Ha for excellent research assistance. All remaining errors are our own.

†Indiana University, Bloomington (yoosoon@indiana.edu)

‡Lehigh University (fabio.gomez-rodriguez@lehigh.edu), Central Bank of Costa Rica

§International Monetary Fund (ghong@imf.org)

1 Introduction

Since the seminal contributions of Friedman (1968) and Phelps (1967), inflation expectations have stood at the core of macroeconomic analysis and monetary policy. A large theoretical and empirical literature highlights their central role in shaping consumption, investment, and wage-setting decisions. Yet, most empirical studies summarize survey-based inflation expectations using a small number of statistics—typically the mean or median—implicitly treating the average belief as representative of the population. Recent research, particularly following the sharp post-pandemic inflation surge, has begun to document the importance of higher moments and tail behavior of inflation expectations across countries (Reis, 2022). A growing body of evidence further underscores substantial heterogeneity in household inflation expectations—driven by differences in demographics, income, and consumption behavior—suggesting that the cross-sectional distribution of expectations contains macroeconomically relevant information beyond its average (D'Acunto et al., 2023; Andre et al., 2022; Doh et al., 2025).

This paper builds on these insights by studying inflation expectations explicitly as a *distributional object*. We ask how different macroeconomic shocks reshape not only the average but the entire cross-sectional distribution of household inflation expectations. Rather than focusing on a few pre-selected summary moments, we analyze the entire cross-sectional distribution of U.S. household inflation expectations using a functional time-series framework. The motivation for this approach is empirical and general: when the shape of the distribution changes independently of its center, conventional moment-based summaries may fail to capture economically meaningful variation. While this issue is present throughout the sample, the post-pandemic period provides a particularly salient illustration.

Figure 1 plots the evolution of one-year-ahead and medium-run inflation expectations of U.S. households from March 2018 to December 2021, using data from the *Survey of Consumers*.[1] Between the pre-pandemic period (blue) and the early months of the pandemic (red), the cross-sectional distributions flatten substantially and the concentration around the mode declines. Despite this pronounced change in shape, the mean one-year-ahead expectation increases only from 2.7 to 2.8 percent, and the medium-run mean from 2.4 to 2.5 percent. These patterns highlight that distributional shifts can be far more informative than movements in averages alone, motivating an empirical framework that captures the full cross-sectional dynamics of inflation expectations.

Figure 1. *Distribution of the U.S. Households' One-Year Ahead and Medium-Run Inflation Expectations: March 2018–December 2021*

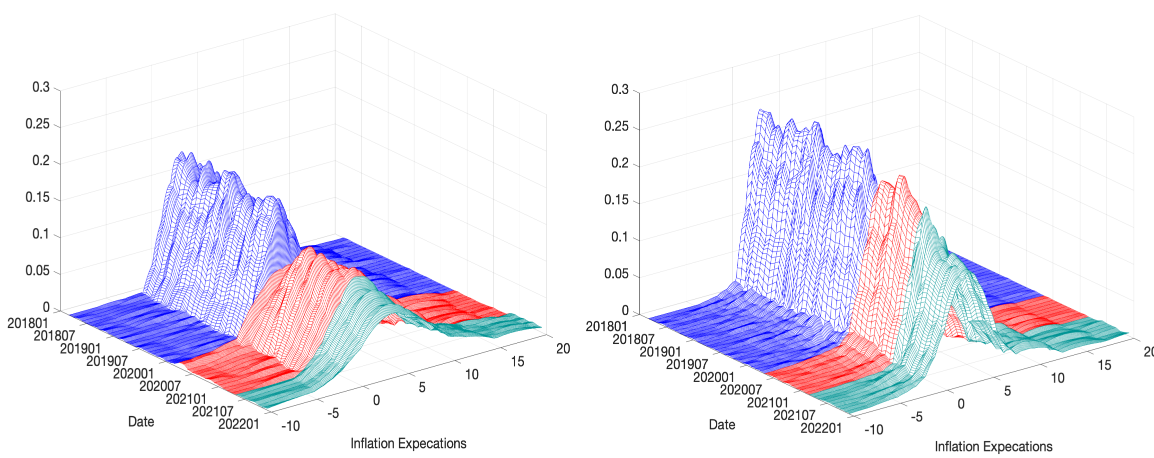

Notes: The distributions of monthly U.S. household inflation expectations are shown for three periods: March 2018–February 2020 (blue), March 2020–March 2021 (red), and April 2021–December 2021 (green). The left (right) panel plots one-year-ahead (medium-run) inflation expectations. Both series are constructed using data from the University of Michigan's *Survey of Consumers*.

Using inflation expectations from the *Survey of Consumers*, we adopt the functional autoregressive framework of Chang et al. (2021) to study how four externally identified

[1]One-year-ahead (short-run) expectations are based on the question: "By about what percent do you expect prices to go (up/down) on average during the next 12 months?" Medium-run expectations refer to the "next 5 to 10 years." Each respondent reports a single percentage, interpreted as an expected annual inflation rate. The 5–10-year series, available monthly, is commonly used to represent medium- to long-term expectations (e.g., Coibion and Gorodnichenko, 2015; Carroll, 2003).

economic shocks—monetary policy, government spending, personal income tax, and gasoline prices—affect the distribution of household inflation expectations.[2] These shock series are drawn from different sources and may contain correlated information. To isolate distinct surprise components, we purge each shock of its linear dependence on the others by regressing it on the remaining shocks and retaining the residual as a purified innovation. This procedure ensures that we attribute distributional responses to distinct sources of economic news rather than to overlapping information.

The key object in our analysis is the *expected inflation distribution* (EID), defined as a time series of estimated density functions, each representing the full cross-sectional distribution of household inflation expectations in a given month.[3] For each horizon—one-year-ahead and medium-run—we apply functional principal component analysis (FPCA) to obtain a low-dimensional representation that captures the dominant modes of variation in the EIDs. These components are then embedded in a vector autoregression and linked to the externally identified shocks to construct *functional impulse responses* that trace how the entire distribution reacts to each shock over time. Modeling expectations as a functional object allows any moment—mean, median, dispersion, or skewness—to be recovered as a by-product, while avoiding the loss of information inherent in moment-restricted approaches.

Our main findings can be summarized as follows. Contractionary monetary policy shocks, measured by Miranda-Agrippino and Ricco (2021), lower both short- and medium-run inflation expectations. A one-percentage-point policy tightening reduces average one-year-ahead and medium-run expectations by 0.10 and 0.04 percentage points, respectively, primarily by reducing the share of households expecting inflation above 5 percent and increasing the share expecting moderate inflation (between -0.6 and 2.4 percent). The

[2]We use the *Survey of Consumers* because its large monthly cross-section enables precise estimation of continuous distributions, unlike smaller-sample surveys such as the *Survey of Professional Forecasters*.

[3]Throughout the paper, the "D" in EID refers interchangeably to "density" or "distribution."

dispersion of expectations declines significantly, indicating stronger anchoring, while skewness remains largely unchanged. Fiscal shocks exhibit asymmetric effects: government spending shocks have no measurable impact on the distribution at any horizon, whereas personal income tax increases shift the medium-run distribution to the right, raising both dispersion and skewness. Finally, consistent with Coibion and Gorodnichenko (2015), gasoline price shocks raise inflation expectations at both horizons. We uncover a novel distributional pattern: while short-run expectations shift uniformly upward, medium-run expectations become more dispersed and asymmetric, suggesting heterogeneity in how households extrapolate energy price shocks into future inflation.

This paper contributes to three strands of the literature. First, it advances the empirical study of inflation expectations by moving beyond moment-based summaries to analyze the entire cross-sectional distribution. Second, it contributes to the literature on the effects of monetary and fiscal shocks on expectations, complementing existing evidence based on average beliefs (Coibion et al., 2020; Grigoli et al., 2020; Coibion et al., 2021). Third, from a methodological perspective, the paper demonstrates how functional autoregressive methods can be applied in a transparent and variance-efficient way to macroeconomic expectation data.[4] While most existing studies focus on average expectations, our results show that important distributional responses—particularly in dispersion and tail behavior—would be missed under such approaches. In this sense, our findings complement Doh et al. (2025) and Guillochon (2024), while documenting a novel response of medium-run inflation expectations to personal income tax shocks.

The remainder of the paper is organized as follows. Section 2 motivates our focus on the full distribution of inflation expectations. Section 3 describes the econometric framework.

[4] For related applications of functional data methods in economics, see Chang et al. (2016), Inoue and Rossi (2019), Bjørnland et al. (2025), Chang et al. (2026) and Chang et al. (2025).

Section 4 details the empirical implementation using data from the *Survey of Consumers*. Section 5 presents the distributional effects of economic shocks, and Section 6 concludes.[5]

2 Why a Functional Approach to the Distribution of Inflation Expectations (EIDs)?

This section motivates our choice of a *functional approach* to analyzing the *entire distribution* of inflation expectations. Rather than appealing to conceptual arguments alone, we provide quantitative evidence on how much information is lost when the distribution is summarized using commonly employed statistics such as the mean or median. We do so through two complementary exercises. First, we examine which distributional features of inflation expectations contain predictive information for key macroeconomic variables using an adaptive LASSO procedure. Second, we assess the trade-off between explanatory power and estimation precision across alternative distributional representations, highlighting the efficiency gains of a functional principal component approach.

We begin by using an adaptive LASSO procedure to identify which distributional properties of one-year-ahead inflation expectations are informative for macroeconomic outcomes[6]. Importantly, this exercise is conducted entirely using distributional statistics computed directly from the raw cross-sectional survey responses in each month, prior to any density estimation or functional modeling. The candidate regressors include a broad set of distri-

[5]The appendices provide supporting material: Appendix A presents the mathematical foundations; Appendix B discusses alternative bases; Appendix C interprets the functional components; Appendix D reports robustness checks; Appendix E details the bootstrap procedures; and Appendix F discusses alternative functional approaches.

[6]The adaptive LASSO (Zou, 2006) is a regularized regression method that performs variable selection and coefficient shrinkage in a data-driven way. It extends the standard LASSO by introducing adaptive weights on the penalty term, improving consistency in identifying relevant predictors. Because our motivating regression includes a large set of distribution-based summary statistics, adaptive LASSO provides an efficient way to isolate informative features while mitigating overfitting and multicollinearity.

butional statistics: the mean, median, standard deviation, skewness, kurtosis, deciles, and specific distributional shares, such as the proportion of households expecting deflation (inflation below 0 percent), the share with anchored expectations (between 1 and 3 percent), and the share expecting high inflation (above 6 or 10 percent). The dependent variables are current and future inflation, the unemployment rate, and real GDP growth.

Table 1. *Variables Selected with Specific Macroeconomic Variable Using LASSO*

	Inflation (π_t)	**Inflation (π_{t+12})**	**Unemployment (u_t)**	**Unemployment (u_{t+12})**	**Real GDP growth (y_t)**
Regressors					
Median		[+,***]			
Skewness	[+,***]	[+]			
Interquartile	[+,***]				
Share of inflation expectations below 0%	[-,***]	[+,]	[+,***]	[+,***]	[-,***]
Share of inflation expectations above 6%	[+,.]	[+,***]			
6th decile	[+,***]	[+,***]			
Share of inflation expectations between 1% and 3%		[+,***]	[-,***]	[-,***]	
Share of inflation expectations above 10%		[-,***]	[+,***]	[+,***]	[-,***]
2nd decile		[+,**]	[-,*]		
3rd decile		[+,***]	[-,***]		
R^2	0.8519	0.6124	0.3577	0.3735	0.1503

Notes: (i) Signs and significance levels are reported as [sign, significance], where '***', '**', '*', and '.' correspond to the 0.001, 0.01, 0.05, and 0.1 levels, respectively. (ii) Real GDP growth is measured at quarterly frequency and matched to the final month of the corresponding quarter in the one-year-ahead inflation expectations data.

Table 1 reports the variables selected by the adaptive LASSO and the sign and significance of the associated coefficients. Notably, commonly used summary statistics such as the mean and standard deviation are not selected as significant predictors. In contrast, higher-order moments and tail features consistently emerge as informative. The median is positively associated with future inflation, skewness is positively correlated with current inflation, and the share of households expecting inflation above 10 percent is positively related to future unemployment while being negatively related to contemporaneous real GDP growth. Similarly, the share of households expecting deflation is strongly correlated with both current inflation and unemployment.

These results indicate that the predictive content of inflation expectations extends well beyond measures of central tendency. Restricting attention to the mean or median obscures systematic variation in the tails and shape of the distribution that is closely linked to macroeconomic outcomes. This evidence motivates an empirical framework that treats the entire distribution of inflation expectations as the relevant object of analysis.

Second, we compare the explanatory power of different distributional representations while considering the trade-off between goodness of fit and estimation precision. This allows us to quantify the information content captured by the functional approach relative to conventional summary statistics. The results are reported in Tables 4 and 5 in Appendix B. In summary, our three functional principal components (FPCs) jointly explain over 95 percent of the total variation in the distribution of inflation expectations (see Section 4 for more details). In contrast, conventional summary statistics capture far less: the mean, standard deviation, and skewness together explain less than 2 percent of the variation, while even including five moments reaches only 4 percent. Using five selected percentiles (10th, 25th, 50th, 75th, 90th) performs substantially better, explaining approximately 72 percent of the variation, but still falls short of the functional approach.

Beyond the loss of information by focusing solely on pre-determined moments, there is an equally critical consideration, which relates to estimation precision. When attempting to capture more distributional information by including additional summary statistics or percentiles, estimation variances increase at an incredibly fast rate. For instance, a three-component representation using our FPCs has a variance measure of 4.7, while a comparable moment-based representation has a variance measure exceeding 3,900—more than 800 times larger. This combination—high explanatory power with manageable estimation variance—makes the functional principal component approach both efficient and practical for analyzing distributional dynamics, a conclusion that is consistent with evidence from other applications that document similar efficiency gains (e.g. Chang et al. (2024b) Bjørnland et al. (2025), Chang et al. (2026)).

Guided by these trade-offs, we approximate each month's expected inflation distribution (EID) using a small number of functional principal components (FPCs) and model the corresponding FPC loadings in a low-dimensional VAR. This approach allows us to trace how the entire distribution responds to externally identified economic shocks. Sections 3 and 4 describe the econometric framework and its implementation; Appendix A provides technical foundations, and Appendix F discusses alternative functional approaches.

3 Econometric Methodology

This section summarizes our econometric framework. We model monthly expected inflation distributions (EIDs) as a functional time series and approximate their dynamics using a low-dimensional VAR in a small number of functional principal component (FPC) loadings. This approach delivers impulse responses for the entire distribution while keeping estimation and inference within familiar VAR tools, without introducing new estimation or identification assumptions.

Our framework follows the functional autoregression (FAR) approach of Chang et al. (2021). Section 4 describes the empirical implementation (including how we link the EID dynamics to externally identified shocks). Technical foundations are collected in Appendix A.

Formally, let (f_t) denote the density representing the EID. We treat (f_t) as a time series of square-integrable functions and write the FAR(2) as

$$f_t = A_1 f_{t-1} + A_2 f_{t-2} + \varepsilon_t, \tag{1}$$

where A_1 and A_2 are linear operators on H and (ε_t) is a functional white noise in H (see Appendix A for precise definitions). We use the second-order FAR, as suggested by the commonly used information criteria, AIC and BIC, for our data.[7]

Following Chang et al. (2021), the infinite-dimensional FAR can be approximated by projecting onto a finite-dimensional subspace spanned by m basis functions $(v_i)_{i=1}^m$. This projection yields an m-dimensional VAR representation:

$$\lfloor f_t \rceil \approx \lfloor A_1 \rceil \lfloor f_{t-1} \rceil + \lfloor A_2 \rceil \lfloor f_{t-2} \rceil + \lfloor \varepsilon_t \rceil, \tag{2}$$

where $\lfloor f_t \rceil = (\langle v_1, f_t \rangle, \ldots, \langle v_m, f_t \rangle)'$ represents the functional observation as an m-dimensional vector of basis coefficients, and $\lfloor A_k \rceil$ denote the corresponding $m \times m$ coefficient matrices. The mapping between the functional and finite-dimensional representations is isometric, preserving the structure of the original FAR (see Appendix A for details).

The effectiveness of this approximation depends crucially on the choice of basis. Following Bosq (2000), Ramsay (2004), Hall and Horowitz (2007), and Park and Qian (2012), among others, we use the functional principal component (FPC) basis (v_i^*), defined as the eigenfunctions of the sample covariance operator

$$\Gamma = \frac{1}{T} \sum_{t=1}^{T} (f_t \otimes f_t). \tag{3}$$

As demonstrated in Section 2, three FPCs capture over 95 percent of the variation in EID while maintaining substantially lower estimation variance than alternative bases such as moments, percentiles, or intervals. Appendix B provides detailed comparisons of functional R-squared and integrated variances across different basis choices (Tables 4 and 5).

[7] A complementary approach to selecting the lag order relies on minimizing out-of-sample mean squared forecast errors (MSFE). In our case, both procedures yield consistent selections.

4 Implementation of Functional Approach using EIDs

In this section, we describe the empirical implementation of the methodology presented in Section 3, with each subsection written in the same order as the implementation.

4.1 Constructing EID Functional Data from Household Survey

Based on the monthly data from the University of Michigan's *Survey of Consumers*, we first estimate the density function representing the underlying distribution of inflation expectations. We use this time series of density functions as our functional data. Our EID sample is from January 1983 to December 2021 for one-year ahead inflation expectations. In the case of medium-run inflation expectations the sample begins on January 1991.

We focus on the following two questions related to households' inflation expectations: (1) "About what percent do you expect prices to go (up/down) on the average, *during the next twelve months*?"; and (2) "About what percent do you expect prices to go (up/down) on the average, *during the next five to ten years*?" The responses to the first question are households' one-year ahead inflation expectations, and those to the second question are households' medium-run inflation expectations. We estimate density functions using a standard kernel density estimation procedure.[8]

The estimated densities represent the underlying distribution of heterogeneous inflation expectations. In Figure 2, we present an example of the estimated density of one-year ahead expected inflation (red line) using the data for April 2011 and the frequency of actual responses (blue bars).

[8] Density functions are estimated using the Epanechnikov kernel function with the rule-of-thumb time-varying bandwidth (Park and Qian, 2012)

__Figure 2.__ Survey Responses and the Estimated Density: An Example

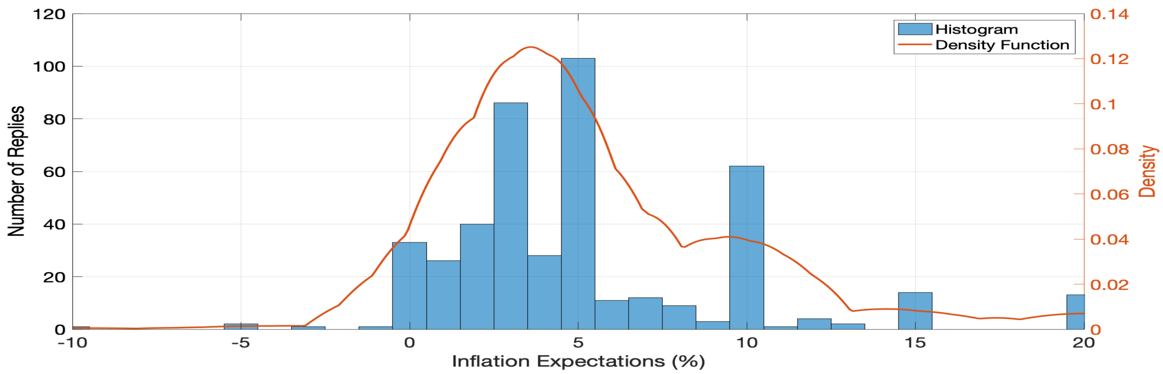

Notes: Survey responses of one-year ahead inflation expectations in April 2011 are used as an example. The left vertical axis represents the number of survey responses corresponding to different levels of inflation expectations reported in the actual survey. The right vertical axis shows the estimated density of inflation expectations.

4.2 Functional Principal Component Analysis

Using the constructed EID, we apply a functional principal component analysis (or FPCA) to obtain the basis of functions that optimally represents the variability of expected inflation distributions as a finite-dimensional process. FPCA generalizes PCA to functional spaces by identifying orthonormal eigenfunctions that capture the main modes, or most common patterns, of variability of the curves. The associated scores quantify how each curve varies along those modes. The eigenfunctions act as functional loadings linking variation across the entire domain.

First, the scree plot in Figure 3 displays how much of the total variation in the EIDs is captured as we increase the number of functional principal components in our basis. For one-year-ahead expectations, the first component alone explains over 74 percent of the variation. Including the second and third components raises the cumulative share to roughly 95 percent. A similar pattern holds for medium-run expectations, indicating that just three components sufficiently summarize the dynamics of both distributions. We, hence, retain the first three components ($m = 3$), which jointly capture about 95 percent of the variation in both EID series.[9]

Figure 3. *Scree Plots of One-Year Ahead and Medium-Run Inflation Expectations*

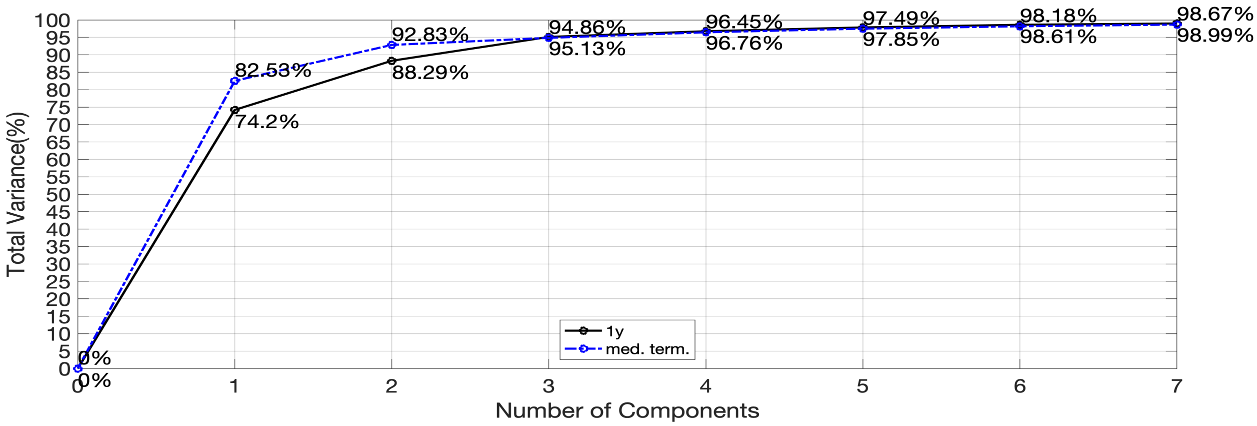

Notes: The horizontal axis shows the number of functional principal components. The vertical axis shows the cumulative proportion of functional variations explained, corresponding to the number of principal components. Black line (blue dotted line) show one-year ahead (medium-run) inflation expectations from January 1978 (from January 1991) to December 2021.

Figure 4 shows the loadings, functional components, and the shape of distribution affected by each of the three components for one-year-ahead expected inflation distributions (EIDs). Specifically, the left column shows the corresponding loadings $(\alpha_{kt}) = \langle v_k^*, f_t \rangle$ for each component. Blue and red stars mark, respectively, the minimum and maximum loadings over time, $\min(\alpha_{kt})$ and $\max(\alpha_{kt})$, for $k = 1, 2, 3$. Furthermore, the middle column presents the three estimated functional principal components (FPCs): the first (v_1^*), second (v_2^*), and third (v_3^*). These functions are mutually orthogonal and have unit norm. The right column depicts the range of distributions implied by each component. Starting from the average EID density,

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

[9]While including more components could increase explained variation, it would substantially inflate the variance of estimated functional coefficients, as documented in Tables 4 and 5 in Appendix B.

Figure 4. *Functional Principal Components and Loadings: one-year ahead Inflation Expectations*

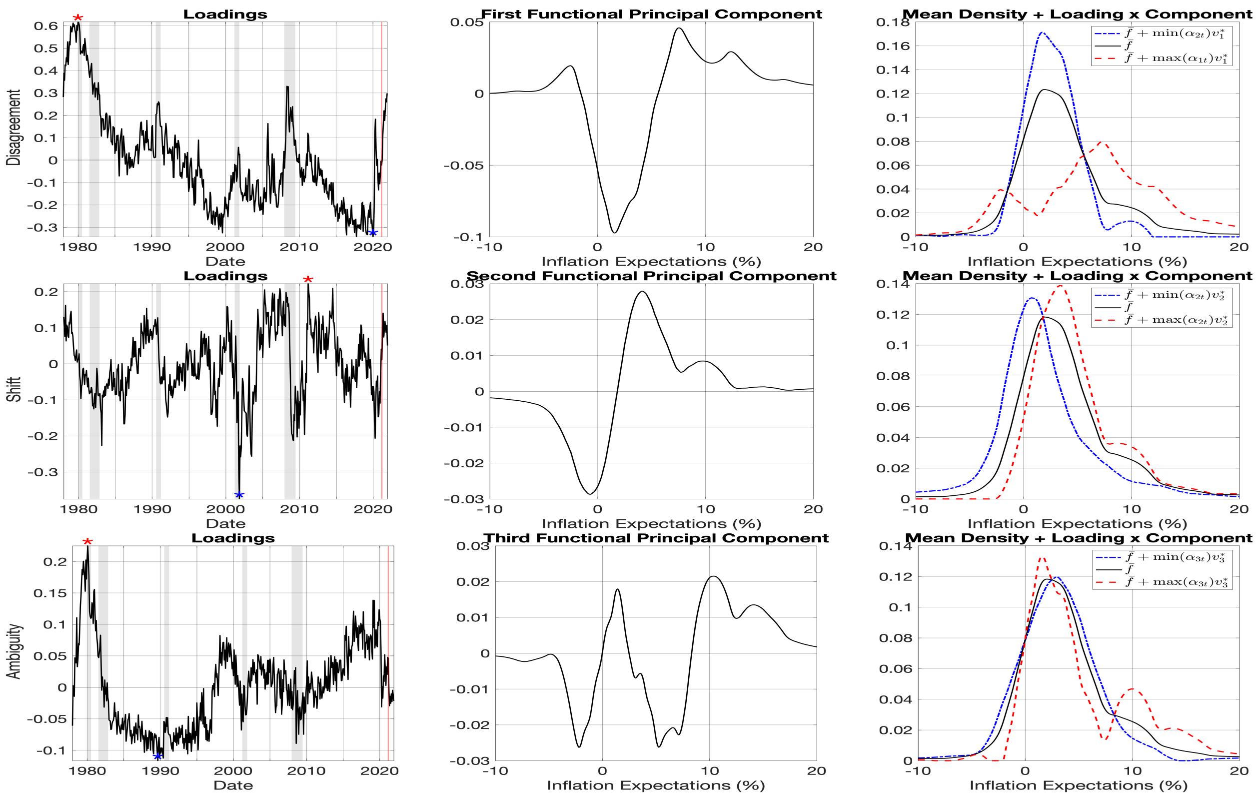

Notes: The left panels show loadings over time for the first three functional principal components (FPCs) for one-year ahead inflation expectations, the first FPC in the top, the second FPC in the middle, and the third FPC in the bottom rows). The center panels show each of the three FPCs as a function. The right panel shows how the sample mean density function (black solid line) changes with maximum (red dashed line) and minimum (blue dotted line) contributions by the maximum (red star) and minimum (blue star) values of the loadings in the left panel for corresponding FPCs.

we add each FPC scaled by its minimum and maximum loading. The black solid line shows $\bar{f}$, the blue dotted line represents $\bar{f} + \min(\alpha_{kt})v_k^*$, and the red dashed line corresponds to $\bar{f} + \max(\alpha_{kt})v_k^*$. Thus, the red (blue) line illustrates the case where the EID is most positively (negatively) affected by the k-th FPC.[10]

As explained above, the first functional principal component explains more than 74 percent of the variation in one-year ahead EIDs. We refer to this component as the *disagreement component* because it shifts probability mass from the center of the distribution

[10] Appendix C provides detailed analyses of how the moments of EIDs vary with the loadings of each functional component, supporting our interpretation and labeling of the three FPCs.

(between -1 and 5 percent) to both negative and very high inflation expectations (below -1 percent and above 4 percent). As illustrated in Figure 4, higher loadings widen the distribution (red dashed line), while lower loadings compress it toward the center (blue dotted line). The loading series exhibits a cyclical pattern, with peaks during major downturns—the early 1980s recession, the Global Financial Crisis, and the COVID-19 period—consistent with disagreement rising in turbulent times. The second component, explaining about 15 percent of the variation, represents a *level-shifting* movement of the entire distribution. Positive loadings shift the density to the right, indicating higher expected inflation across the cross-section, whereas negative loadings shift the density to the left. Finally, the third component captures *ambiguity in tail behavior*. Positive loadings produce fatter right tails and develop multiple bumps at elevated expectation levels (above 8 percent), while negative loadings smooth the upper tail. This component therefore reflects uncertainty concentrated in extreme inflation outcomes—particularly around the Federal Reserve's 2 percent target and in the higher tail of the distribution.

4.3 Measuring Functional EID Responses to External Shocks

As a next step, by applying a recursive normalization to the corresponding approximate vector autoregression (VAR), we obtain *functional shocks* driving the expected inflation distributions (EIDs) and their impacts on expected inflation distributions. The recursive scheme used here serves purely as a normalization.

The error term (ε_t) in the approximate second-order VAR in (2) refers to $\varepsilon_t = (\varepsilon_{1t}, \varepsilon_{2t}, \varepsilon_{3t})'$, where (ε_{1t}), (ε_{2t}), and (ε_{3t}) are innovations in our approximate VAR. These correspond to the loadings of the three leading FPCs, which we label as the *disagreement*, *shifting*, and *ambiguity* components, respectively. Let

$$\varepsilon_t = Be_t, \tag{4}$$

where B is a 3×3 matrix and $e_t = (e_{1t}, e_{2t}, e_{3t})'$ is a vector of functional shocks defined below.

We identify three functional shocks $(e_{1t}), (e_{2t}), (e_{3t})$ by imposing a recursive structure among them, again purely for convenience. The matrix B in (4) is defined as a unique lower triangular matrix satisfying

$$BB' = \Sigma,$$

where Σ is the estimated covariance matrix of ε_t. Each column β_i of B represents the at-impact response of the three-dimensional vector of FPC scores to the i-th functional shock e_{it}. We define

$$b_i = \pi^{-1}(\beta_i), \tag{5}$$

as the at-impact functional response of f_t to e_{it}. Thus, b_i is the *baseline EID response* to the i-th functional shock, for $i = 1, 2, 3$.

Figure 5 presents these baseline responses, computed for different sample periods to assess stability across time. Recall that β_i are columns of the lower triangular matrix B, implying that b_3 depends solely on the third FPC (v_3^*), b_2 on a linear combination of (v_2^*, v_3^*), and b_1 on (v_1^*, v_2^*, v_3^*). The shapes of b_1, b_2, b_3 thus closely resemble those of the corresponding FPCs, reflecting the dominance of leading components in our empirical model. We refer to these as the EID responses in *disagreement*, *shifting*, and *ambiguity*, respectively.

It is important to clarify that the retrieved functional shocks $(e_{1t}), (e_{2t}), (e_{3t})$ are not intended to be interpreted in structural economic terms. Instead, they provide an or-

Figure 5. At-impact Baseline EID Responses to Functional Shocks Using Different Samples

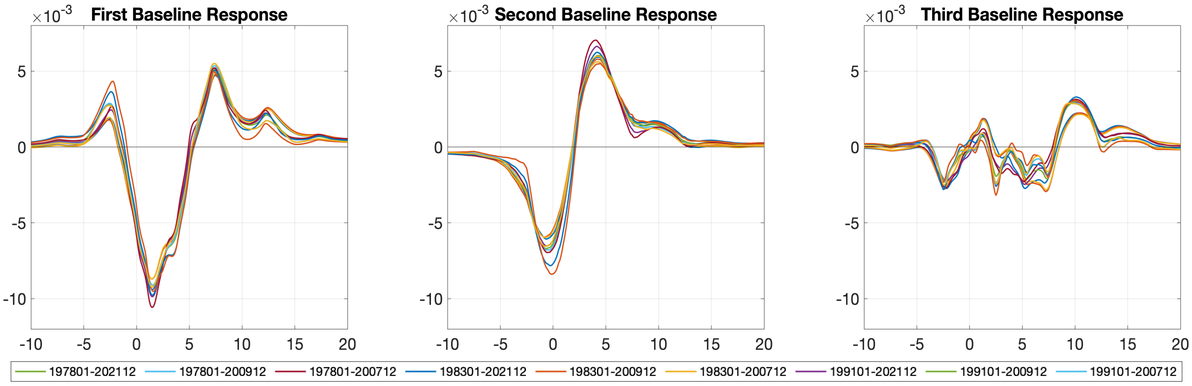

Notes: The figure shows the at-impact baseline impulse responses of one-year ahead EIDs to the three functional shocks using different sample periods, including the baseline sample from January 1983 to December 2021 for Figure 4.

thogonal basis that efficiently summarizes the dominant sources of variation in the EID dynamics. To study economically meaningful transmission channels, we therefore link these functional shocks to externally identified structural shocks.

Following Chang et al. (2025) and Bjørnland et al. (2025), we treat monetary policy, government spending, personal income tax, and gasoline price shocks as *external instruments*. Each of these shock series is standardized to have unit variance and is projected onto the space spanned by the orthogonal functional shocks.

Formally, let x_t denote an external shock. We compute the correlation vector

$$\rho_x = (\rho_{x1}, \rho_{x2}, \rho_{x3})', \qquad \rho_{xi} = \text{corr}(e_{it}, x_t),$$

which is equivalent to the vector of regression coefficients because both x_t and (e_{it}) have been normalized. Since the functional shocks are orthogonal, the magnitude $\|\rho_x\|$ measures how strongly the external shock loads on the EID-relevant subspace.

The at-impact functional impulse response (IRF) of the expected inflation distribution to x_t is then a variance-efficient linear combination of the three baseline EID responses:

$$\phi_x = \rho_{x1} b_1 + \rho_{x2} b_2 + \rho_{x3} b_3, \tag{6}$$

where b_i are the baseline responses defined in (5). Because the external shock series are standardized, ϕ_x represents the EID's response to a one–standard deviation shock to x_t.[11] This construction preserves the interpretation of b_1, b_2, b_3 as the fundamental modes of distributional adjustment while enabling direct study of structural transmission mechanisms.

Let $b_i(h)$ denote the horizon-h response to the i-th functional shock e_i. For any external shock x_t with correlation weights $\rho_x = (\rho_{x1}, \rho_{x2}, \rho_{x3})'$, the functional IRF is

$$\phi_x(h) = \rho_{x1} b_1(h) + \rho_{x2} b_2(h) + \rho_{x3} b_3(h) \in H,$$

which at impact ($h = 0$) coincides with equation (6).

Moment responses follow directly by taking inner products with test functions—for instance, the mean response is $\Delta\mu_x(h) = \langle \phi_x(h), u \rangle$, where $u(\cdot)$ is the identity map.

Discrepancy in frequencies. A practical complication is the mismatch in data frequency: monetary-policy and gasoline-price shocks are available monthly, like the EID functional shocks, whereas government-spending and tax shocks are quarterly. We therefore measure the response of the monthly EID in the third month of each quarter.[12] In this setting, the quarterly response of EID to a shock x_s is the cumulative sum of the three monthly

[11] Responses to alternative shock magnitudes follow by linear scaling. For example, if one standard deviation of the monetary policy shock corresponds to a 30 bp funds rate move, then a 25 bp shock is obtained by scaling ϕ_x by 25/30.

[12] An alternative would be to estimate a quarterly FAR using every third monthly observation, but we retain the monthly model to exploit the full sample.

responses within quarter s:

$$\Phi_x = \phi_{x,0} + \phi_{x,1} + \phi_{x,2}.$$

Intuitively, the EID response in the final month of a quarter combines the contemporaneous effect with the propagated responses from the preceding two months.

For robustness, we also estimate a proxy-VAR following Chang et al. (2026), where the external shock enters as the first variable in an $(m+1)$-dimensional VAR. The EID response is then derived from the impulse responses of the remaining m components. Appendix D presents the proxy-VAR results, which closely mirror those obtained from the functional VAR.

In all cases, the VAR approximation to the dynamics of the expected inflation distribution allows us to apply well-established tools for inference. For example, the confidence bands shown in the empirical results are constructed using bootstrap methods (see Appendix E).

Finally, note that all transformations from the VAR representation to the functional impulse responses are linear. Hence, any orthogonal rotation of the baseline functional shocks—including the recursive normalization used here—produces identical functional IRFs for the EID. This invariance underscores that the recursive scheme serves purely as a normalization device rather than an identifying restriction.

5 Distributional Effects of Economic Shocks on EID

This section presents our empirical findings. We start with an overview of the shocks we implement. As we borrow the shocks from the literature, a purification procedure is needed to ensure that the chosen shocks are uncorrelated. To ensure a common sample across shocks in the purification step, we use the same sample period, 1991Q1 to 2006Q4, for all four economic shocks. Detailed results for each individual shock using the largest sample size possible are available in Appendix D.

5.1 Description of Economic Shocks

We examine four external shocks, all borrowed from the literature: monetary policy, government spending, personal income tax and gasoline price shocks.[13] We briefly discuss the shocks and their possible links to inflation expectations.

First, for monetary policy shock, we use the series constructed by Miranda-Agrippino and Ricco (2021), using a high-frequency identification strategy to build an instrumental variable that identifies monetary policy shocks robust to the presence of informational frictions. For government spending shocks, we refer to Auerbach and Gorodnichenko (2012), who use a trivariate VAR model with real government spending, real government receipts – direct and indirect taxes, net transfers to businesses and individuals – and real gross domestic product (GDP). For personal income tax shocks, we use the series constructed by Mertens and Ravn (2013) based on instrumental variable and a narrative approach. Finally, for gasoline price shocks, we use retail gasoline price, following Kilian and Zhou (2022) who show that gasoline price, rather than oil price, is more salient from the perspective of consumers and therefore is more influential in households' inflation expectations. Anderson et al. (2013) shows that households participating in the *Survey of Consumers* treat the real price of gasoline approximately as a random walk. Given that the change in the real price of gasoline is approximately the same as that in the nominal price of gasoline, we use the monthly changes in the nominal gasoline price as our gasoline price

[13]Except for gasoline price shocks, economic shocks are obtained using the replication codes and data made available by the authors of the referred literature.

shocks, drawn from FRED's "Consumer Price Index for All Urban Consumers: Gasoline (All Types) in the U.S. City Average."

5.2 Purification of the Shocks

The externally identified shocks we use may not be entirely pure surprises, which may compromise the conventional interpretation of the response of EID to these external shocks. To mitigate concerns that these series may share correlated information, we purify each external shock by purging its linear dependence on the other shocks we consider.

The purification process consists in estimating the following regression

$$x_t^i = \beta X_t^{-i} + \tilde{x}_t^i$$

where x_t^i are the shocks to be purified, and X_t^{-i} is a set of all shocks excluding shock i. We consider the residual of this regression ($\tilde{x}_t^i$) as a purified shock. This process requires us to consider, for practically all shocks, a subsample of the data available for each shock. This requirement results from the need to have a common sample for all shocks in the purification process. That sample becomes from 1991Q1 to 2006Q4.[14]

5.3 Linking EID Functional Shocks to Economic Shocks

As previously discussed, we link external shocks to EID dynamics by projecting them onto the functional shocks and constructing functional impulse responses from the resulting weights. Tables 2 and 3 summarize the estimated correlations between each type of external shock and the three functional shocks for one-year ahead and medium-run expectations.

[14]Note that for series of shocks available at a monthly frequency, such as the MP shock, the corresponding quarterly shock results from adding the three monthly observations for the quarter.

Three patterns emerge. First, monetary policy surprises transmit meaningfully into EID variation, with correlations around 10–18 percent for the short run and roughly 13 percent for the medium run. These responses occur throughout the quarter, consistent with a rapid pass-through of monetary policy news into inflation beliefs. Second, government spending shocks correlate nontrivially with the functional shocks—in some cases above 20 percent—yet ultimately produce negligible changes in the EID. This disconnect highlights that correlation alone does not imply an economically meaningful functional response: the linear combination of baseline components may lead to offsetting effects in the distribution. Third, personal income tax and gasoline price shocks display the strongest and most systematic correlations across horizons. For personal income taxes, the relationship is concentrated in the first month of the quarter, while gasoline price shocks exhibit persistent influence throughout the quarter. These patterns foreshadow the clear functional responses documented below. All correlations are precisely estimated, and confidence intervals for both the correlation coefficients and the resulting impulse responses are constructed using a bootstrap procedure detailed in Appendix E.

5.4 Results

The primary contribution of this paper is to examine inflation expectations as a full distribution rather than through a small set of sample statistics. At the same time, working with the density representation allows us to recover familiar moments such as the mean, standard deviation, and skewness as by-products.

5.4.1 Monetary Policy Shocks

Figure 6 shows the impact of contractionary monetary policy shocks on selected moments of the distribution of inflation expectations. A one-percentage-point policy tighten-

Table 2. *Correlations between Functional Shocks and Economic Shocks: One-Year Ahead EIDs*

	Monetary Policy	Govt. Spending	Personal Income Tax	Gasoline Price
corr($x_s, e_{1,t}$)	$-$**0.1062** (0.0979)	0.0229 (0.1266)	**0.1774** (0.1120)	$-$**0.1645** (0.1321)
corr($x_s, e_{2,t}$)	-0.0607 (0.1050)	-0.0651 (0.1297)	**0.2744** (0.1427)	**0.3002** (0.1506)
corr($x_s, e_{3,t}$)	0.0183 (0.1159)	-0.0971 (0.1292)	0.1067 (0.1154)	**0.2361** (0.1518)
corr($x_s, e_{1,t+1}$)	$-$**0.0998** (0.0895)	$-$**0.2467** (0.1220)	0.0180 (0.1058)	**0.1522** (0.1210)
corr($x_s, e_{2,t+1}$)	-0.0405 (0.1647)	0.0176 (0.1281)	-0.0191 (0.1463)	**0.3714** (0.1164)
corr($x_s, e_{3,t+1}$)	**0.1242** (0.1152)	0.0817 (0.1354)	-0.0485 (0.1351)	-0.0189 (0.1153)
corr($x_s, e_{1,t+2}$)	-0.0207 (0.0882)	**0.1355** (0.1264)	-0.0730 (0.1217)	**0.2275** (0.2101)
corr($x_s, e_{2,t+2}$)	$-$**0.1805** (0.1662)	-0.0873 (0.1308)	-0.0406 (0.1402)	**0.1611** (0.1570)
corr($x_s, e_{3,t+2}$)	-0.0264 (0.1748)	**0.1391** (0.1118)	0.0443 (0.1162)	-0.0269 (0.1473)

Notes: Correlations between the three functional shocks of one-year ahead EIDs (first, second, and third month of the quarter) and each of the following four economic shocks are considered: (i) contractionary monetary policy shock (Miranda-Agrippino and Ricco (2021)); (ii) government spending (Auerbach and Gorodnichenko (2012)); (iii) personal income tax increase (Mertens and Ravn (2013)); and (iv) gasoline price changes (FRED). Bootstrapped standard errors are reported in parentheses.

ing reduces both the average and the dispersion of expected inflation across short- and medium-run horizons (first two columns), while the effect on skewness is negligible. These moment-based results suggest that monetary policy lowers the overall level of inflation expectations and compresses their cross-sectional spread.

To examine the full distributional effects, Figure 7 plots the impulse responses of the expected-inflation distributions (EIDs) to a contractionary monetary policy shock, for one-year ahead and medium-run horizons. For short-term expectations (top-left panel), the distribution becomes more concentrated around zero following the policy shock, as indicated by the bulge at the center. The probability mass in the upper tail—households expecting inflation above five percent—declines substantially, implying that fewer households anticipate high inflation. Medium-term expectations (top-right panel) display a similar but milder pattern: the distribution shifts leftward and narrows, though the in-

Table 3. Correlations between Functional Shocks and Economic Shocks: Medium-run EIDs

	Monetary Policy	Govt. Spending	Personal Income Tax	Gasoline Price
$\mathrm{corr}(x_s, e_{1,t})$	-0.0292 (0.0855)	-0.0213 (0.1171)	**0.2198** (0.1150)	-0.0175 (0.1265)
$\mathrm{corr}(x_s, e_{2,t})$	**0.1370** (0.1147)	**0.2352** (0.1146)	$-$**0.3121** (0.1066)	$-$**0.1460** (0.1225)
$\mathrm{corr}(x_s, e_{3,t})$	**0.1322** (0.0915)	0.0134 (0.1454)	$-$**0.1882** (0.1355)	$-$**0.1895** (0.1048)
$\mathrm{corr}(x_s, e_{1,t+1})$	-0.0670 (0.1388)	-0.0092 (0.1230)	-0.0741 (0.1278)	**0.1524** (0.1151)
$\mathrm{corr}(x_s, e_{2,t+1})$	0.1193 (0.1244)	0.0895 (0.1062)	$-$**0.1387** (0.1180)	0.0742 (0.1208)
$\mathrm{corr}(x_s, e_{3,t+1})$	0.1054 (0.1430)	-0.0394 (0.1303)	-0.0734 (0.1157)	$-$**0.1570** (0.1319)
$\mathrm{corr}(x_s, e_{1,t+2})$	-0.0491 (0.1042)	0.1368 (0.1418)	-0.0043 (0.1353)	0.0963 (0.1890)
$\mathrm{corr}(x_s, e_{2,t+2})$	-0.0327 (0.1091)	-0.0408 (0.1409)	$-$**0.2942** (0.1173)	-0.0474 (0.1420)
$\mathrm{corr}(x_s, e_{3,t+2})$	-0.0155 (0.0964)	**0.1623** (0.1183)	-0.0565 (0.1010)	$-$**0.1595** (0.1142)

Notes: Correlations between the three functional shocks of medium-run EIDs (first, second, and third month of the quarter) and each of the following four economic shocks are considered: (i) contractionary monetary policy shock (Miranda-Agrippino and Ricco (2021)); (ii) government spending (Auerbach and Gorodnichenko (2012)); (iii) personal income tax increase (Mertens and Ravn (2013)); and (iv) gasoline price changes (FRED). Bootstrapped standard errors are reported in parentheses.

crease in mass occurs primarily at low positive values rather than at negative inflation expectations.

To quantify these shifts, we divide the range of inflation expectations into ten intervals based on pre-shock percentiles of the EID (the 10^{th}, 20^{th}, $\cdots$, 90^{th} percentiles). By construction, each decile initially contains ten percent of the households. The bottom panels of Figure 7 display the median post-shock change in the frequency of each decile (dots) with 68-percent confidence intervals (vertical lines). The horizontal reference line at ten percent marks the pre-shock share of each bin, allowing a visual assessment of statistically significant deviations.

Following a 100-basis-point monetary tightening, the shares of households in the 2^{nd} through 4^{th} deciles—corresponding roughly to inflation expectations between -5.4 and 2.4 percent —increase significantly. In contrast, the upper deciles (7^{th} to 10^{th}), representing

Figure 6. *Impact of Monetary Policy Shocks on Specific Moments of One-Year Ahead and Medium-Run EIDs*

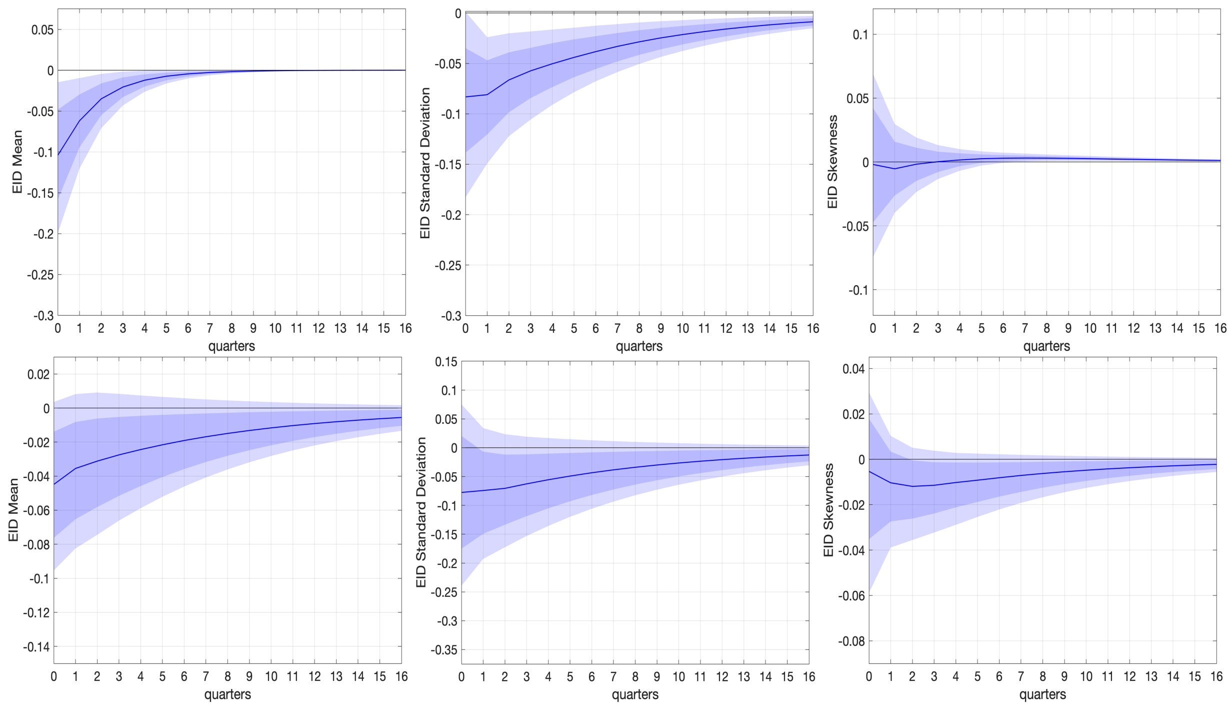

Notes: Each column shows the impact of contractionary monetary policy shocks of 1 percent on the mean (left), standard deviation (center), and skewness (right) of expected inflation distributions (EIDs). The top (bottom) row shows the impact on one-year ahead (medium-run) inflation expectations. The sample runs from 1991Q1 to 2006Q4. The IRFs include two confidence bands (68 and 90 percent) estimated using bootstrap methods.

inflation expectations above 4.7 percent, exhibit a clear and statistically significant decline. These shifts explain the moment-based results in Figure 6: the mean of inflation expectations falls because a larger fraction of households now anticipate low inflation, while fewer expect high inflation. The contraction of the upper tail also reduces the standard deviation, although not sufficiently to alter skewness, which remains broadly unchanged.

For medium-run expectations (bottom-right panel), the distributional adjustments are qualitatively similar but quantitatively smaller. The share of households expecting high inflation (above 7.1 percent) declines, while the shares at low-positive values increase slightly. The resulting pattern again reflects a downward and narrowing shift in the distribution of medium-run inflation expectations, consistent with the interpretation that

Figure 7. *One-year Ahead and Medium-Run EID Responses to Monetary Policy Shocks*

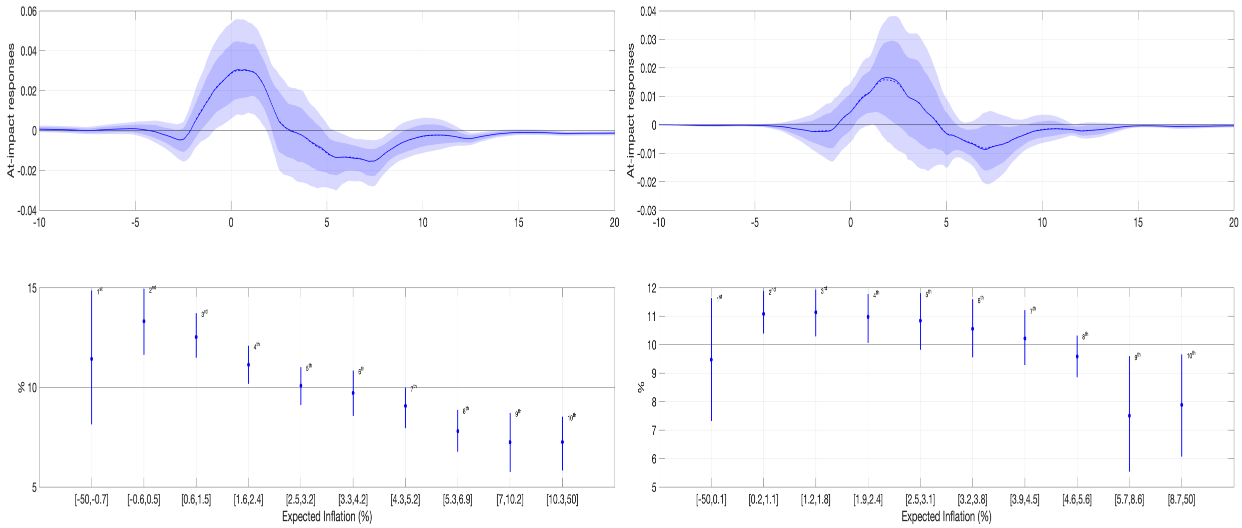

Notes: The top charts show the at-impact response of one-year ahead expected inflation expectations (left) and medium-run inflation expectations (right) following a contractionary monetary policy shock of 1 percent. The sample runs from 1991Q1 to 2006Q4. The IRFs include two confidence bands (68 and 90 percent) estimated using bootstrap methods. The bottom charts show the estimated post-shock probabilities of each decile following the same monetary policy shock for one-year ahead (left) and medium-run inflation expectations (right), compared to the reference pre-shock probability of 10 percent signified by a black horizontal line. For each decile, the dot represents the average post-shock probability, and the vertical line indicates a 68 percent confidence band.

contractionary monetary policy not only lowers expected inflation but also compresses belief heterogeneity among households.

Consistent with Doh et al. (2025), we find that contractionary monetary policy shocks reduce household inflation expectations most strongly in the upper part of the distribution, encompassing the upper middle and upper tail. The decile responses indicate a statistically and economically meaningful decline in the share of households with high inflation expectations following a policy tightening. However, as one moves toward the far right end of the distribution—particularly for medium-run expectations—the estimated effects become less precise, with wider confidence bands. This pattern suggests that while monetary policy does affect households with high inflation expectations, beliefs at the extreme upper tail are more heterogeneous and respond less uniformly, potentially reflecting expectation formation driven by factors beyond contemporaneous monetary signals.

5.4.2 Fiscal Policy Shocks: Government Spending and Personal Income Tax Shocks

We next examine the effects of fiscal policy shocks on expected inflation distributions (EIDs), beginning with government spending shocks identified following Auerbach and Gorodnichenko (2012). Figure 8 reports the responses of the mean (left column), standard deviation (center column), and skewness (right column) of inflation expectations to these shocks. Across both one-year ahead and medium-run horizons, no statistically significant effects are detected for any of the moments, indicating that changes in government spending have little contemporaneous influence on either the level or dispersion of household inflation expectations.

Figure 8. *Impact of Government Spending Shocks on Specific Moments of One-Year Ahead and Medium-Run EIDs*

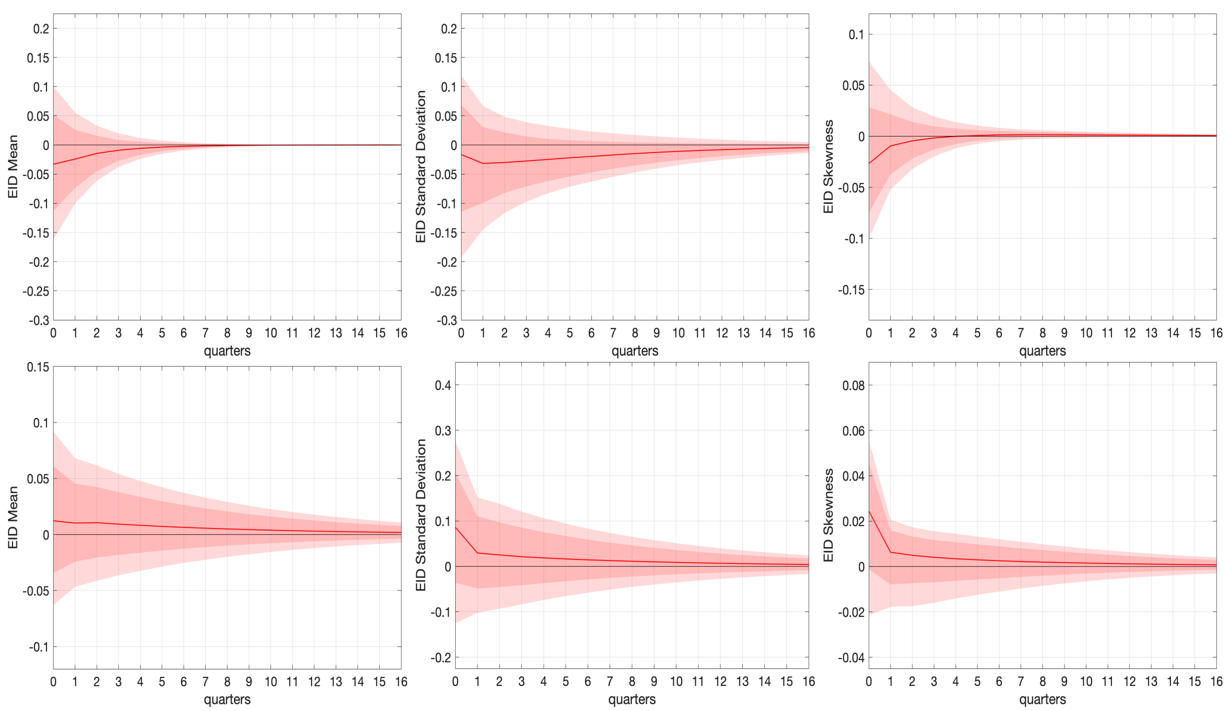

Notes: Each column shows the impact of government spending shocks on the mean (left), standard deviation (center), and skewness (right). The top row reports for the results for one-year ahead expected inflation distributions (EIDs), using the sample from 1991Q1 to 2006Q4. The bottom row reports the results for medium-run inflation expectations, using the sample from 1991Q1 to 2006Q4.

Figure 9 further explores the full distributional effects by plotting the at-impact responses of the entire EIDs for one-year ahead (left) and medium-run (right) inflation expectations. Consistent with the moment-based evidence, the distributional responses are small and statistically indistinguishable from zero across all inflation-expectation levels. In short, government spending shocks appear not to meaningfully alter the cross-sectional distribution of inflation expectations, suggesting that households do not view variations in public spending as informative about future inflation dynamics.

Figure 9. *One-year Ahead EID and Medium-Run EID Responses to Government Spending Shocks*

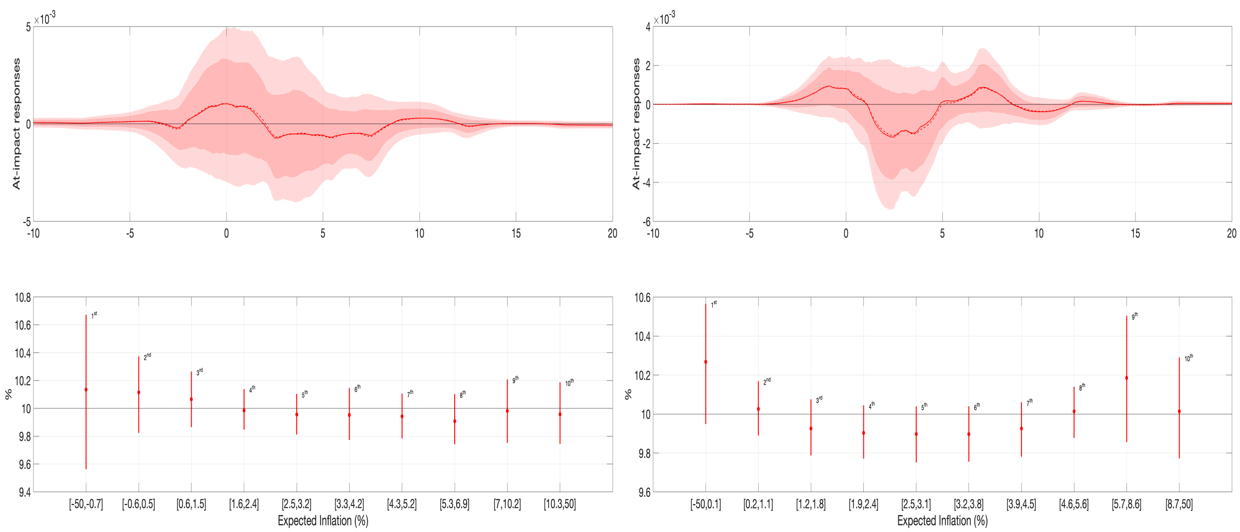

Notes: The top charts show the at-impact responses of one-year ahead (left) and medium-run inflation expectations (right) to government spending shocks using the sample from 1991Q1 to 2006Q4. The bottom charts shows the estimated post-shock probabilities of decile bins for one-year ahead (left) and medium-run inflation expectations (right) following the same shock, compared to the reference 10 percent probability line for each decile before the shock. Footnote otherwise identical to the bottom chart of Figure 7.

The absence of any measurable response of inflation expectations to government spending shocks may raise questions of statistical power. However, our estimated standard errors imply that even relatively small shifts in the mean—on the order of a few hundredths of a percentage point—would have been detectable at conventional significance levels. Thus, the lack of a significant response is unlikely to reflect insufficient precision but rather indicates that household inflation expectations are largely unresponsive to this type of fiscal intervention.

We then turn to personal income tax shocks. Figure 10 displays the impulse responses of the mean, standard deviation, and skewness of inflation expectations to such shocks. Unlike government spending shocks, personal income tax shocks produce economically and statistically significant effects, particularly for medium-run expectations. Specifically, a positive tax shock raises the mean of medium-run inflation expectations on impact and is accompanied by an increase in both dispersion and skewness, implying that the right tail of the distribution expands as some households anticipate higher inflation in response to tax increases.

Figure 10. *Impact of Personal Income Tax Shocks on Specific Moments of One-Year Ahead and Medium-Run EIDs*

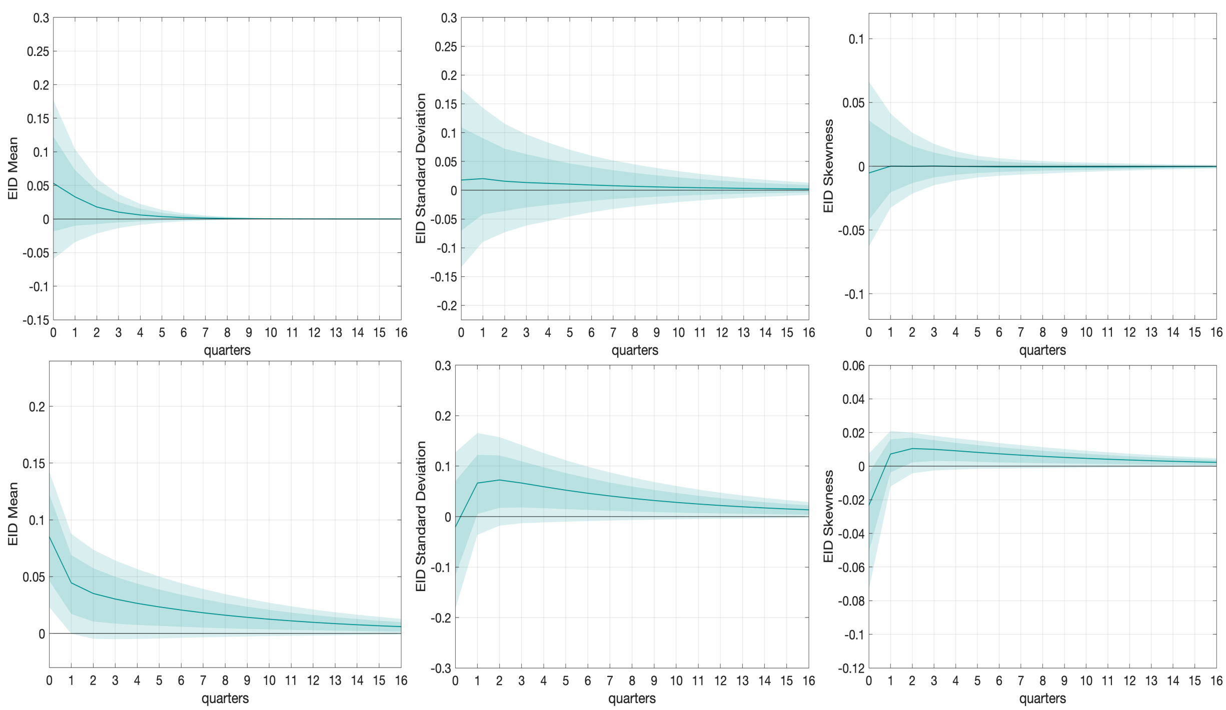

Notes: Each column shows the impact of personal income tax shock on the mean (left), standard deviation (center), and skewness (right). The top (bottom) row reports the results for one-year ahead (medium-run) expected inflation distributions, using the sample from 1983Q1 to 2006Q4.

Figure 11 illustrates these results for the full distributions. For one-year ahead expectations, the effects remain statistically insignificant. By contrast, for medium-run expectations, the EID shifts to the right: the frequency of responses associated with negative or low-to-moderate inflation (around 2–3 percent) declines, while the frequency of higher expected inflation (above 4 percent) rises. The decile analysis confirms this pattern—post-shock, fewer households fall in the lower half of the distribution (first to fifth deciles), while more are found in the sixth decile and above.

Figure 11. *One-year Ahead and Medium-Run EID Responses to Personal Income Tax Shocks*

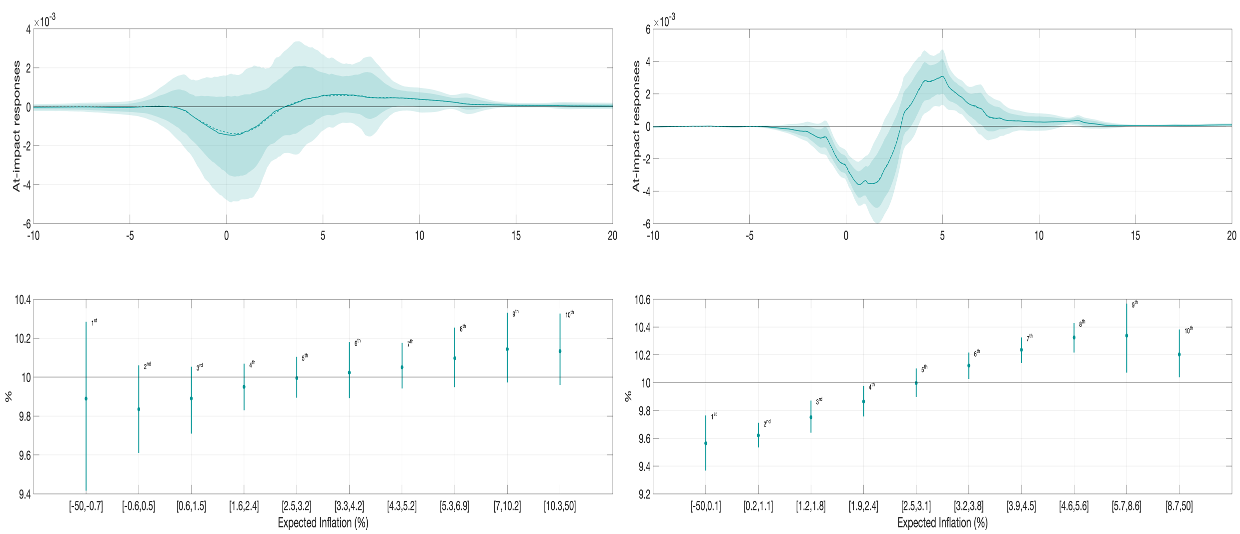

Notes: The top charts show the at-impact response of one-year ahead (left) and medium-run (right) inflation expectations to personal income tax shock using the sample from 1991Q1 to 2006Q4. The bottom charts shows the estimated post-shock probabilities of decile bins for one-year ahead (left) and medium-run (right) EIDs following the same shock, compared to the reference 10 percent probability line for each decile before the shock. Footnote otherwise identical to the bottom chart of Figure 7.

Together, these findings indicate that different fiscal instruments have distinct expectation effects. Government spending shocks do not materially affect inflation beliefs, whereas personal income tax shocks raise and skew the distribution of medium-run expectations upward, consistent with households associating tax changes with broader price-level pressures or long-run fiscal concerns.

5.4.3 Gasoline Price Shocks

Finally, we examine the impact of gasoline price shocks on the expected inflation distributions (EIDs). These shocks serve as a useful benchmark, since it is a well-established stylized fact that increases in gasoline prices raise inflation expectations. Figure 12 displays the responses of selected moments to a positive gasoline price shock. A surprise increase in gasoline prices leads to higher average inflation expectations for both horizons, with a substantially larger effect for one-year ahead expectations (left column). The at-impact response of short-term expectations is nearly three times stronger than that of medium-run expectations. For the medium-run horizon (bottom center), the standard deviation rises, indicating greater disagreement about future inflation, whereas no significant change in dispersion is observed for the short-term horizon. Skewness remains broadly unchanged for one-year ahead expectations but exhibits a small decline in the first quarter followed by a gradual increase after three quarters for the medium-run expectations.

Figure 13 reports the corresponding distributional responses of the EIDs. Consistent with the moment-based results and with prior findings in the literature (Harris et al., 2009; Coibion and Gorodnichenko, 2015), gasoline price shocks shift inflation expectations upward. For one-year ahead EIDs (top-left panel), the frequency of inflation expectations below three percent declines, while that of expectations above three percent increases, confirming that gasoline price shocks raise the average level of short-term inflation expectations. In contrast, medium-run expectations (top-right panel) respond more heterogeneously: both the lower and upper tails gain mass, indicating that some households expect the inflationary effect to persist, while others anticipate a temporary impact or future reversal.

The decile analysis (bottom panels) reinforces this interpretation. For one-year ahead expectations (bottom-left panel), the share of households expecting inflation below 1.5

Figure 12. *Impact of Gasoline Price Shocks on Specific Moments of One-Year Ahead and Medium-Run EIDs*

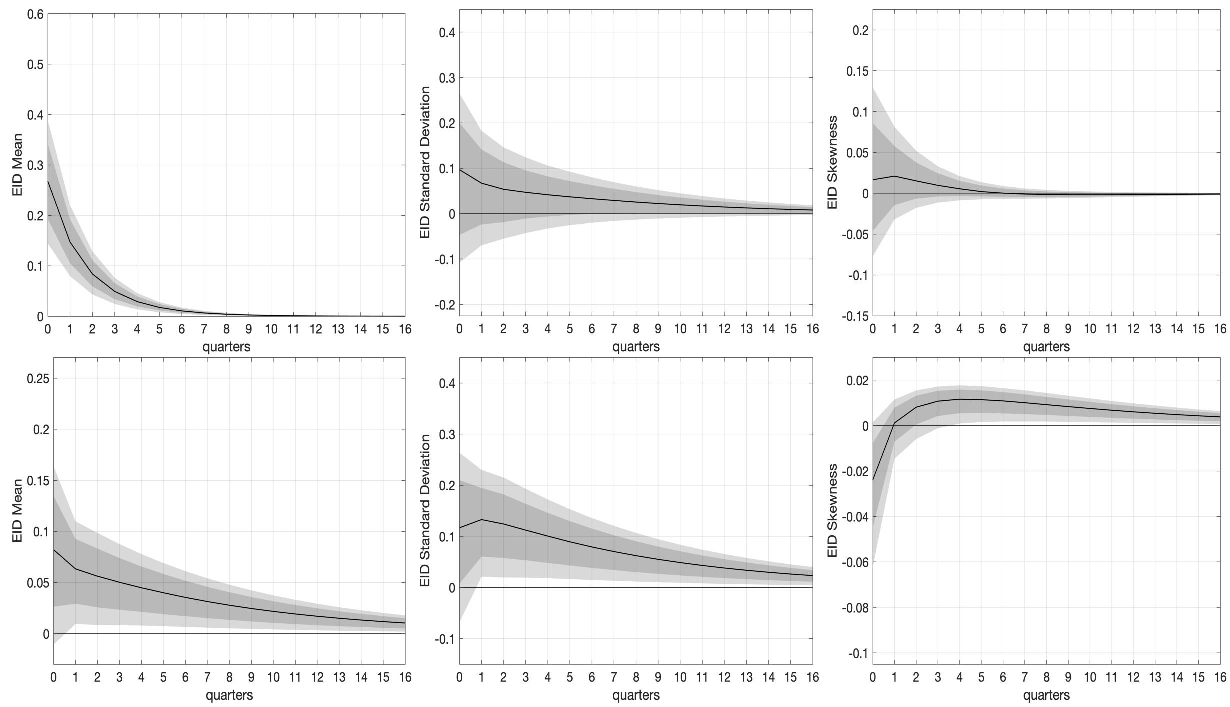

Notes: Each column shows the impact of gasoline price shock on the mean (left), standard deviation (center), and skewness (right). The top (bottom) row reports the results for one-year ahead (medium-run) expected inflation distributions (EIDs), using the sample from 1991Q1 to 2006Q4.

percent declines significantly, while the share above 2.5 percent increases. For medium-run expectations (bottom-right panel), however, the share of households expecting very low inflation (below 0.1 percent) rises alongside an increase in higher deciles. This pattern points to widening disagreement among households about the long-term implications of gasoline price shocks for inflation.

Taken together, these results validate our empirical approach: the estimated EIDs reproduce a well-documented fact—gasoline price increases are inflationary in the short run—while also revealing a novel distributional dimension (e.g., Harris et al., 2009; Coibion and Gorodnichenko, 2015). Short-term inflation expectations respond relatively homogeneously upward, whereas medium-run expectations display greater heterogeneity, sug-

Figure 13. *One-year Ahead EID and Medium-Run EID Responses to Gasoline Price Shocks*

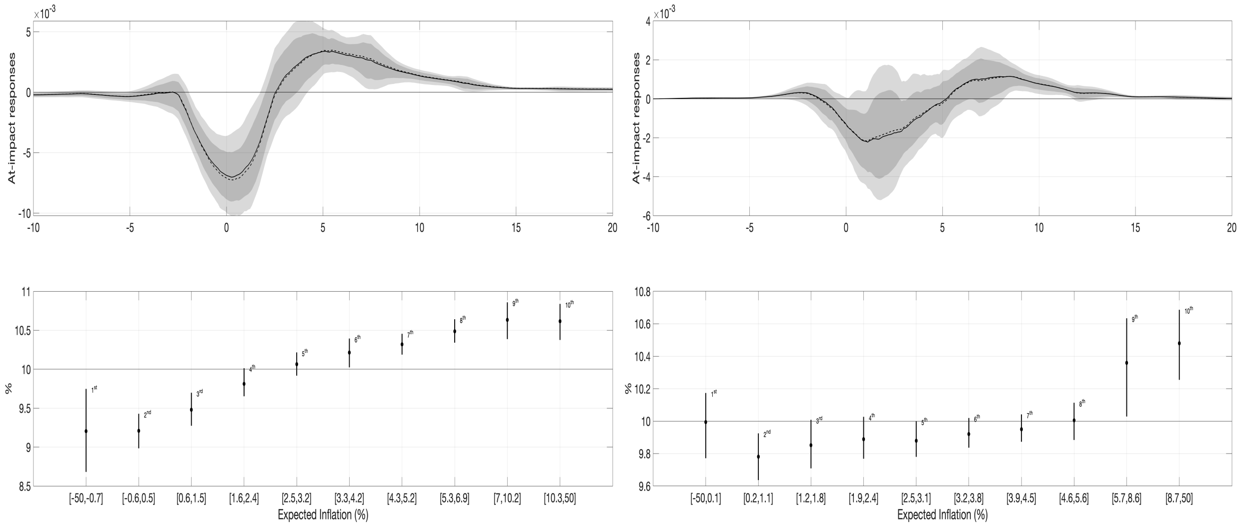

Notes: The top charts show the at-impact response of one-year ahead (left) and medium-run (right) expected inflation distributions (EIDs) to gasoline price shocks using the sample from 1991Q1 to 2006Q4. The bottom charts shows the estimated post-shock probabilities of decile bins for one-year ahead (left) and medium-run inflation expectations (right), following the same shock, compared to the reference 10 percent probability line for each decile before the shock. Footnote otherwise identical to the bottom charts of Figure 7.

gesting that households differ in how they map energy-price shocks into future inflation dynamics.

5.5 Interpretation and Plausibility of Findings

Taken together, the results across policy and energy shocks reveal that different macroeconomic disturbances shape household inflation expectations through distinct channels.

We now attempt to provide plausible explanations for these findings using established transmission mechanisms in the literature. This task is inherently challenging, given the limited empirical evidence on how economic shocks affect the *distribution* of inflation expectations rather than their average. The results underscore the need for further research into the behavioral and informational foundations of household expectations.

Monetary policy shocks. Contractionary monetary policy compresses the distribution and anchors beliefs toward moderate inflation, largely affecting the center and upper middle of the EID but leaving the extreme right tail unchanged. A priori, tightening monetary shocks could either raise or lower inflation expectations. While the standard channel of dampening aggregate demand would indicate lower inflation expectations associated with monetary policy tightening, empirical studies have shown that contractionary monetary policy can, under some conditions, generate price puzzles or neo-Fisherian effects (Uribe, 2022; Sims, 1992; Hanson, 2004; Rusnák et al., 2013). In relation to the most recent evidence, studies show that U.S. monetary tightening during the 2021–22 inflation episode led certain households—especially those with optimistic economic outlooks—to develop deflationary expectations (Armantier et al., 2022). Our results support the conventional Phillips-Curve mechanism: tighter policy is expected to slow economic activity, reduce inflationary pressures, and lower expected inflation. At the same time, the unresponsiveness of the extreme right tail suggests that a subset of households forms expectations through non-monetary channels—possibly linked to real-side or fiscal factors—an interpretation that complements the findings of Guillochon (2024), who documents that news about energy and taxes rather than monetary developments predominantly drives household inflation expectations in France.

Government-spending shocks. Our analysis shows that government-spending innovations elicit almost no response. The unresponsiveness of household inflation expectations to government-spending shocks can stem from several mechanisms. A general lack of awareness of fiscal details, combined with confidence in the central bank's commitment to price stability, may dampen reactions to fiscal expansions. Moreover, the delayed transmission of government-spending effects and the credibility of monetary authorities can keep expectations anchored despite fiscal stimulus (Leeper et al., 2010). Historical experience—where past spending increases did not yield visible inflation—may also shape expectations. In open economies, global slack and external price conditions further moderate how households perceive the inflationary consequences of domestic fiscal policy.

Personal-income-tax shocks. Our results suggest that a higher personal income tax shock shifts and leads to a higher skewness of the medium-run distribution. The literature documents that an increase in taxes typically has contractionary effects, accompanied by a significant drop in inflation (Alesina and Ardagna, 2010; Romer and Romer, 2010). To the extent that any change in personal income tax directly affects household disposable income, one can conjecture that household inflation expectations will be adjusted in response to a change in personal income tax. Our finding of the asymmetric and horizon-dependent effects of personal-income-tax shocks could be explained by considering near- and long-term channels jointly. In the near term, higher taxes may restrain consumption and reduce inflationary pressures. Over longer horizons, however, households may interpret tax increases as precursors to future fiscal expansion or as signals of debt stabilization that could eventually prompt monetary accommodation or higher public spending (Leeper and Nason, 2010). This forward-looking interpretation aligns with models of imperfect common knowledge: agents form expectations not only about fundamentals but also about others' beliefs (Angeletos and Lian, 2018; García-Schmidt and Woodford, 2019). The resulting heterogeneity of forecasts produces the observed upward skewness in medium-run EIDs.

Gasoline-price shocks. Finally, gasoline-price shocks validate the empirical framework by reproducing the well-known positive response of inflation expectations to energy-price increases (Harris et al., 2009; Coibion and Gorodnichenko, 2015). Short-term expectations shift uniformly upward, whereas medium-run expectations exhibit wider disagreement, as some households expect transitory effects while others foresee persistent inflation. This pattern highlights the salience of energy prices as a real-side anchor of expectations and their role in amplifying heterogeneity.

Overall, the differentiated responses across monetary, fiscal, and real shocks suggest that households form inflation expectations through multiple, partly segmented channels. Monetary policy primarily affects belief anchoring; fiscal policies influence asymmetric, longer-term expectations; and salient real-side shocks like gasoline prices increase disagreement. These results call for theoretical models that explicitly link heterogeneous information and behavioral updating to the full distribution of inflation expectations.

6 Conclusion

This paper provides new evidence on how different types of economic shocks shape the entire distribution of household inflation expectations over both short- and medium-term horizons, using a functional time-series framework. We show that contractionary monetary policy shocks lower the average level of inflation expectations and compress their distribution, reducing the share of households expecting high inflation while slightly increasing the share anticipating deflation. Fiscal shocks, in contrast, exhibit heterogeneous effects: government spending shocks have little impact, whereas increases in personal income taxes significantly raise medium-term inflation expectations and widen their dispersion. These results suggest that fiscal and monetary policies operate through distinct channels in shaping household beliefs about inflation.

Gasoline price shocks, which serve as a benchmark for salient real-side disturbances, also exert strong and persistent effects on inflation expectations. In line with the literature, higher gasoline prices raise one-year ahead expectations, but we further show that they increase disagreement in the medium-run, as households differ in how they extrapolate energy-price shocks into future inflation.

Our findings contribute to the policy debate on the anchoring of household inflation expectations. The evidence highlights that monetary, fiscal, and real shocks influence not only the average level of expected inflation but also its dispersion and asymmetry. Understanding these distributional effects is essential for assessing how expectations shape macroeconomic dynamics and the effectiveness of stabilization policy. Future research should aim to identify the precise mechanisms—informational, behavioral, and institutional—that drive these heterogeneous responses, and to examine how shifts in the distribution of inflation expectations feed back into key macroeconomic outcomes such as inflation, unemployment, and growth.

References

Alesina, A. and Ardagna, S. (2010). Large changes in fiscal policy: Taxes versus spending. *Tax policy and the Economy*, 24(1):35–68.

Anderson, S. T., Kellogg, R., and Sallee, J. M. (2013). What do consumers believe about future gasoline prices? *Journal of Environmental Economics and Management*, 66(3):383–403.

Andre, P., Pizzinelli, C., Roth, C., and Wohlfart, J. (2022). Subjective models of the macroeconomy: Evidence from experts and representative samples. *The Review of Economic Studies*, 89(6):2958–2991.

Angeletos, G.-M. and Lian, C. (2018). Forward guidance without common knowledge. *American Economic Review*, 108(9):2477–2512.

Armantier, O., Koşar, G., Somerville, J., Topa, G., Van der Klaauw, W., and Williams, J. C. (2022). The curious case of the rise in deflation expectations. *Federal Reserve Bank of New York Staff Report*, 1037.

Auerbach, A. J. and Gorodnichenko, Y. (2012). Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy*, 4(2):1–27.

Bjørnland, H. C., Chang, Y., and Cross, J. L. (2025). Oil and the stock market revisited: A mixed functional var approach. *Quantitative Economics*. Forthcoming.

Bosq, D. (2000). Linear processes in function spaces, volume 149 of lecture notes in statistics.

Carroll, C. D. (2003). Macroeconomic expectations of households and professional forecasters. *the Quarterly Journal of economics*, 118(1):269–298.

Chang, M., Chen, X., and Schorfheide, F. (2024a). Heterogeneity and aggregate fluctuations. *Journal of Political Economy*, 132(12):4021–4067.

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

Chang, Y., Kim, S., and Park, J. Y. (2025). How do macroaggregates and income distribution interact dynamically? A novel structural mixed autoregression with aggregate and functional variables. *CAEPR Working Paper 2025-002*.

Chang, Y., Miller, J. I., and Park, J. Y. (2024b). What drives temperature anomalies: A functional svar approach. CAEPR Working Paper 2024-007, Indiana University.

Chang, Y., Park, J. Y., and Pyun, D. (2021). From functional autoregressions to vector autoregressions. Mimeo, Indiana University.

Coibion, O. and Gorodnichenko, Y. (2015). Is the phillips curve alive and well after all? inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics*, 7(1):197–232.

Coibion, O., Gorodnichenko, Y., Kumar, S., and Pedemonte, M. (2020). Inflation expectations as a policy tool? *Journal of International Economics*, 124:103297.

Coibion, O., Gorodnichenko, Y., and Weber, M. (2021). Fiscal policy and households' inflation expectations: Evidence from a randomized control trial. *NBER Working Paper*, 28485.

D'Acunto, F., Malmendier, U., and Weber, M. (2023). What do the data tell us about inflation expectations? *Handbook of Economic Expectations*, pages 133–161.

Doh, T., Lee, J. H., and Park, W. Y. (2025). Heterogeneity in household inflation expectations and monetary policy. *Journal of Financial Econometrics*, 23(1):nbae034.

Friedman, M. (1968). The role of monetary policy. *American Economic Review*, 58(1):1–17.

García-Schmidt, M. and Woodford, M. (2019). Are low interest rates deflationary? a paradox of perfect-foresight analysis. *American Economic Review*, 109(1):86–120.

Grigoli, F., Gruss, B., and Lizarazo, S. (2020). Monetary policy surprises and inflation expectation dispersion. *IMF Working Paper 252*, International Monetary Fund, Washington DC.

Guillochon, J. (2024). News-driven household macroeconomic expectations: Regional vs. national telecast information. *Journal of Money, Credit and Banking*.

Hall, P. and Horowitz, J. L. (2007). Methodology and convergence rates for functional linear regression. *Annals of Statistics*, 35:70–91.

Hanson, M. S. (2004). The "price puzzle" reconsidered. *Journal of Monetary Economics*, 51(7):1385–1413.

Harris, E. S., Kasman, B. C., Shapiro, M. D., and West, K. D. (2009). Oil and the macroeconomy: Lessons for monetary policy. In *US Monetary Policy Forum Report*, volume 23, page 2015.

Inoue, A. and Rossi, B. (2019). The effects of conventional and unconventional monetary policy on exchange rates. *Journal of International Economics*, 118:419–447.

Kilian, L. and Lütkepohl, H. (2017). *Structural Vector Autoregressive Analysis*. Cambridge University Press.

Kilian, L. and Zhou, X. (2022). The impact of rising oil prices on us inflation and inflation expectations in 2020–23. *Energy Economics*, 113:106228.

Leeper, E. M. and Nason, J. M. (2010). Government budget constraint. In *Monetary Economics*, pages 108–117. Springer.

Leeper, E. M., Walker, T. B., and Yang, S.-C. S. (2010). Government investment and fiscal stimulus. *Journal of Monetary Economics*, 57(8):1000–1012.

Mertens, K. and Ravn, M. O. (2013). The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States. *American Economic Review*, 103(4):1212–1247.

Miranda-Agrippino, S. and Ricco, G. (2021). The transmission of monetary policy shocks. *American Economic Journal: Macroeconomics*, 13(3):74–107.

Park, J. Y. and Qian, J. (2012). Functional regression of continuous state distributions. *Journal of Econometrics*, 167(2):397 – 412. Fourth Symposium on Econometric Theory and Applications (SETA).

Phelps, E. S. (1967). Phillips curves, expectations of inflation and optimal unemployment over time. *Economica*, pages 254–281.

Plagborg-Møller, M. and Wolf, C. K. (2021). Local projections and vars estimate the same impulse responses. *Econometrica*, 89(2):955–980.

Ramsay, J. O. (2004). Functional data analysis. *Encyclopedia of Statistical Sciences*, 4.

Reis, R. (2022). Losing the inflation anchor. *Brookings Papers on Economic Activity*, 2021(2):307–379.

Romer, C. D. and Romer, D. H. (2010). The macroeconomic effects of tax changes: estimates based on a new measure of fiscal shocks. *American Economic Review*, 100(3):763–801.

Rusnák, M., Havranek, T., and Horváth, R. (2013). How to solve the price puzzle? a meta-analysis. *Journal of Money, Credit and Banking*, 45(1):37–70.

Sims, C. A. (1992). Interpreting the macroeconomic time series facts: The effects of monetary policy. *European Economic Review*, 36(5):975–1000.

Uribe, M. (2022). The neo-fisher effect: Econometric evidence from empirical and optimizing models. *American Economic Journal: Macroeconomics*, 14(3):133–62.

Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American Statistical Association*, 101(476):1418–1429.

For Online Publication

A Hilbert Space Framework for Functional Data Analysis

This appendix provides the mathematical foundations underlying the functional autoregression (FAR) methodology used in the main text. The framework follows Chang et al. (2021), which we summarize here for completeness. Readers familiar with functional data analysis may proceed directly to Section 4.

A.1 Basic Hilbert Space Concepts

The functional time series (f_t) represents a time series of random elements taking values in the Hilbert space H, where H is endowed with the inner product $\langle \cdot, \cdot \rangle$ and the norm $\| \cdot \|$ given by

$$\langle f, g \rangle = \int f(r)g(r)dr \quad \text{and} \quad \|h\|^2 = \int h^2(r)dr$$

for all f, g and h in H.

In addition to the inner product and norm, we introduce the tensor product in H. The tensor product $f \otimes g$ with any given f and g in H is a linear operator on H defined as

$$(f \otimes g)v = \langle v, g \rangle f$$

for all v in H. If $H \equiv \mathbb{R}^n$, we have $f \otimes g = fg'$, i.e., $f \otimes g$ reduces to the outer product, in contrast to the inner product $\langle f, g \rangle = f'g$, where f' and g' are the transposes of f and g.

For a random function f taking values in H, we define $\mathbb{E}f$ to be a function in H such that

$$\mathbb{E}\langle v, f\rangle = \langle v, \mathbb{E}f\rangle$$

for all $v \in H$, whose existence is guaranteed by the Riesz representation theorem. If f and g are random functions taking values in H, then their covariance operator $\mathbb{E}(f \otimes g)$ is generally defined as a linear operator satisfying

$$\langle u, [\mathbb{E}(f \otimes g)]v\rangle = \mathbb{E}\langle u, f\rangle\langle v, g\rangle$$

for all u and v in H. In particular, for the functional error (ε_t) assumed to be white noise in equation (1), we let $\mathbb{E}\varepsilon_t = 0$ for all $t \geqslant 1$, and (ε_t) be serially uncorrelated with $\mathbb{E}(\varepsilon_t \otimes \varepsilon_t) = \Sigma$ for all $t \geqslant 1$.

A.2 Finite-Dimensional Approximation via Basis Projection

For a given orthonormal basis (v_i) of H, we write f in H as

$$f = \sum_{i=1}^{\infty}\langle v_i, f\rangle v_i,$$

and approximate it as

$$f \approx \sum_{i=1}^{m}\langle v_i, f\rangle v_i, \tag{7}$$

where m is the truncation number that must be selected appropriately.[15] We let V be the subspace of H spanned by the sub-basis $(v_i)_{i=1}^m$, and denote by P the Hilbert space projection on the m-dimensional subspace V. The approximation in (7) can be viewed as the projection of f onto V, given by $Pf = \sum_{i=1}^m\langle v_i, f\rangle v_i$. This compact representation

[15]The choice of m can complement the selection of the lag order p in a joint framework in which (m, p) is chosen by minimizing the out-of-sample MSFE in a rolling-window forecasting exercise. In our case, this procedure yielded selections consistent with those obtained using the information criteria.

follows directly from the orthonormality of the basis functions. Once f is approximated by an m-dimensional element Pf in V, we may represent it as an m-dimensional vector. Consequently, it is well expected that the FAR in equation (1) can be represented by a VAR upon this approximation, as we now explain.

We approximate the FAR in equation (1) as

$$\begin{aligned} f_t &= A_1 P f_{t-1} + A_2 P f_{t-2} + A_1(1-P) f_{t-1} + A_2(1-P) f_{t-2} + \varepsilon_t \\ &\approx A_1 P f_{t-1} + A_2 P f_{t-2} + \varepsilon_t, \end{aligned} \tag{8}$$

where $1-P$ is the Hilbert space projection defined as $(1-P)f = f - Pf$ for all f in H. The approximation error terms $\big(A_k(1-P)f_{t-k}\big)$ for $k=1$ and 2 are asymptotically negligible under suitable regularity conditions, if we set $m \to \infty$ as $T \to \infty$ at an appropriate rate. The required conditions are not very stringent and they are expected to hold generally. Our empirical analysis is based on the approximate FAR in (8).

A.3 Isometric Mapping to Finite-Dimensional VAR

Define the mapping

$$\pi : f \mapsto (f) \equiv \begin{pmatrix} \langle v_1, f \rangle \\ \vdots \\ \langle v_m, f \rangle \end{pmatrix} \tag{9}$$

for any f in H, and

$$\pi : A \mapsto (A) \equiv \begin{pmatrix} \langle v_1, Av_1 \rangle & \cdots & \langle v_1, Av_m \rangle \\ \vdots & \vdots & \vdots \\ \langle v_m, Av_1 \rangle & \cdots & \langle v_m, Av_m \rangle \end{pmatrix} \tag{10}$$

for any linear operator A on H. Then we may represent the approximate FAR in (8) as

$$\lgroup f_t \rgroup \approx \lgroup A_1 \rgroup \lgroup f_{t-1} \rgroup + \lgroup A_2 \rgroup \lgroup f_{t-2} \rgroup + \lgroup \varepsilon_t \rgroup, \tag{11}$$

a conventional m-dimensional VAR, which is referred to as the approximate VAR of our FAR. Note that $\big(\lgroup f_t \rgroup\big)$ and $\big(\lgroup \varepsilon_t \rgroup\big)$ are m-dimensional time series and $\lgroup A_1 \rgroup$ and $\lgroup A_2 \rgroup$ are $m \times m$ matrices. The approximate VAR in (11) is readily derived from the approximate FAR in (8), since we have $\lgroup APf \rgroup = \lgroup A \rgroup \lgroup f \rgroup$ for any f in H and any operator A on H, and $\lgroup f + g \rgroup = \lgroup f \rgroup + \lgroup g \rgroup$ for all f and g in H. The approximate FAR in (8) is therefore equivalent to the approximate VAR in (11), which implies that the original FAR in equation (1) may be analyzed by the approximate VAR in (11) if we let $m \to \infty$ as $T \to \infty$ as mentioned earlier. Indeed, Chang et al. (2021) show that the use of the VAR in (11) is valid under mild conditions for the general structural analysis of the FAR in equation (1) relying on general sample and bootstrap asymptotic theories.

Although π's in (9) and (10) are defined for any f in H and for any linear operator A on H, we interpret them as their restricted versions on the linear subspace V spanned by the sub-basis $(v_i)_{i=1}^m$ whenever necessary. The restricted versions of π's are one-to-one so that their inverses exist and are well-defined. We may indeed easily show that

$$\pi^{-1}\big(\lgroup f \rgroup\big) = Pf \quad \text{and} \quad \pi^{-1}\big(\lgroup A \rgroup\big) = PAP.$$

Consequently, from the estimates $\widehat{\lgroup A_1 \rgroup}$ and $\widehat{\lgroup A_2 \rgroup}$ of the autoregressive coefficient matrices $\lgroup A_1 \rgroup$ and $\lgroup A_2 \rgroup$ and the fitted values $\big(\widehat{\lgroup \varepsilon_t \rgroup}\big)$ of the residuals $\big(\lgroup \varepsilon_t \rgroup\big)$ in (11), we may easily obtain the corresponding estimates $\widehat{A}_1$ and $\widehat{A}_2$ as linear operators on V and the fitted functional residuals $(\widehat{\varepsilon}_t)$ as a time series taking values in V.

This mathematical framework establishes the equivalence between the infinite-dimensional FAR and its finite-dimensional VAR approximation, enabling standard econometric methods to be applied to functional time series data.

B Choice of Basis

The VAR representation in equation (2) may be obtained using any orthonormal basis (v_i) of H. However, the effectiveness of the resulting approximation depends crucially on the choice of basis. As discussed in Section 2, we use the functional principal component (FPC) basis (v_i^*), which maximizes the explained variation in the data. This appendix compares the FPC basis to three alternative bases—moment, histogram, and quantile—documenting the superiority of FPCs in terms of both explained variation and estimation precision. For detailed theoretical justification of the FPC basis choice, including isometry properties and variance-bias decomposition, see Chang et al. (2021).

B.1 Alternative Basis Definitions

We consider three alternatives to the FPC basis:

Moment basis. The moment basis $(v_i)_{i=1}^m$ is obtained by the Gram-Schmidt orthogonalization procedure from the pre-basis defined as $u_i(r) = r^i$ for $i \geqslant 1$ over the interval $[p, q)$ with $p = -0.5$ and $q = 0.5$. We call (u_i) the moment basis, since

$$\langle u_i, f_t \rangle = \int r^i f_t(r) dr$$

and $(\langle u_i, f_t \rangle)$ represents the i-th moments of the EID given by the densities (f_t) for $i \geqslant 1$.

Histogram basis. The histogram basis $(v_i)_{i=1}^m$ is given by

$$v_i(r) = \frac{1}{\sqrt{q_i - p_i}} 1\{p_i \leqslant r < q_i\},$$

where $\big([p_i, q_i)\big)$ is a partition of the support $[p, q)$ of the densities (f_t) into $(m+1)$ sub-intervals of equal length. We let $p = -0.5$ and $q = 0.5$ and take only m indicators as a basis, ignoring the first sub-interval to avoid linear dependence.

Quantile basis. The quantile basis $(v_i)_{i=1}^m$ is defined similarly as indicators over a different partition $\big([p_i, q_i)\big)$, where the $(m+1)$ sub-intervals are obtained by defining (q_i) as the $i/(m+1)$-th sample quantiles of all observations for $i = 1, \ldots, m+1$. As with the histogram basis, we include only m indicators.

B.2 Empirical Comparison

Tables 4 and 5 compare the four bases using one-year ahead expected inflation distributions from January 1978 to December 2021. Table 4 reports the functional R-squared (FR^2), which measures the proportion of total variation explained by the first m basis functions. Table 5 reports $\mathrm{trace}(Q\Lambda Q)^+$, a measure asymptotically proportional to the variance of the autoregressive operator estimators $\widehat{A}_1$ and $\widehat{A}_2$.

The tables demonstrate the clear superiority of the FPC basis. For $m = 3$, the FPC basis captures over 95 percent of total EID variation while maintaining a variance measure of only 4.7. In sharp contrast, alternative bases perform poorly on both dimensions. The moment basis is especially ineffective, capturing only 1.7 percent of variation at $m = 3$ while producing a variance measure of 3,916—over 800 times larger than the FPC basis. The quantile basis performs somewhat better, explaining 55.7 percent of variation, but still

Table 4. *FR2 for Four Choices of Basis*

m	FPC Basis	Histogram Basis	Quantile Basis	Moment Basis
1	0.7419	0.0095	0.0322	0.0048
2	0.8829	0.0123	0.3707	0.0064
3	0.9513	0.0400	0.5573	0.0172
4	0.9676	0.1103	0.5611	0.0220
5	0.9785	0.0735	0.7181	0.0353
6	0.9861	0.2789	0.7309	0.0461
7	0.9898	0.1430	0.7580	0.0561
8	0.9921	0.3938	0.7731	0.0678
9	0.9936	0.2482	0.7737	0.0771
10	0.9946	0.4142	0.7849	0.0878

Notes: The FR^2 values are reported for four different choices of basis including the functional principal component (FPC) basis, histogram basis, quantile basis and moment basis. The FR^2 is expected to strictly increase as m gets large. However, this is not the case for the histogram basis, since it is defined differently for different values of m. The time series of one-year ahead expected inflation distributions from January 1978 to December 2021 are used to compute the reported FR^2 values.

Table 5. *trace$(Q\Lambda Q)^+$ for Four Choices of Basis ($\times 10^4$)*

m	FPC Basis	Histogram Basis	Quantile Basis	Moment Basis
1	0.613	16.084	7.990	69.714
2	1.475	2230.758	10.867	639.162
3	4.708	8905.797	27.642	3916.064
4	8.939	13743.106	52.248	11926.281
5	16.930	17148.747	57.780	16819.808
6	24.676	21877.429	209.583	21002.513
7	39.782	30720.005	161.306	27026.392
8	60.503	41978.159	487.056	35492.106
9	88.282	53705.816	666.356	42583.255
10	124.784	62145.588	1079.753	59256.713

Notes: The values of trace$(Q\Lambda Q)^+$ are asymptotically proportional to the variances of the autoregressive operator estimators in the Hilbert-Schmidt norm for four different choices of basis including the functional principal component (FPC) basis, histogram basis, quantile basis and moment basis. The time series of one-year ahead expected inflation distributions from January 1978 to December 2021 are used.

falls far short of the FPC basis and exhibits substantially higher estimation variance (27.6 versus 4.7).

These results confirm the findings presented in Section 2: the FPC basis is uniquely effective for representing EID dynamics, offering both superior explanatory power and substantially lower estimation variance compared to moment-based, histogram-based, or quantile-based alternatives. The choice of basis is therefore critically important for our functional methodology.

C Interpretation of Functional Principal Components

Here, we examine how each of the three FPC affects the first four standardized moments, namely, the mean, standard deviaion, skewness, and kurtosis, computed from the EID. Each panel in Figure 14 shows the change in these four moments with respect to loadings for each FPC.

Figure 14. *Range of variation of the mean, standard deviation, skewness and kurtosis due to the functional principal components*

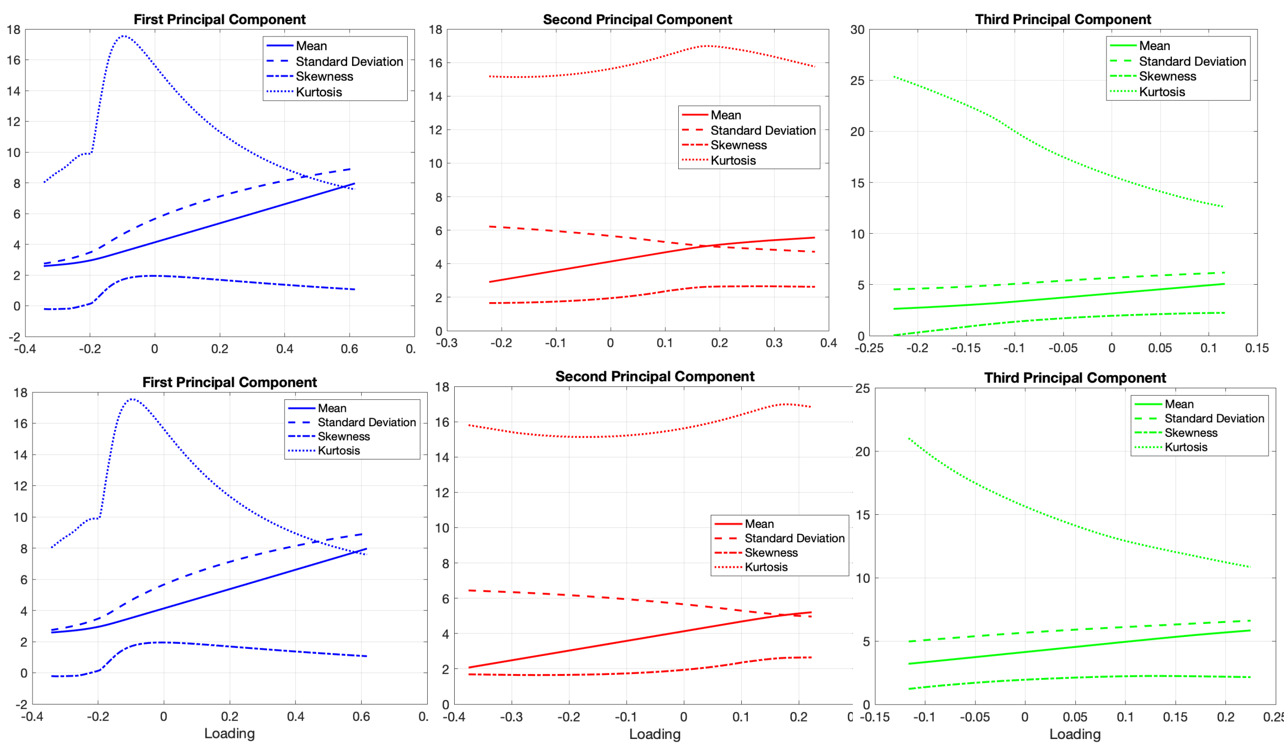

Notes: Each figure reports the results for each FPC. The blue/red/green colors signify the first/second/third FPCs. The ordered loadings are on the horizontal axis and the values of the four standardized moments are shown on the vertical axis.

Overall this exercise is the basis of our labeling of the three FPCs as disagreement, shift, and inflationary ambiguity.

(a) First FPC (blue): The second moment varies more than the first moment across all loadings. The variations in the third and fourth moments are implied by those in the first and the second. Therefore it makes sense to name the first FPC and call it as the disagreement component. While we observe that a significant variation in the fourth moment (kurtosis) with loadings, the movement in the second moment has a clear economic interpretation.

(b) Second FPC (red): It is the first moment (mean) that shows the strongest relationship with the loading. Therefore, we label the second functional component as the 'shift' component.

(c) Third FPC (green): The first three moments are pretty stable across all loadings, while the fourth moment shows a clear relationship with the loadings. As it is difficult to interpret the economic meaning of the fourth moment, we label this functional component as the 'ambiguity' component.

D Robustness Check: Measuring Responses of EID to External Shocks

When measuring the impulse response of expected inflation distributions (EID) to external shocks, we face two main methodological choices. First, government spending and personal income tax shocks are available only quarterly, while our EID data are monthly, creating a frequency mismatch. Second, we can measure EID responses using either a correlation-based method (as in the main text) or a proxy-VAR approach. This appendix describes how we handle the frequency mismatch under both methods and presents robustness checks comparing the resulting impulse responses.

D.1 Handling Mixed-Frequency External Shocks

We use two methods to measure EID responses to external shocks:

Correlation method: Links the functional shocks (e_{it}) identified from our FAR directly to the external shock x_t through correlations, as described in the main text.

Proxy-VAR method: Embeds the external shock in an augmented VAR with the EID loadings, $(x_t, (f_t))$, and identifies the shock recursively. This approach is a simple form of Mixed Autoregression (MAR), where the external shock enters as an aggregate variable alongside the functional variable (Chang et al., 2026).[16] Below we describe each method and how we adapt them for quarterly shocks.

Monthly shocks (no frequency mismatch). When the external shock is monthly—as with monetary policy and gasoline price shocks—both methods are implemented exactly

[16] We could also use local projections (Plagborg-Møller and Wolf, 2021), which would yield similar results to the proxy-VAR approach.

as described in the main text, since the monthly EID and monthly shock can be matched one-to-one.

Quarterly shocks (frequency mismatch). When the external shock is quarterly—as with government spending and personal income tax shocks—we must account for the mismatch with monthly EID data. We use the following approach:

Rather than aggregate monthly EID to quarterly frequency (which would discard two-thirds of the data), we retain the monthly FAR and construct the quarterly response as the cumulative effect across all three months of the quarter. A quarterly shock occurring in quarter s (spanning months t, $t+1$, $t+2$) can affect EID in any of those three months. We therefore measure the full quarterly response in the third month of the quarter by summing the contributions from all three months.

Correlation method with quarterly shocks. For a quarterly shock x_s, we correlate it with the monthly functional shocks in each of the three months of the quarter. Specifically, for each functional component $i = 1, \ldots, m$ and each month $j = 0, 1, 2$ within the quarter, we estimate:

$$e_{i,t+j} = \rho_{xi,j}\, x_s + \text{residuals},$$

where $\rho_{xi,j}$ measures how strongly the quarterly shock x loads onto functional component i in month j of the quarter (with $j = 0$ denoting the first month, $j = 1$ the second, and $j = 2$ the third). This notation extends the correlation ρ_{xi} from the main text (equation 6) by adding the monthly index j, since quarterly shocks can have different impacts in each month of the quarter.

The EID response in each month j within the quarter is then constructed as:

$$\phi_{x,j} = \sum_{i=1}^{m} \rho_{xi,j}\, b_i(j),$$

where $b_i(j)$ is the baseline EID response to functional shock i at horizon j. The cumulative quarterly response is the sum across the three months:

$$\Phi_x = \phi_{x,0} + \phi_{x,1} + \phi_{x,2} = \sum_{i=1}^{m} \left(\rho_{xi,0}\, b_i(0) + \rho_{xi,1}\, b_i(1) + \rho_{xi,2}\, b_i(2)\right),$$

as defined in the main text. This expression generalizes equation (6) to accommodate quarterly shocks by allowing the correlations to vary across the three months of the quarter and summing the resulting monthly responses.

Proxy-VAR method with quarterly shocks. The implementation is more straightforward. We construct a quarterly VAR using:

$$\begin{pmatrix} x_s \\ \llbracket f_s \rrbracket \end{pmatrix},$$

where s indexes quarters and $\llbracket f_s \rrbracket$ denotes the EID loadings from the third month of each quarter. The external shock x_s is ordered first and identified recursively. Since x_s is externally identified, contemporaneous feedback from EID to x_s is ruled out by construction, validating the recursive identification.

The EID response to the external shock is obtained from the impulse responses of $\llbracket f_s \rrbracket$ to the first structural shock, mapped back to the functional space using the isometric transformation described in the main text.

D.2 Comparison of Methods and Samples

Figures 15 through 20 present at-impact EID responses to the four external shocks under different methods and sample choices. Each figure displays responses at either the one-year or medium-run horizon. The figures compare responses estimated over the maximum available sample for each shock (row 1) versus the maximum common sample across all shocks (row 2).

For monthly shocks (monetary policy and gasoline prices), these two specifications are shown at quarterly frequency in Figures 15 through 18, where the quarterly version aggregates the three monthly observations within each quarter. Figures 19 and 20 additionally display these responses at their native monthly frequency.[17]

The results demonstrate that our main findings are robust across both methodological approaches and sample periods. The correlation-based and proxy-VAR methods yield qualitatively similar responses, and the patterns remain consistent whether estimated over the maximum available sample or restricted to the common sample period.

D.3 Responses at Impact

This subsection presents at-impact EID responses to the four external shocks under alternative methods and sample specifications. The purpose is to demonstrate the robustness of our main findings to different methodological and sampling choices.

Figures 15 through 20 display EID responses comparing the correlation-based and proxy-VAR methods across different sample periods and frequencies. The correlation method directly links functional shocks to external shocks through correlations, while

[17] We also estimated responses over the common sample period (1991Q1–2006Q4) and using purified shocks (orthogonalized with respect to other external shocks). These specifications yield results qualitatively similar to those reported here and are available upon request.

the proxy-VAR approach embeds the external shock in an augmented VAR with EID loadings—a simple form of Mixed Autoregression (MAR) where the external shock enters as an aggregate variable.

Figures 15 and 16 show responses using the correlation method for one-year ahead and medium-run horizons, respectively. Figures 17 and 18 present the corresponding proxy-VAR results. Each figure displays responses to all four shocks in a grid format: columns correspond to monetary policy, government spending, personal income tax, and gasoline price shocks (from left to right). Rows compare responses estimated over the maximum available sample for each shock (row 1) versus the maximum common sample across all shocks (row 2). All responses are reported at quarterly frequency.

For the two monthly shocks (monetary policy and gasoline prices), Figures 19 and 20 additionally show responses at monthly frequency for one-year ahead and medium-run horizons, respectively. These figures facilitate direct comparison between the correlation and proxy-VAR methods at the native monthly frequency of these shocks.

The figures reveal two key patterns. First, the correlation-based and proxy-VAR methods yield qualitatively similar EID responses across all four shocks, confirming that our findings are not driven by the specific identification approach. Second, responses remain consistent across the different sample specifications, indicating that our results are robust to the choice of sample period. These robustness checks support the reliability of the distributional responses documented in the main text.

E Bootstrapping Methods

We use bootstrap methods to construct confidence intervals for the impulse response functions (IRFs) of EID to external shocks. This appendix describes the bootstrap procedures

Figure 15. *One-year EID response estimated using the correlation method*

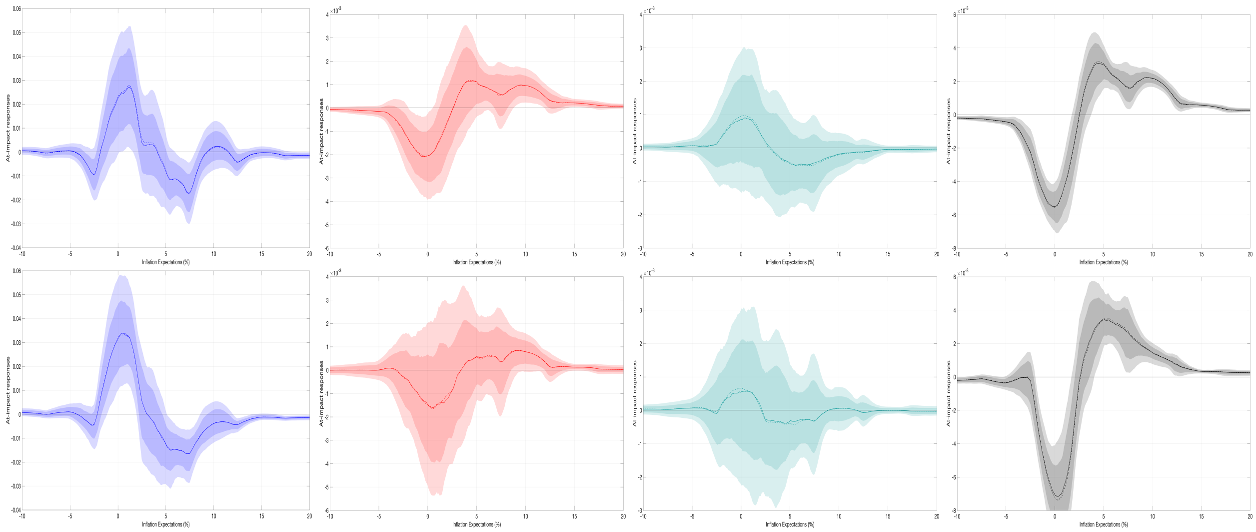

Notes: Each column corresponds to a different external shock: monetary policy, government spending, personal income tax, and gasoline price shocks (from left to right). All responses are reported at a quarterly frequency. The first row shows quarterly EID responses estimated using the maximum available sample for each shock, while the second row shows responses estimated over the maximum common sample across all shocks. Shaded areas represent 68 and 90 percent confidence bands constructed using bootstrap methods.

Figure 16. *Medium-run EID response estimated using the correlation method*

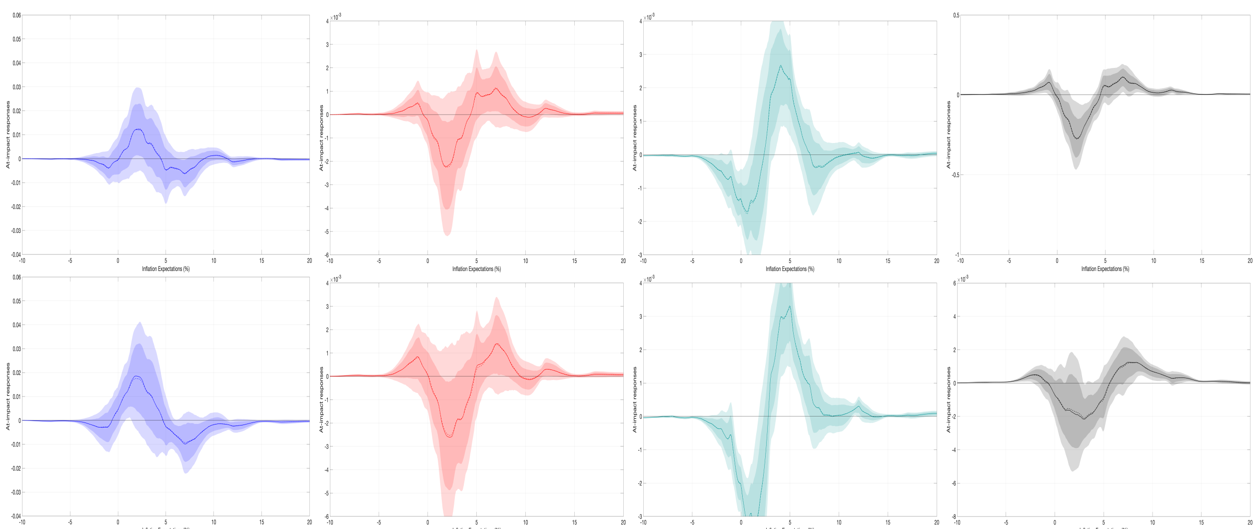

Notes: Each column corresponds to a different external shock: monetary policy, government spending, personal income tax, and gasoline price shocks (from left to right). All responses are reported at a quarterly frequency. The first row shows quarterly EID responses estimated using the maximum available sample for each shock, while the second row shows responses estimated over the maximum common sample across all shocks. Shaded areas represent 68 and 90 percent confidence bands constructed using bootstrap methods.

Figure 17. *One-year EID response estimated using the proxy-VAR method*

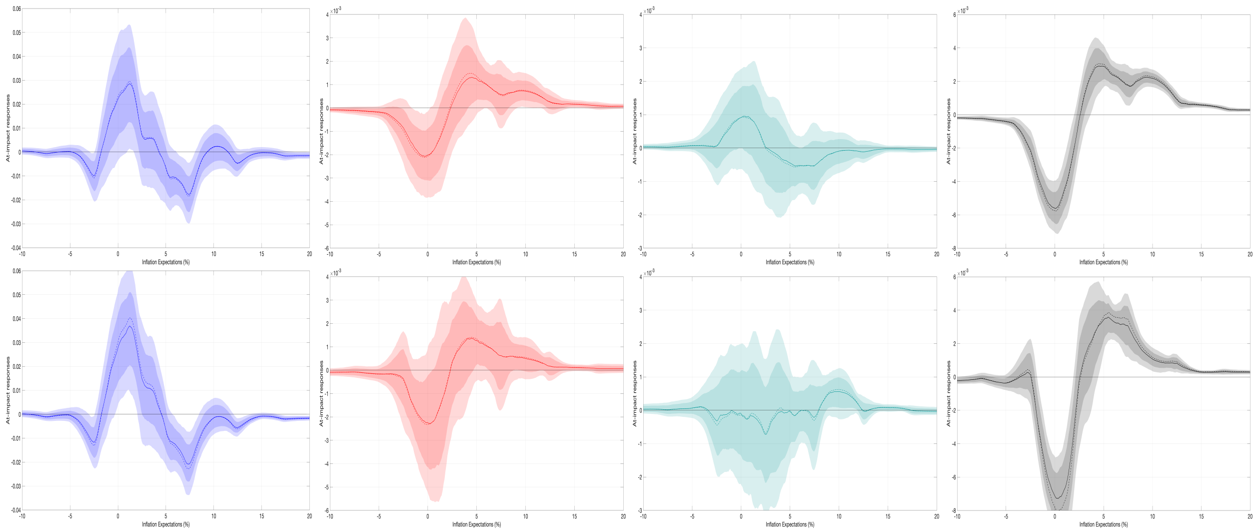

Notes: Each column corresponds to a different external shock: monetary policy, government spending, personal income tax, and gasoline price shocks (from left to right). The first row shows responses estimated using the maximum available sample, and the second row shows responses estimated over the maximum common sample across all shocks. Shaded areas represent 68 and 90 percent confidence bands constructed using bootstrap methods.

Figure 18. *Medium-run EID response estimated using the proxy-VAR method*

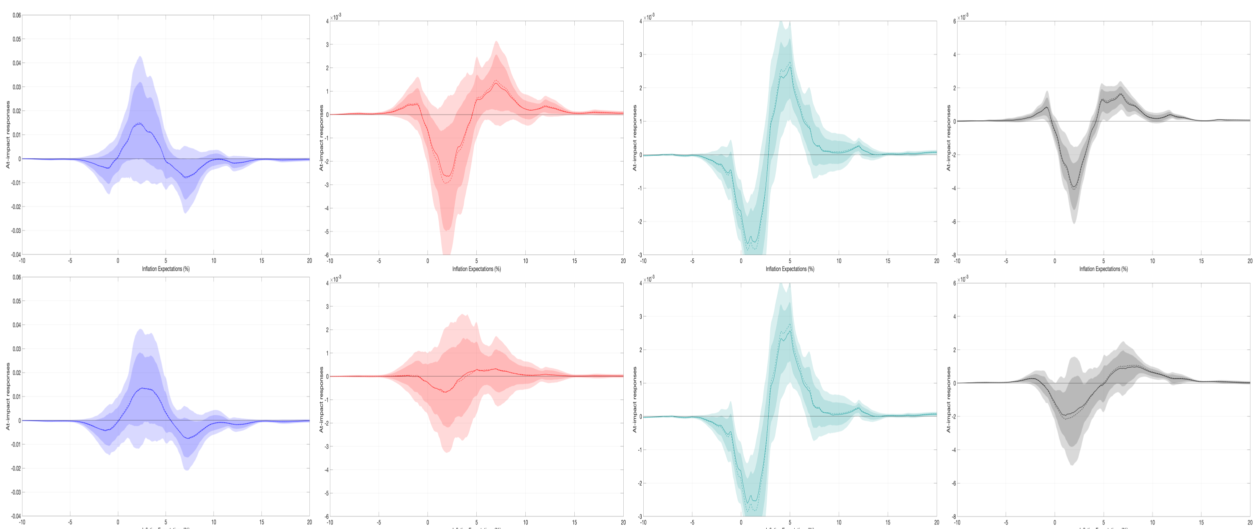

Notes: Each column corresponds to a different external shock:monetary policy, government spending, personal income tax, and gasoline price shocks (from left to right). The first row shows responses estimated using the maximum available sample, and the second row shows responses estimated over the maximum common sample across all shocks. Shaded areas represent 68 and 90 percent confidence bands constructed using bootstrap methods.

Figure 19. *One-year EID responses to monthly external shocks*

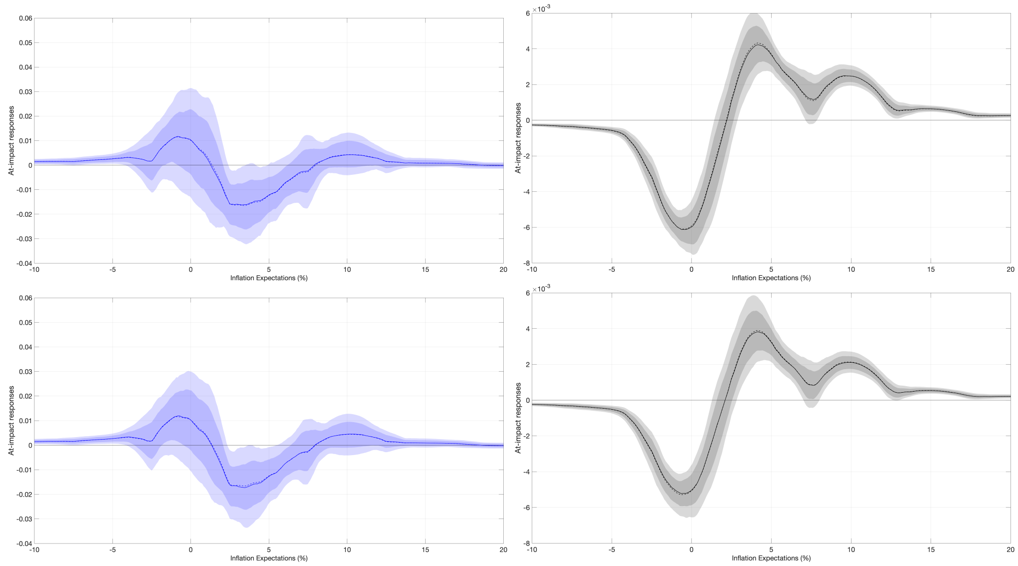

Notes: Columns correspond to monetary policy (left) and gasoline price shocks (right). Responses are reported at a monthly frequency. The top row displays correlation-based EID responses, and the bottom row displays proxy-VAR-based responses. Shaded areas denote 68 and 90 percent bootstrap confidence bands.

Figure 20. *Medium-run EID responses to monthly external shocks*

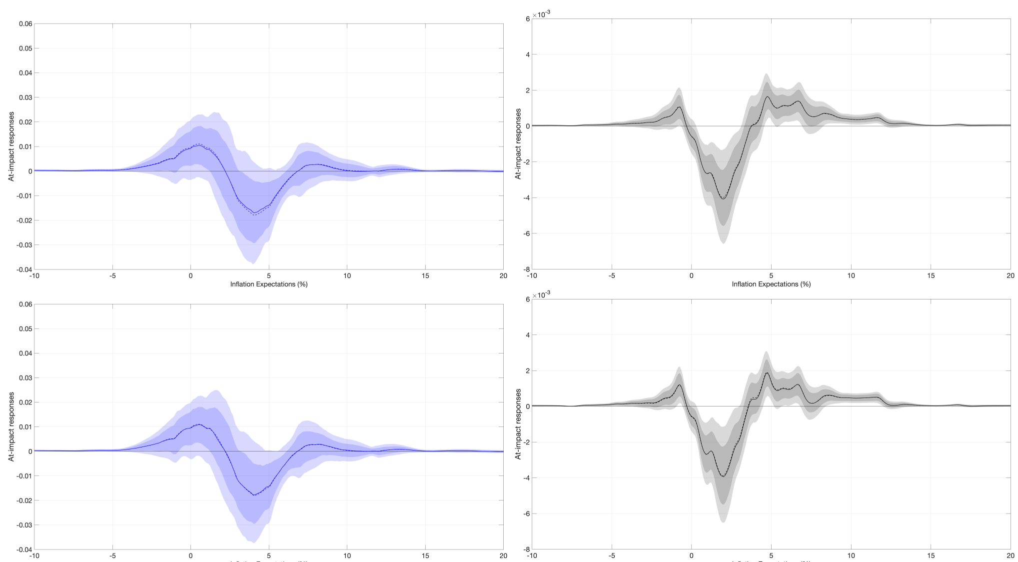

Notes: Columns correspond to monetary policy (left) and gasoline price shocks (right). Responses are reported at a monthly frequency. The top row displays correlation-based EID responses, and the bottom row displays proxy-VAR-based responses. Shaded areas denote 68 and 90 percent bootstrap confidence bands.

for both the correlation-based method (our primary approach) and the proxy-VAR method (used for robustness).

E.1 Bootstrap for Correlation-Based Method

As described in the main text, the EID response to an external shock x_t is constructed as a linear combination of baseline responses to the three functional shocks (e_{1t}, e_{2t}, e_{3t}), weighted by correlations:

$$\phi_x = \rho_{x1} b_1 + \rho_{x2} b_2 + \rho_{x3} b_3,$$

where $\rho_{xi} = \text{corr}(e_{it}, x_t)$ and b_i denotes the baseline response of EID to functional shock i.

The bootstrap procedure accounts for sampling uncertainty in the estimated correlations ρ_{xi} through synchronized resampling of the external shock and functional shocks:

(a) **Synchronized resampling.** Randomly draw with replacement a bootstrap sample of size T from the original sequence of pairs (x_t, e_t) for $t = 1, \ldots, T$, where T denotes the length of the common sample period for which both the external shock and functional shocks are available. This synchronized resampling preserves the joint distribution between the external shock and functional shocks.

(b) **Compute bootstrap correlations.** Calculate the correlation vector between the resampled external shock and functional shocks:

$$\rho_x^* = (\rho_{x1}^*, \rho_{x2}^*, \rho_{x3}^*), \quad \rho_{xi}^* = \text{corr}(e_{it}^*, x_t^*).$$

(c) **Construct bootstrap IRF.** Using the bootstrap correlations as weights, form the bootstrap IRF by combining the original baseline responses:

$$\phi_x^* = \rho_{x1}^* b_1 + \rho_{x2}^* b_2 + \rho_{x3}^* b_3.$$

We repeat steps 1–3 for $B = 1000$ bootstrap replications. Confidence bands are constructed as quantiles of the resulting bootstrap distribution $\{\phi_x^*\}_{b=1}^B$.

This procedure preserves the joint distribution of the external shock and functional shocks while allowing the correlations to vary across bootstrap samples, thereby capturing sampling uncertainty in the estimated weights.

E.2 Bootstrap for Proxy-VAR Method

As a robustness check, we also compute distributional impulse responses using a proxy-VAR approach following Chang et al. (2026). This augments the approximate VAR in equation (2) by including the external shock as the first variable.

Let the stacked vector be

$$y_t = \begin{pmatrix} x_t \\ (f_t) \end{pmatrix},$$

where x_t is the external shock (standardized to unit variance) and $(f_t) \in \mathbb{R}^m$ denotes the vector of FPC scores from equation (2). For clarity, we describe the bootstrap for a VAR(1), though the procedure extends immediately to higher orders.

Bootstrap Algorithm

(a) **Estimate the VAR.** Fit

$$y_t = \hat{\Phi} y_{t-1} + \hat{u}_t$$

and store the residuals $\{\hat{u}_t\}_{t=1}^T$.

(b) **Residual resampling.** Draw with replacement T bootstrap residuals $\{u_t^*\}_{t=1}^T$ from $\{\hat{u}_t\}$ and recenter them to enforce zero mean.

(c) **Recursive simulation.** Generate a bootstrap sample $\{y_t^*\}_{t=1}^T$ via

$$y_t^* = \hat{\Phi} y_{t-1}^* + u_t^*,$$

initialized at $y_0^* = y_0$.

(d) **Re-estimate and construct IRFs.** Using $\{y_t^*\}$, re-estimate the VAR and compute structural impulse responses as in the main text, identifying the shock in x_t recursively. The functional IRFs are obtained by reconstructing the EID through the FPC basis:

$$\Theta_s^*(h, \cdot) = \sum_{k=1}^{m} \mathrm{IRF}_{k,s}^*(h)\, v_k^*(\cdot),$$

where $\mathrm{IRF}_{k,s}^*(h)$ is the bootstrap impulse response of the k-th FPC score.

Repeating steps 2–4 for $B = 1000$ replications, we construct pointwise confidence bands using empirical quantiles of the bootstrap distribution.

This residual-based bootstrap preserves the dependence structure implied by the estimated VAR, ensuring valid inference under covariance-stationary innovations. The procedure follows standard VAR bootstrap methods (Kilian and Lütkepohl, 2017) adapted to the functional setting.

F Relation to Alternative Functional Approaches

This appendix positions our empirical framework relative to other functional approaches used to analyze time-varying distributions in macroeconomics and applied econometrics.

Our methodology follows the functional autoregression (FAR) framework developed by Chang et al. (2021) (CPP), which provides asymptotic theory for functional time series and a convenient bridge to finite-dimensional VAR representations. This framework has been applied to a variety of distributional objects, including income distributions (Chang et al., 2025), climate temperature distributions (Chang et al., 2024b), and stock return distributions (Bjørnland et al., 2025). The present paper extends this approach to household inflation expectations, representing a novel application to survey-based distributional data.

Within the broader class of functional methods, our CPP-based approach differs from functional state-space frameworks that represent distributions using pre-specified sieve bases embedded in linear-Gaussian state-space systems. Such approaches are desirable when cross-sectional samples are small, distributions are censored, or when explicit modeling of density estimation uncertainty through measurement equations is required, as in the functional state-space framework of Chang et al. (2024a).

In contrast, we project expected inflation distributions (EIDs) onto a data-driven functional principal component (FPC) basis and approximate the FAR using a low-dimensional VAR in FPC loadings. This choice reflects key features of our application. The large monthly cross-sections of the Michigan Survey allow EIDs to be estimated precisely, while the observed distributional dynamics—such as abrupt shifts around salient shocks, episodic tail thickening, and time-varying disagreement—favor adaptive, data-driven bases over fixed parametric forms. As documented in Appendix B, a small number of FPCs capture over 95 percent of distributional variation with substantially lower estimation variance than moment-based or percentile-based representations.

These approaches are complementary. State-space methods with sieve bases are advantageous when the research objective is to track pre-specified distributional features or groups (e.g., particular percentiles or demographic segments). By contrast, the CPP framework with a data-driven FPC basis is well suited for uncovering the dominant patterns of distributional evolution implied by the data itself. For research questions focused on how the overall shape of the distribution responds to economic shocks—rather than on pre-defined subgroups—the FPC-based FAR approach offers gains in parsimony and variance efficiency. Given the characteristics of the Michigan Survey data and our focus on shock-driven distributional responses, this data-driven framework is well suited to our application.