The Effects of Economic Shocks on Heterogeneous Inflation Expectations*

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

This version: June 29, 2023

Abstract

We use a functional approach to study the effects of economic shocks on the distribution of U.S. household inflation expectations. Contractionary monetary policy shocks have no effect on average short- and medium-run inflation expectations. The impact on the average, however, masks distributional effects, as more households expect negative inflation and fewer expect moderate inflation. Government spending shocks raise average short-run inflation expectations by increasing the proportion of households with high inflation expectations. Personal income tax shocks raise only medium-run inflation expectations. Gasoline price shocks increase average inflation expectations for both short-run and medium-run, but with different distributional consequences.

JEL *classification*: E52, E61, E62, E65, H11, H30

Keywords: Inflation expectations, household survey, functional autoregression, functional impulse responses, transmission of economic shocks

*We would like to thank Chaewon Baek, 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 and the KAEA seminar series. 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 (gomezrj@bccr.fi.cr). The views expressed herein are those of the authors and do not necessarily represent the views of the Central Bank of Costa Rica.

¶International Monetary Fund (IMF) (ghong@imf.org). The views expressed herein are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

1 Introduction

> *"Perhaps more importantly, we need to know more about the manner in which inflation expectations are formed and how monetary policy influences them."* (Janet Yellen, 2016)

The seminal works of Friedman (1968) and Phelps (1967) have placed inflation expectations at the heart of macroeconomics and monetary policy. A well-established empirical and theoretical literature underscores the significance of inflation expectations in economic decision-making. In this extensive literature, researchers tend to limit their analysis to sample statistics of survey responses, such as mean, median or standard deviation, and consider *average* inflation expectations at a given time as *the* expected inflation.[1] At the same time, recent studies document significant heterogeneity in inflation expectations across households associated with demographic factors or consumption patterns, emphasizing the distributional aspect of inflation expectations (D'Acunto et al., 2023; Andre et al., 2022).

To demonstrate the limitations of average inflation expectations in fully representing heterogeneous inflation expectations, we plot the evolution of cross-sectional distributions of one-year ahead and medium-run (5 to 10 years ahead) inflation expectations for U.S. households from March 2018 to December 2021 using the University of Michigan's *Survey of Consumers* (Figure 1).[2] The comparison spans across three time periods: pre-pandemic months (from March 2018 to February 2020), pandemic months before inflation increases (from March 2020 to March 2021), and months with higher inflation (from April 2021 to December 2021). There is a visible change in the shape of the distributions of one-year ahead inflation expectations from the pre-pandemic period (blue) to the early pandemic months (red) – the curve flattened and the concentration of survey responses at the modal

[1] Recent work by Reis (2022), which examines different moments and tail behaviors of inflation expectations in several countries, is a notable exception. Here again, higher moments were explored rather than the entire distribution.

[2] In this paper, we use one-year ahead inflation expectations and short-run inflation expectations interchangeably.

value declined sharply. Despite such clear distributional changes, however, the average inflation expectations increased only marginally from 2.7 percent (pre-pandemic period) to 2.8 percent (early pandemic months). A similar observation can be made about medium-run inflation expectations where the change in the average inflation expectations from 2.4 percent to 2.5 percent does not correspond to more pronounced changes in the shape of the distributions from the pre-pandemic period to early pandemic years. These findings cast doubt on using the average inflation expectations to represent the entire distribution of inflation expectations.

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

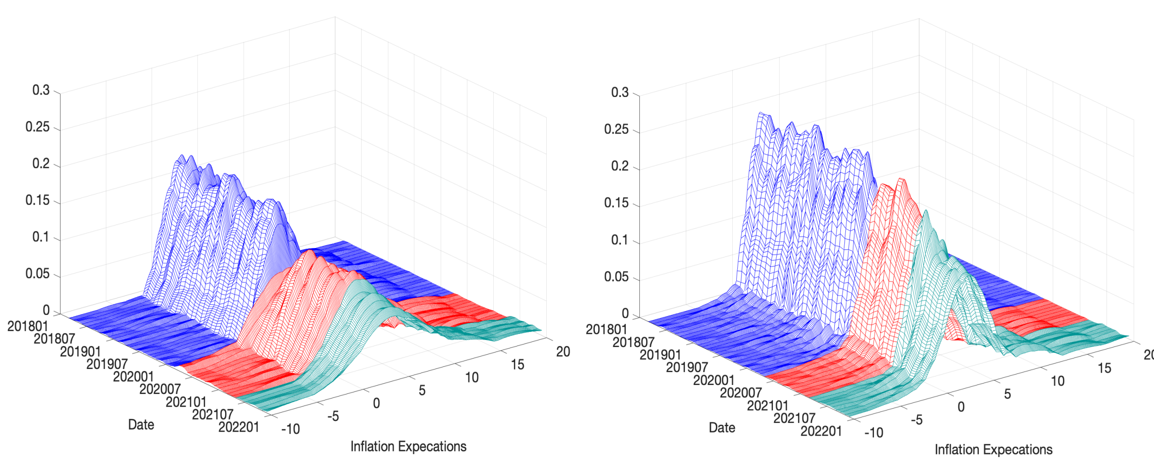

Notes: The distributions of the monthly U.S. household inflation expectations are shown for three periods: from March 2018 to February 2020 (blue), from March 2020 to March 2021 (red), and from April 2021 to December 2021 (green). The left (right) panel plots one-year ahead (medium-run) inflation expectations.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

In this paper, we adopt a functional approach to analyze U.S. household inflation expectations by broadening the scope of interest from specific sample statistics, such as the average, to the entire distribution of inflation expectations. Specifically, we examine how economic shocks affect the distribution of inflation expectations, focusing on four economic shocks: monetary policy, government spending, personal income tax, and gasoline prices. To examine the effects of one-dimensional economic shocks on the infinitedimensional distributions of households' inflation expectations, we use a methodology developed by Chang et al. (2021) which approximates functional autoregression (FAR) to a conventional vector autoregression (VAR) with an adequate representation of functions as finite-dimensional vectors. To implement this methodology, we first construct a functional time series of density functions, with each density function summarizing the monthly cross-section of households' inflation expectations from the University of Michigan's *Survey of Consumers*. We refer to these estimated density functions as 'EID', or the expected inflation distribution. Then, separately for one-year ahead and medium-run inflation expectations, we apply the functional principal component analysis (FPCA) to the functional time series of EIDs and identify three distinctive functional shocks driving the EID dynamics. Finally, we link these functional shocks to the four economic shocks to quantify the EID's impulse responses to each of these individual economic shocks.

We find that economic shocks studied in the paper affect the distributions of household inflation expectations differently for one-year ahead and medium-run inflation expectations, as we will explain below. To complement the distributional analysis, we analyze how economic shocks affect household inflation expectations summarized by commonly-used statistics in the literature, such as the average and standard deviations, and less commonly-used statistics, such as skewness, a measure of asymmetry of a distribution.[3] Furthermore, we examine how economic shocks affect the distribution of household inflation expectations by looking at the changes in probabilities of household inflation expectations before and after each shock, broken down by decile. This exercise reveals that changes (or lack thereof) in the selected sample statistics conceal significant distributional changes in inflation expectations caused by economic shocks. To our knowledge, no studies have looked at

[3]Skewness provides useful information about the distributional changes: a decline (an increase) in skewness suggests the shift of the distribution to the left (right) in an asymmetric way. In the context of our paper, a decline (an increase) in skewness implies a higher share of households with low (high) levels of inflation expectations.

the different effects of economic shocks on household inflation expectations, as a funciton of the initial level of household inflation expectations prior to economic shocks.

Using monetary policy shocks developed by Miranda-Agrippino and Ricco (2021), we find that contractionary monetary policy shocks have no effect on the average one-year ahead and medium-run inflation expectations. Nonetheless, monetary policy shocks affect the distribution of household inflation expectations. Following contractionary monetary policy shocks, the share of households with negative inflation expectations (below -0.7 percent) increases, while the share of households with moderate levels of pre-shock inflation expectations, between about 1.6 and 7 percent, decreases. A similar pattern emerges for medium-run inflation expectations, with a decrease in the share of households with pre-shock inflation expectations ranging from 3 to about 6 percent. For both forecast horizons, changes in the distribution are not accompanied by changes in the standard deviation (or dispersion) of inflation expectations. Skewness declines for one-year ahead inflation expectations, consistent with the increase in the share of households with negative inflation expectations.

Government spending shocks increase the average one-year ahead inflation expectations, consistent with the finding in Andre et al. (2022), albeit only weakly for medium-run inflation expectations. For both forecast horizons, we do not observe a significant increase in disagreements among households, measured by standard deviation. Importantly, unlike monetary policy shocks, government spending shocks affect households with high levels of inflation expectations, increasing the frequency of households with high pre-shock inflation expectations above 7 percent. Aside from an increase in the share of households with inflation expectations above about 9 percent, government spending shocks have no significant effects on medium-run inflation expectations. On the other hand, personal income tax shocks increase only the average medium-run inflation expectations, with no significant distributional or average effects on the short-run inflation expectations. This finding is consistent with a recent finding that shows that consumers tend to associate negative economic news with an inflationary shock (Coibion et al., 2020), as consumers may interpret personal income tax increases as having a negative impact on their personal finances, though other channels may be at work.

Finally, a surprise hike in gasoline prices increases the average inflation expectations for both horizons of inflation expectations, as found in the literature (Coibion and Gorodnichenko, 2015). However, the distributional implications of short-run and medium-run inflation expectations differ significantly. While gasoline price shocks reduce disagreements among households about short-run inflation expectations, medium-run inflation expectations show more dispersion. Households' differing views on the future path of gasoline prices could be one plausible explanation for this finding. The skewness for one-year ahead inflation expectations increases, mirroring the rise in the share of households expecting higher inflation in one year. The skewness of medium-run inflation expectations, on the other hand, decreases as the share of households with low inflation expectations compared to pre-shock levels grows.

To the extent that the existing studies focus on average inflation expectations, the only direct comparison we can make vis-à-vis the existing literature is the results on average inflation expectations. The limited impact of monetary and government spending shocks on medium-run inflation expectations is in line with the findings in Coibion et al. (2021) and Coibion et al. (2022), while our finding that a surprise increase in personal income tax increases the average medium-run inflation expectations is novel.

The contribution of our work to the literature is three-fold. First, it contributes to the rich empirical literature using survey data on inflation expectations to understand the formation of inflation expectations (see Coibion et al. (2018) for a recent survey), by broadening the scope of interest from sample statistics to the entire distribution of household inflation expectations. Second, the paper contributes to the literature on the impact of economic shocks on household inflation expectations for four different commonly-used shocks in the literature. Households are shown to be inattentive to monetary policy shocks when inflation expectations are well-anchored (Coibion et al., 2020), while others find that monetary policy shocks increase the disagreement due to information rigidities (Grigoli et al., 2020). Our work shows that monetary policy shocks affect household inflation expectations differently for households, depending on their initial level of inflation expectations before the shock. The transmission of fiscal policy shocks to inflation expectations is less studied, except to understand how the fear of fiscal dominance potentially undermines a central bank's independence in the context of emerging economies (Favero and Giavazzi, 2004). One notable exception is the work by Coibion et al. (2021), showing that household inflation expectations are affected by the news on public debt, but not on fiscal deficits. In this paper, we use widely-used fiscal policy shocks, namely, government spending and personal income tax shocks, and show that these fiscal shocks affect household inflation expectations in a different manner with varying distributional consequences. To the rich literature on the effects of gasoline price shocks on household inflation expectations, we contribute by highlighting the distributioanl effects, with varying effects on short-run and medium-run inflation expectations. Finally, we analyze the dynamics of the entire distribution of inflation expectations using a functional autoregressive model which can be well approximated by a conventional VAR as shown in Chang et al. (2021). Variations of this functional approach are used in Hu et al. (2016), Chang et al. (2020b) and Chang et al. (2020a).[4]

Our paper is organized as follows. Section 2 motivates our approach in understanding the entire distribution of inflation expectations. Section 3 describes the econometric methodology which approximates a functional autoregression (FAR) model to a vector autoregression (VAR) model. Section 4 describes the steps to implement our functional

[4]The framework and methodology in the paper are indeed widely applicable to study various functional data in many different areas of economics. For other related approaches and issues in analyzing economic models with functional variables, see Chang et al. (2016), Inoue and Rossi (2019), and Bjørnland et al. (2023).

approach using the *Survey of Consumers* data. Section 5 presents the findings on the distributional effects of economic shocks on EID and possible interpretations of our findings, and Section 6 concludes the paper. Appendices describe the econometric methods used for our empirical analyses and provide supplemental materials motivating our functional approach.

2 Distributions of Inflation Expectations (EIDs)

In this section, we highlight the importance of looking at heterogeneous households' inflation expectations.

Figure 2 shows the time-series variation of the households' inflation expectations in our sample. The top panels plot one-year inflation expectations, and the bottom panels show medium-run inflation expectations. In the left-most column, we select six months with the highest values of the average monthly inflation expectations and six months with the lowest values within our sample. We observe that the months with the highest average inflation expectations were in 1979 and 1980, years preceding the Volcker disinflation. The months exemplified by low inflation expectations, for both one-year ahead and medium-run, correspond to the periods of low inflation, stretching from the early 2000s until the pandemic. In the center column, we select six months with the highest values of standard deviations and the other six months with the lowest values. The months chosen with high standard deviations tend to coincide with the months chosen for high average inflation expectations. The right-most column plots the average distributions of inflation expectations as dark solid lines from January 2018 to December 2021 split into three periods as in Figure 1, where the colors also match the shaded regions of Figure 1. For one-year ahead inflation expectations, the distributions from early pandemic months have fatter tails compared to those from pre-pandemic months. The distributions of the medium-run inflation expectations also shifted to the right over time, albeit more moderately.

Figure 2. *Selected Observations of EIDs*

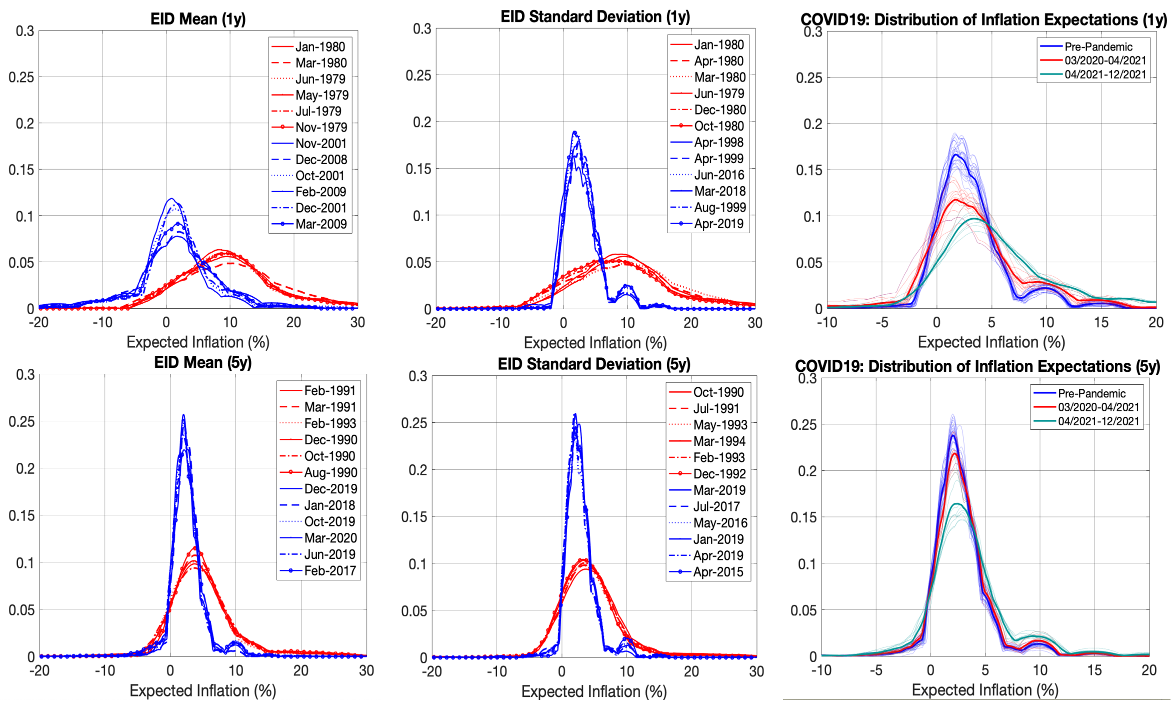

Notes: The top (bottom) panel represents one-year ahead (medium-run) household inflation expectations. The left column plots the densities of six observations with the highest monthly average vaules (red) and six observations with the lowest values (blue) from January 1978 to December 2021. The center column shows the densities of six observations with the highest (red) and and lowest (blue) standard deviations for the same sample period. The right column shows EIDs for each of the three different sub-periods mentioned in Figure 1, with the darker line showing the average EID for each sub-sample.

To underline the importance of evaluating the entire distribution of inflation expectations, we further examine the correlations between different features of the distributions of inflation expectations and key macroeconomic variables. The objective is to confirm whether or not there are specific sample moments that exhibit a clear and meaningful relationship with economic indicators. Table 1 displays the regression results of three key economic indicators, namely, real GDP growth, inflation, and the change in the unemployment rate, on different distributional aspects of one-year ahead inflation expectations. These aspects include the mean, standard deviation, skewness, and kurtosis of households'

Table 1. *Correlations between Selected Macroeconomic Indicators and Sample Moments of One-year Ahead Inflation Expectations*

	GDP growth	Inflation	Δ Unemployment
Intercept	2.12 (0.00)	−4.13 (0.00)	−2.15 (0.00)
Mean		1.54 (0.00)	0.41 (0.00)
Standard Deviation	1.04 (0.00)		−0.28 (0.00)
Skewness		−0.20 (0.00)	
Kurtosis		0.03 (0.00)	
Frequency at 0%		6.23 (0.00)	4.77 (0.00)
Frequency below 0%	−57.72 (0.00)		33.41 (0.00)
R^2	0.502	0.828	0.359

Notes: The table reports the coefficients of separate regressions of *real GDP growth*, *CPI inflation*, and changes in *unemployment* on different aspects of the distribution of one-year ahead U.S. household inflation expectations from the University of Michigan's *Survey of Consumers* from 1983Q1 to 2021Q4. Standard errors are reported in parentheses.

one-year ahead inflation expectations for each month. In addition, we look at the share of households who expect future inflation to be at 0 percent (or "Frequency at 0 percent") and the share of households who expect future inflation to be below zero (deflation expectations). The results reported here are chosen among combinations of the moments that result in the largest R^2 for each regression. The chosen macroeconomic indicators show varying correlations with different features of the distribution of inflation expectations. Notably, the average inflation expectation is not always selected in the group of EID properties explaining the largest portion of variations in macroeconomic variables, as is the case with real GDP growth.

3 Econometric Methodology

Our analysis of EID relies on a functional autoregression (FAR). In this section, we present a theoretical framework on how we formulate and implement the FAR to study the dynamics of EID. In what follows, we denote by (f_t) the density representing the EID, which we formally view as a functional time series taking values in the Hilbert space H of square-integrable functions defined in R generated 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, which will be defined more precisely below. We use the second order FAR, as suggested by both of the commonly used information criteria, AIC and BIC, for our data.

We begin with some basic concepts of the Hilbert space used in our model. The functional time series (f_t) represents a time series of random elements taking values in 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 the norm, we also need to 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 $H \equiv R^n$, 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

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

for all u and v in H. In particular, for the functional error (ε_t) assumed to be a white noise in (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 $t \geqslant 1$.

To study the dynamics of EID, we use the econometric methodology recently developed by Chang et al. (2021), which will be referred to as CPP for short. Other existing methodologies may also be used to estimate our FAR in (1). None of them, however, provides a framework to identify shocks and subsequently analyze the responses of EID to them, which is indeed the main exercise of our empirical study. Besides, as of now, only the methodology in CPP allows us to bootstrap the confidence bands for the responses of EID consistently. Below the methodology in CPP will be introduced concisely, but only briefly just to make our paper self-contained. The interested reader is referred to CPP for more details.

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

$$f \approx \sum_{i=1}^{m} \langle v_i, f \rangle v_i, \tag{2}$$

where m is the truncation number which should be chosen appropriately. We let V be the subspace of H spanned by a sub-basis $(v_i)_{i=1}^m$, and denote by P the Hilbert space projection on the m-dimensional subspace V. Then the approximation in (2) may simply be regarded as a projection of f on V yielding $Pf = \sum_{i=1}^m \langle v_i, f \rangle v_i$. 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 (1) can be represented by a VAR upon this approximation, as will be explained in what follows.

We approximate the FAR in (1) as

$$\begin{aligned} f_t &= A_1Pf_{t-1} + A_2Pf_{t-2} + A_1(1-P)f_{t-1} + A_2(1-P)f_{t-2} + \varepsilon_t \\ &\approx A_1Pf_{t-1} + A_2Pf_{t-2} + \varepsilon_t, \end{aligned} \tag{3}$$

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 (3). Subsequently, we will explain how we may represent this approximate FAR as a finite-dimensional VAR.

Define a mapping

$$\pi : f \mapsto (f) \equiv \begin{pmatrix} \langle v_1, f \rangle \\ \vdots \\ \langle v_m, f \rangle \end{pmatrix} \tag{4}$$

for any f in H, and

$$\pi : A \mapsto (A) \equiv \begin{pmatrix} \langle v_1, Av_1 \rangle & \cdots & \langle v_1, Av_m \rangle \\ \vdots & \vdots & \vdots \\ \langle v_m, Av_1 \rangle & \cdots & \langle v_m, Av_m \rangle \end{pmatrix} \tag{5}$$

for any linear operator A on H. Then we may represent the approximate FAR in (3) as

$$(f_t) \approx (A_1)(f_{t-1}) + (A_2)(f_{t-2}) + (\varepsilon_t), \tag{6}$$

a conventional m-dimensional VAR, which is referred to as the approximate VAR of our FAR. Note that $(\lbag f_t \rbag)$ and $(\lbag \varepsilon_t \rbag)$ are m-dimensional time series and $\lbag A_1 \rbag$ and $\lbag A_2 \rbag$ are $m \times m$ matrices. The approximate VAR in (6) is readily derived from the approximate FAR in (3), since we have $\lbag APf \rbag = \lbag A \rbag \lbag f \rbag$ for any f in H and any operator A on H, and $\lbag f + g \rbag = \lbag f \rbag + \lbag g \rbag$ for all f and g in H. The approximate FAR in (3) is therefore equivalent to the approximate VAR in (6), which implies that the original FAR in (1) may be analyzed by the approximate VAR in (6) if we let $m \to \infty$ as $T \to \infty$ as mentioned earlier. Indeed, CPP shows that the use of the VAR in (6) is valid under mild conditions for the general structural analysis of the FAR in (1) relying on the general sample and bootstrap asymptotic theories.

Although π's in (4) and (5) 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

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

Consequently, from the estimate $\widehat{\lbag A_1 \rbag}$ and $\widehat{\lbag A_2 \rbag}$ of the autoregressive coefficient matrices $\lbag A_1 \rbag$ and $\lbag A_2 \rbag$ and the fitted values $\big(\widehat{\lbag \varepsilon_t \rbag}\big)$ of the residuals $\big(\lbag \varepsilon_t \rbag\big)$ in (6), 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.

The VAR representation in (6) may be obtained for any choice of an orthonormal basis (v_i) of H. The effectiveness of the resulting approximation, however, depends crucially on the choice of basis. Following Bosq (2000), Ramsay and Silverman (2005), Hall and Horowitz (2007) and Park and Qian (2012), among others, we use the functional principal component basis (v_i^*). For $i = 1, \ldots, T$, we define v_i^* as the eigenfunction of the sample variance operator of EID given by

$$\Gamma = \frac{1}{T}\sum_{t=1}^{T}(f_t \otimes f_t) \tag{7}$$

associated with the i-th largest eigenvalue. The functional principal component basis (v_i^*) is known to most effectively approximate the temporal variations of functional time series, although other bases can also be used. In Appendix A, we demonstrate that the use of other bases such as the moment and quantile bases is much less effective. They not only explain much less EID variations over time but also yield the estimators of autoregressive operators A_1 and A_2 in (1) with unacceptably large variances. See Tables 4 and 5 in Appendix A for a comparison of functional R-squared and variances of $\hat{A}_1$ and $\hat{A}_2$ based on four different choices of basis, including FPC basis, which is our choice, histogram basis, quantile basis and moment basis, computed using the one-year ahead EID.

4 Implementation of Functional Approach using EIDs

In this section, we describe the implementation of the methodology presented in Section 3 to analyze the effects of economic shocks on the distribution of the U.S. households' inflation expectations from the University of Michigan's *Survey of Consumers*. There are several steps involved. First, based on the monthly responses from the *Survey of Consumers*, we estimate the density function representing the underlying distribution of inflation expectations each month, and use this monthly time series of density functions as our functional data. Then, we apply the functional principal component analysis described in Section 3 to extract the basis to obtain the optimal finite-dimensional representation of expected inflation distributions. As a next step, by applying a recursive identification to the corresponding approximate vector autoregression (VAR), we obtain *functional shocks* driving the expected inflation distributions, and their impacts on expected inflation distributions. We call the responses of the functional variable EID to these functional shocks as *baseline EID responses*. Finally, we compute the correlations between these functional shocks and an external economic shock and use them as weights to combine the baseline EID responses. We interpret them as the functional response of the EID to the economic shocks. Each of these steps will be explained in greater detail below.

4.1 Constructing EID Functional Data from Household Survey

Our reference sample is monthly data from January 1983 for one-year ahead (from January 1991 for medium-run inflation expectations) to December 2021 from the University of Michigan's *Survey of Consumers*. 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 summarize the monthly survey responses to a density function using a standard kernel density estimation method.[5] The estimated densities represent the underlying distribution of heterogeneous inflation expectations. In Figure 3, we present an example of the estimated density of one-year ahead expected inflation (red line) using the data for April 2011 and the distribution based on actual responses (blue bars).

[5]Density functions are estimated using the Epanechnikov kernel function with the rule-of-thumb time-varying bandwidth.

Figure 3. Survey Responses and the Estimated Density: An Example

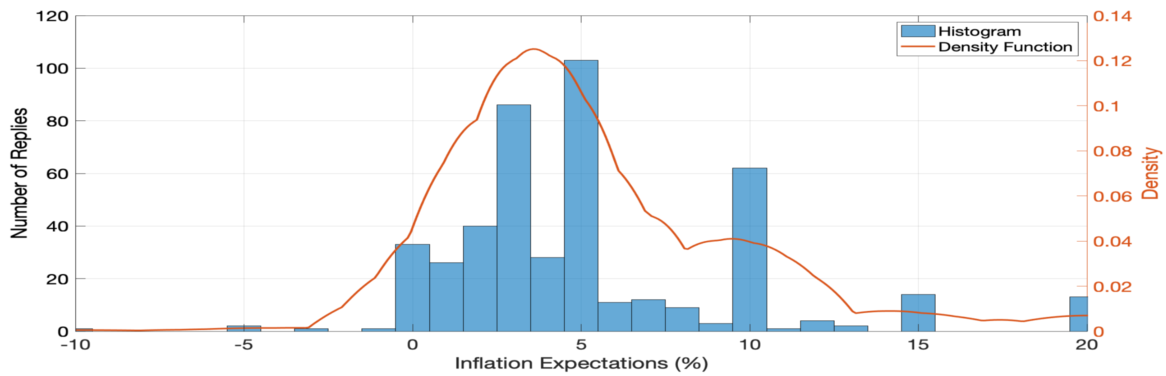

Notes: Survey responses of one-year ahead inflation expectations in April 2011 are used as an example. The left vertical axis represents the number of survey responses corresponding to different levels of inflation expectations reported in the actual survey. The right vertical axis shows the estimated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

4.2 Functional Principal Component Analysis

We apply the functional principal component analysis (FPCA) to these estimated densities to extract the leading components of one-year ahead and medium-run expected inflation distributions. The scree-plot presented in Figure 4 shows the cumulative share of the total variation of EIDs explained corresponding to the number of functional principal components for both one-year ahead and medium-run inflation expectations. The first principal component explains more than 74 percent of the total variation in one-year inflation expectations. In turn, the first three components combined account for about 95 percent of the total variation in both one-year and medium-run inflation expectations. Based on this observation, we choose the first three components ($m = 3$) to implement the functional approach described in the previous section, as it appears that these three components can sufficiently capture most of the variations of EIDs over time.[6]

[6]While a larger value of m may explain a greater variation of EIDs, having more components tend to increase the variances of functional coefficient estimates drastically, as shown in Tables 4 and 5 in Appendix.

Figure 4. *Scree Plots of One-Year Ahead and Medium-Run Inflation Expectations*

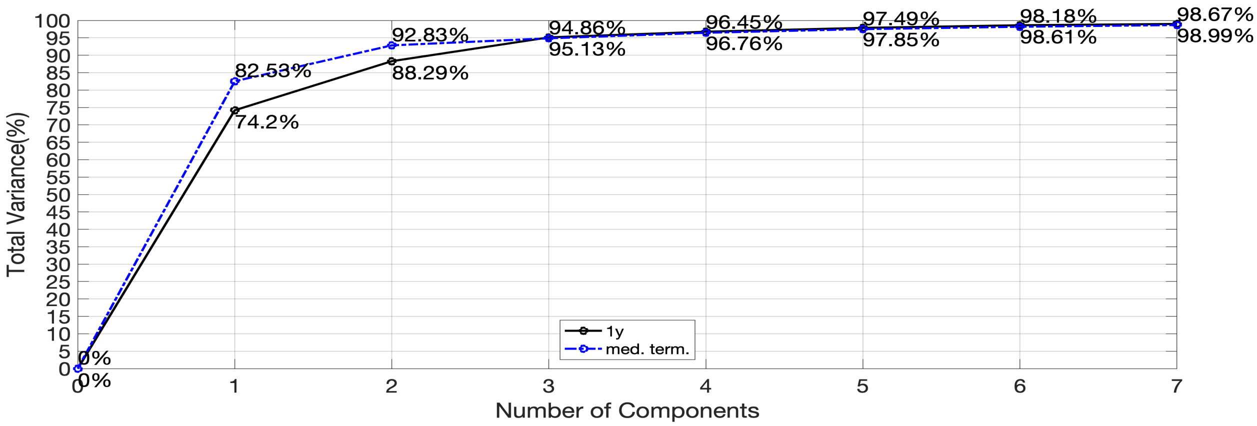

Notes: The horizontal axis shows the number of functional principal components. The vertical axis shows the cumulative proportion of functional variations explained, corresponding to the number of principal components. Black line (blue dotted line) show one-year ahead (medium-run) inflation expectations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

Figure 5 presents the results of FPCA for one-year ahead expected inflation distributions, with functional principal components (FPCs) shown in the center column and their loadings in the left column. In the center column, each row corresponds to the first FPC (v_1^*), second FPC (v_2^*), and third FPC (v_3^*). In the left column, the corresponding loadings for each FPC, $(\alpha_{kt}) = (\langle v_k^*, f_t \rangle)$, are plotted, where the blue and red stars in the figures indicate the minimum and maximum values of loadings, or $\min(\alpha_{kt})$ and $\max(\alpha_{kt})$, over time $t = 1, \ldots, T$ for each $k = 1, 2, 3$. In the right column of Figure 5, we plot the combined effects of FPC and loadings on the shape of EID. To do so, we multiply each FPC by its minimum and maximum loadings and add them to the average EID density over time, which is calculated as $\bar{f} = (1/T) \sum_{t=1}^{T} f_t$. The black solid, blue dotted, and red dashed lines represent the average density $(\bar{f})$, average density plus the FPC scaled by the minimum loading $(\bar{f} + \min(\alpha_{kt}) v_k^*)$ and the average density plus the FPC scaled by the maximum loading $(\bar{f} + \max(\alpha_{kt}) v_k^*)$ for $k = 1, 2, 3$, respectively. The red-dashed (blue dotted) line,

Figure 5. Functional Principal Components and Loadings: One-Year-Ahead Inflation Expectations

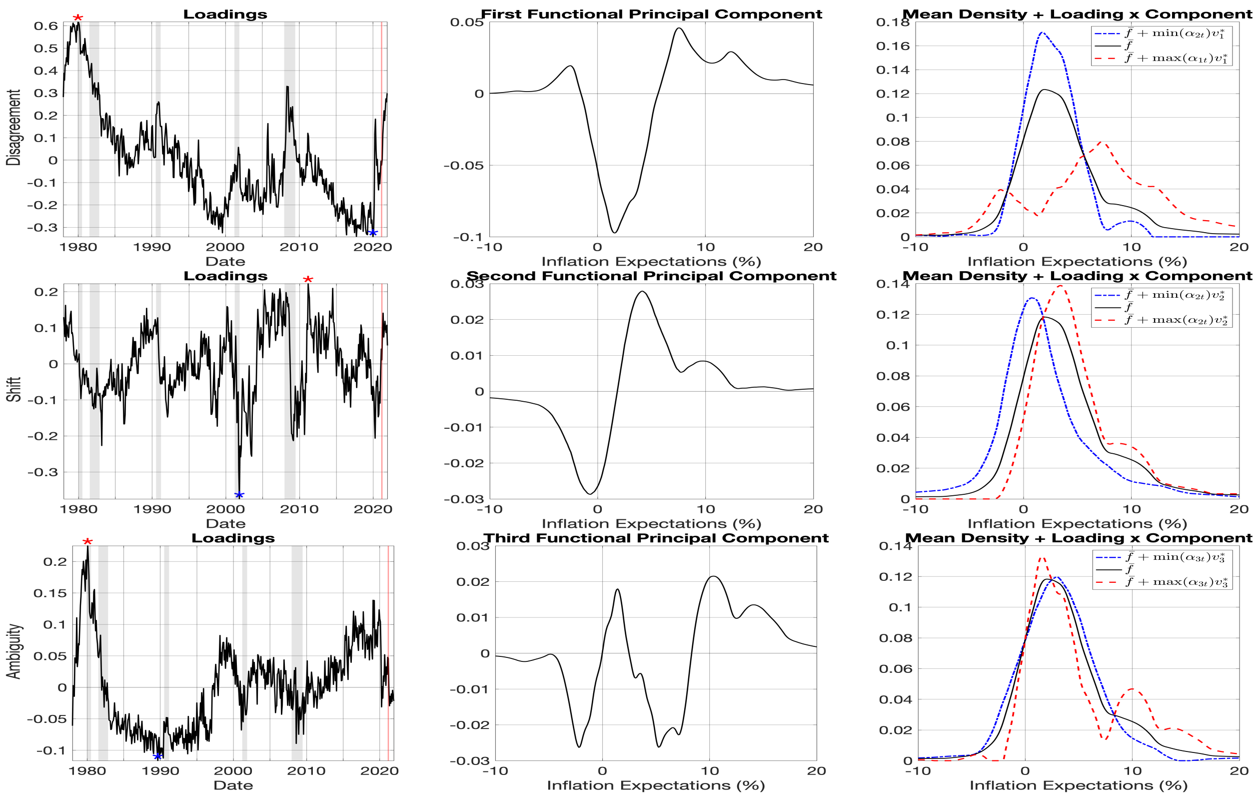

Notes: The left panels show loadings over time for the first three functional principal components (FPCs) for one-year ahead inflation expectations, the first FPC in the top, the second second FPC in the middle, and the third FPC in the bottom rows). The center panels show each of the three FPCs as a 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.

therefore, illustrates the case in which the EID is affected most positively (negatively) by the FPC.[7]

The first principal component alone, shown in the first row of Figure 5, explains more than 74 percent of the total variance of one-year ahead EID. We refer to this component as the *disagreement component* for the following reasons. This FPC accentuates the bi-modal distribution, with an increase in the densities of negative inflation expectations between -5

[7] To substantiate our interpretation and labeling of the three functional components, we provide in Appendix B detailed analyses on how moments of EID change as loadings of each functional principal component vary.

and -1 percent, as well as an increase in the densities of high (above 5 percent) inflation expectations. At the same time, the component decreases the densities in the middle range, with inflation expectations between -1 and 5 percent. As shown in the top right figure, the dispersion of the EID increases with the maximum loading (red dashed line), while the EID becomes more concentrated with the minimum loading (blue dashed line). Interestingly, the loading for the *disagreement component* demonstrates cyclical features to some degree, as the periods with high values of the loadings overlap with the times of economic recessions such as the beginning of the 80's, the Global Financial Crisis and the COVID-19 pandemic.

The second component, shown in the middle row of Figure 5, explains around 15 percent of the variations in the inflation expectation distributions over time. We call this component the *shifting component*, as the EID shifts to the right with positive loadings and to the left with negative loadings, as shown by the red dashed and blue dotted lines in the rightmost figure in the middle row.

The third component, shown in the bottom row of Figure 5, explains about 7 percent of the variation in the inflation expectation distributions over time, and we label this component as the *ambiguity component*. This component increases the frequency of moderate inflation expectations at around 2 percent, the Fed's current (average) inflation target, but also very high expectations above 8 percent. When its loading takes a positive value, it generates fat tails on the right side of the distribution, creating several bumps for high levels of inflation expectations, as shown by the red dashed line in the bottom panel in the last column of Figure 5. In contrast, when the ambiguity component interacts with negative loadings, it leads to a smoother tail (blue dotted line).

4.3 Functional Shock Identification

To identify shocks driving the expected inflation distributions, we write the error term $((\varepsilon_t))$ in the approximate second-order VAR in (6) as $(\varepsilon_t) = (\epsilon_{1t}, \epsilon_{2t}, \epsilon_{3t})$ for $t = 1, \ldots, T$, where (ϵ_{1t}), (ϵ_{2t}) and (ϵ_{3t}) are innovations in our approximate VAR. Note that (ϵ_{1t}), (ϵ_{2t}) and (ϵ_{3t}) are the loadings of three leading FPCs, which we labeled as the disagreement, shifting and ambiguity components, respectively, in our three-dimensional approximation $(\lceil \varepsilon_t \rfloor)$ of (ε_t). Let

$$\lceil \varepsilon_t \rfloor = Be_t \tag{8}$$

where B is a three-by-three matrix and (e_t), a three-dimensional vector representing shocks to the EID, which we define explicitly below.

In what follows, we let $e_t = (e_{1t}, e_{2t}, e_{3t})'$ for $t = 1, \ldots, T$ and tentatively identify three functional shocks (e_{1t}), (e_{2t}) and (e_{3t}) by assuming a recursive structure among them. The recursive structure we impose here is purely for convenience and inconsequential in our subsequent analysis. The matrix B in (8) is then defined uniquely as a three-by-three lower triangluar matrix satisfying the relationship

$$BB' = \lceil \Sigma \rfloor,$$

where $\lceil \Sigma \rfloor$ is the estimated covariance matrix of (Be_t) such that $Be_t = \lceil \varepsilon_t \rfloor$ for $t = 1, \ldots, T$.

Each column β_i of the matrix B thus defined represents at-impact response of the three dimensional vector $(\lceil f_t \rfloor)$ to functional shock (e_{it}) for $i = 1, 2, 3$. Correspondingly, we define

$$b_i = \pi^{-1}(\beta_i) \tag{9}$$

to be at-impact response of (f_t) to functional shock (e_{it}). Note that $\lceil b_i \rfloor = \beta_i$, and therefore, β_i and b_i are the vector and functional versions of at-impact response of the EID to functional shock (e_{it}), respectively, for $i = 1, 2, 3$. In our subsequent discussions, b_i will be referred to as the *baseline EID response* to functional shock (e_{it}) for $i = 1, 2, 3$.

Figure 6. At-impact Baseline EID Responses to Functional Shocks Using Different Samples

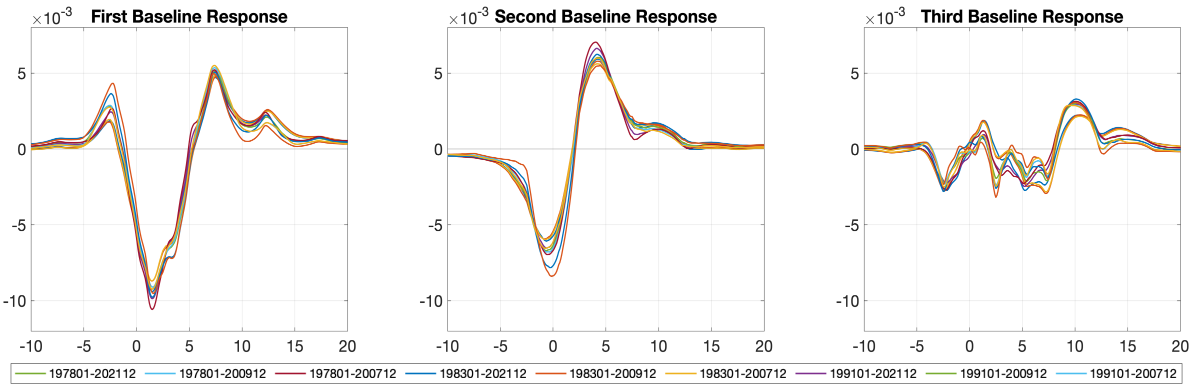

Notes: The figure shows the at-impact baseline impulse responses of one-year ahead EIDs to the three functional shocks using different sample periods, including the baseline sample from January 1978 to December 2021 for Figure 5.

Figure 6 presents the baseline EID responses to each of the three functional shocks. In addition to the baseline sample, they are calculated for different sample periods to demonstrate the stability of EID responses in terms of shapes and magnitudes across different samples. Recall that β_i's for $i = 1, 2, 3$ are three columns of a lower triangular matrix B. Therefore, b_3 is defined exclusively by a constant multiple of the third FPC (v_3^*), b_2 by a linear combination of the second and third FPCs (v_2^*, v_3^*), and b_1 by a linear combination of all three FPCs (v_1^*, v_2^*, v_3^*). The third baseline EID response b_3 thus has exactly the same shape as the third FPC (v_3^*). The first and second EID responses, b_1 and b_2, also have very similar, though not exactly the same, shapes as the first and second FPCs (v_1^*) and (v_2^*), respectively. This is because the loadings of leading FPCs truly dominate those of subsequent FPCs in our empirical model. For this reason, we will call these three baseline EID responses (b_1, b_2, b_3) as responses in disagreement, shifting and ambiguity, respectively, using the same labels that we used for the FPCs (v_1^*, v_2^*, v_3^*).

The three functional shocks (e_{1t}), (e_{2t}) and (e_{3t}) are not identified in any usual structural sense, since they are not economically interpretable. To be able to identify structural shocks meaningfully, we need to introduce additional identifying restrictions. Chang et al.

(2023) and Bjørnland et al. (2023), for example, impose restrictions in the responses of functional variables such as income distributions and stock return distributions to identify structurally interpretable shocks. The interested readers are referred to their papers for details. We may similarly identify structural shocks defined as linear combinations of the the functional shocks (e_{1t}), (e_{2t}) and (e_{3t}) that satisfy certain restrictions in the responses of EID to them.

In the paper, however, we take a different path. We identify economically interpretable shocks by linking externally identified economic shocks to the three functional shocks (e_{1t}), (e_{2t}) and (e_{3t}). As externally identified economic shocks, we will consider monetary policy (mp), government spending (gs), personal income tax (it), and gasoline price change (gp) shocks. The economic shocks we consider are available at different frequencies. Monetary policy and gasoline price shocks are available monthly, the same frequency as EID functional shocks, while government spending and personal income tax shocks are available only at quarterly. To be able to compare subsequent results on EID responses to different economic shocks, we set the frequency of all our data, both EID functional shocks and economic shocks, at the lower quarterly frequency.[8]

To determine the response of EIDs to an external shock (x_t), we compute the correlations between the external shock (x_t) and three functional shocks, (e_{it}), $i = 1, 2, 3$. Denote the correlation vector by $\rho_x = (\rho_{x1}, \rho_{x2}, \rho_{x3})'$, with $\rho_{xi} = \text{corr}(e_{it}, x_t)$ for $i = 1, 2, 3$. Since the individual functional shocks, (e_{it}) for $i = 1, 2, 3$, are uncorrelated, we can interpret the norm $\|\rho_x\|$ of the correlation vector as the percentage of the economic shock (x_t) that is transmitted to the EID. We use these correlation coefficients $(\rho_{x1}, \rho_{x2}, \rho_{x3})'$ as weights to

[8]We use the last observation in each quarter to obtain quarterly EID data and use them to compute quarterly EID functional shocks. For monetary policy and gasoline price shocks, we use the sum of all three monthly values to obtain quarterly time series.

define the at-impact response of EID to the given external shock x_t as

$$\phi_x = \rho_{x1} b_1 + \rho_{x2} b_2 + \rho_{x3} b_3, \tag{10}$$

where b_i is the at-impact baseline response of EID to the i-th functional shock for $i = 1, 2, 3$, defined in (9) and presented in Figure 6. Note that the external shock (x_t) has been normalized to have unit variance so that we may interpret ϕ_x as the response of EID to a one standard deviation external shock (x_t).[9]

We may readily obtain functional impulse responses of the EID to each economic shock x_t at a future horizon $h > 1$ as we did for the impact date in (11) as

$$\phi_{x,h} = \rho_{x1} b_{1,h} + \rho_{x2} b_{2,h} + \rho_{x3} b_{3,h}, \tag{11}$$

using the correlations ρ_{xi} as weights to combine the h-period ahead baseline EID response $b_{i,h}$ to the economic shock x_t. We present the three dimensional impulse response surfaces of EIDs to each of the four economic shocks whose two-dimensional slices present functional EID responses at each horizon $h = 0, 1, 2, \ldots, 60$ quarters in Appendix C.

5 Distributional Effects of Economic Shocks on EID

This section presents the results on the effects of externally identified economic shocks on the distributions of one-year ahead and medium-run inflation expectations using the methodology described above.[10]

[9] When a different shock size is desired, we may simply adjust the scale of the EID response as $\theta_x \phi_x$ with an appropriately chosen factor θ_x, viz., $\theta_x \phi_x = \rho_{\theta x1} b_1 + \rho_{\theta x2} b_2 + \rho_{\theta x3} b_3$, where $\rho_{\theta xi} = \theta_x \rho_{xi}$ are scale-adjusted new weights. For example, when a one-standard deviation monetary policy shock represents an increase in 30 bp of the federal funds rate, we may compute the response of EID to a 25 bp increase in the policy rate simply by setting the factor $\theta_{mp} = 25/30 = 5/6$.

[10] The IRF results are broadly similar using different samples. The baseline sample runs from January 1983 to December 2021 for one-year ahead inflation expectations, and from January 1991 to December 2021 for medium-run inflation expectations.

5.1 Economic Shocks

The four external shocks studied in this paper are monetary policy, government spending, personal income tax and gasoline price shocks.[11] To begin, for the monetary policy shock, we use the series constructed by Miranda-Agrippino and Ricco (2021), who introduce a high-frequency identification strategy to build an instrumental variable that makes the identification of monetary policy shocks robust in the presence of informational frictions. How contractionary monetary policy shocks affect inflation expectations is ambiguous, *a priori*. On the one hand, an increase in interest rates can dampen aggregate demand, lowering inflation expectations reflecting a lower output in the future. On the other hand, an increase in interest rates can be interpreted as a signal that the central bank is anticipating inflationary pressures (Nakamura and Steinsson, 2018), which can consequently lead households to increase inflation expectations.

For government spending shocks, we refer to Auerbach and Gorodnichenko (2012), based on 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). To the extent government spending shocks may lead to higher future public debt, government spending shocks can lead to higher inflation expectations. At the same time, if households expect an increase in future taxes to finance the current government spending, the net effect on inflation expectations is not clear.

As for our study for personal income tax shocks, we use the series constructed by Mertens and Ravn (2013) based on a narrative approach. The literature documents that an increase in taxes has contractionary effects, accompanied by a significant drop in inflation (Alesina and Ardagna, 2010; Romer and Romer, 2010). However, as documented in Mertens and Ravn (2013), different tax items may have varying impact on inflation. The

[11] Except for gasoline price shocks, economic shocks and the estimation codes are made available by the original authors of these shocks.

effect of these specific tax changes, such as personal income tax or corporate income tax, on household inflation expectations is less clear. 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. *A priori*, the direction of the adjustment is ambiguous. Households may increase inflation expectations if they interpret personal income tax increases as negative economic news. At the same time, an increase in personal tax income may result in lower future public debt, which may cause households' inflation expectations to be revised downward.

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 shocks, drawn from FRED's "Consumer Price Index for All Urban Consumers: Gasoline (All Types) in the U.S. City Average." The literature shows that gasoline price hikes, as salient shocks, play a crucial role in shaping households' inflation expectations. However, there is relatively sparse literature on the distributional effects of gasoline prices on households' inflation expectations.

5.2 Linking EID Functional Shocks to Economic Shocks

Because the individual functional shocks, $(e_{it}),\ i = 1, 2, 3$, are uncorrelated with each other, the norm of the correlation vector, $\rho_x = (\rho_{x1}, \rho_{x2}, \rho_{x3})'$, can be interpreted as percentage of the economic shock x that is transmitted to expected inflation distribution. The correlation

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

	Disagreement	Shifting	Ambiguity	Norm
Monetary Policy: ρ_{mp}	0.040 (0.063)	−0.087 (0.070)	0.084 (0.068)	0.128 (0.050)
Government Spending: ρ_{fp}	0.029 (0.114)	0.166 (0.104)	0.110 (0.100)	0.202 (0.090)
Personal Income Tax: ρ_{it}	−0.019 (0.096)	−0.035 (0.109)	0.015 (0.098)	0.043 (0.071)
Gasoline Prices: ρ_{gp}	0.057 (0.080)	0.384 (0.045)	0.195 (0.047)	0.434 (0.041)

Notes: Correlations between the three functional shocks of one-year ahead EIDs 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). Boostrapped standard errors are reported in parentheses. The sample size depends on the availability of economic shocks.

of each functional shock with each economic shock for the one-year ahead inflation expectations is reported in Table 2 from columns 2 to 4. The last column reports the norm for each economic shock (x_t), which ranges from about 4 percent (personal income tax shocks) to about 43 percent (gasoline price shocks).[12] Monetary policy shocks have the largest and negative correlations with the shifting element. Government spending shocks have positive correlations with the three functional shocks, but these correlations are not statistically significant. On the other hand, the correlations between the functional shocks and personal income tax shocks are negative except with the ambiguity factor, but without statistical significance. Gasoline price shocks, on the other hand, have a positive relationship with both shifting and disagreement shocks.

Table 3 repeats the same exercise for the medium-run inflation expectations. The results show the varying impact of economic shocks on one-year ahead and medium-run inflation expectations. For instance, the norm for personal income tax shocks increases from 4 percent for one-year ahead inflation expectations to 45 percent for medium-run inflation expectations. Based on the norms, the correlations between EIDs and monetary

[12] Since the shocks we have used here come from different sources, it is important to interpret these results individually. We do not attempt or claim to identify the effect of these shocks simultaneously. Future work relating the EID to other economic aggregates such as inflation, output, or unemployment, for instance, will address the issue of simultaneous identification.

and gasoline price shocks are lower for medium-run inflation expectations than for one-year ahead inflation expectations. Correlations between monetary shocks and all three functional shocks are shown to be relatively weak, while personal income tax shocks are correlated with the shifting functional shock with statistical significance. The norm does not change much for government spending shocks for the two different horizons of inflation expectations, and the correlations with the three functional shocks remain weak. Compared with one-year ahead inflation expectations, the correlation between disagreement shocks and gasoline price shocks becomes more prominent for medium-run inflation expectations.

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

	Disagreement	Shifting	Ambiguity	Norm
Monetary Policy: ρ_{mp}	0.011 (0.049)	−0.109 (0.064)	−0.009 (0.057)	0.110 (0.053)
Government Spending: ρ_{fp}	0.100 (0.138)	0.011 (0.130)	−0.145 (0.125)	0.177 (0.107)
Personal Income Tax: ρ_{it}	0.080 (0.121)	0.430 (0.089)	−0.120 (0.121)	0.453 (0.078)
Gasoline Prices: ρ_{gp}	0.131 (0.061)	0.049 (0.051)	−0.017 (0.057)	0.141 (0.054)

Notes: Correlation between the three functional shocks of the medium-run expected inflation distributions and each of the four economic shocks are considered. Economic shocks are identical to those in Table 2.

5.3 Effects of Economic Shocks on Expected Inflation Distributions

In the following subsections, we document the results for each of the four shocks [13].

5.3.1 Monetary Policy Shocks

We first examine the effects of contractionary monetary policy shocks on EIDs. The at-impact responses of one-year ahead and medium-run inflation expectations to contractionary monetary policy shocks using the baseline sample from 1991Q1 to 2009Q4

[13] The entire impulse response surfaces of EIDs to each shock can be found in Appendix C

are shown in the top charts of Figure 7. Contractionary monetary policy shocks increase the frequency of negative one-year ahead inflation expectations, while decreasing that of moderate levels of inflation expectations (top left chart). The shape of the at-impact response for medium-run inflation expectations is largely similar (top right chart).

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

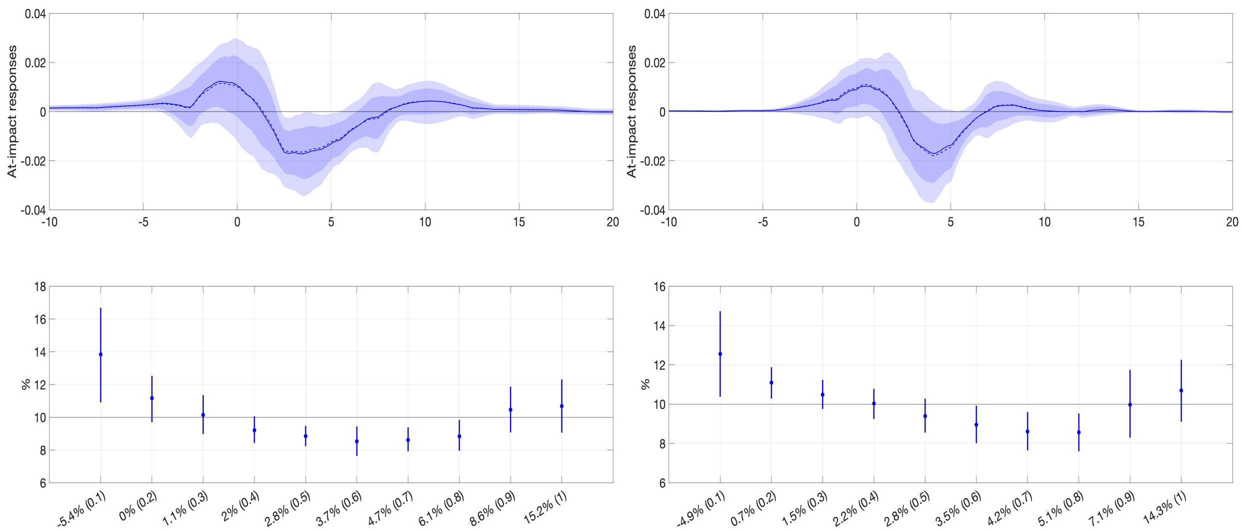

Notes: The top charts show the at-impact response of one-year ahead expected inflation expectations (left) and medium-run inflation expectations (right) following a contractionary monetary policy shock of 1 percent. The sample runs from 1991Q1 to 2009Q4. The IRFs include two confidence bands (68 and 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. For each decile, a dot represents the average post-shock probability, while a line indicates the confidence band of 68 percent.

Furthermore, we examine the effects of monetary policy shocks on the distribution of households' inflation expectations, by showing how the frequency in each decile bin based on households' initial levels of inflation expectations changes after a contractionary monetary policy shock. The results are show in the bottom charts of Figure 7. We divide the range of expected inflation along the horizontal axis into 10 intervals of equal probability, i.e., splitting the observations with cutoffs at $10^{th}, 20^{th}, \ldots, 90^{th}$ percentile values of the EID prior to the shock. This results in ten bins, each with a 10 percent probability. The average impact of contractionary monetary shocks on the frequency of each decile is displayed by a dot at the center, within a line representing a 68 percent confidence interval. Here, the horizontal line drawn at 10 percent represents the pre-shock probability of all ten bins, serving as the reference line to determine the statistical significance.

Following a contractionary monetary policy shock of 100 basis points, the probability in the first decile of one-year ahead EID increases by 3.5 percentage points to 13.5 percent from 10 percent (bottom left chart). That is, contractionary monetary policy shock increases the share of households in the first decile, with inflation expectations below negative 0.7 percent, to 13.5 percent from the pre-shock 10 percent probability. This change is statistically significant. The probability in the second decile, with inflation expectations ranging between -0.6 to 0.5 percent prior to the shock, increases by about 1 percentage point, but the result is not significantly different from the pre-shock 10 percent probability. On the other hand, there is an evident and significant decline in the share of households with inflation expectations between 4^{th} and 8^{th} deciles, with inflation expectations ranging from 1.6 to 6.9 percent. Contractionary monetary policy shocks do not appear to affect the share of households in the top two deciles, corresponding to inflation expectations above 7 percent.

The distributional impact of contractionary monetary policy shocks on medium-run inflation expectations is broadly similar (bottom right chart), but the magnitude of the changes is generally smaller for medium-run inflation expectations compared to that of one-year ahead inflation expectations. For instance, the probability in the first decile increases to 12.8 percent from 10 percent, corresponding to households with medium-run inflation expectations below 0.1 percent, somewhat lower than the 3.5 percentage points increase for one-year ahead inflation expectations. As in the case for one-year ahead inflation expectations, the share of households with moderate levels of medium-run inflation expectations between about 3 to 6 percent declines following contractionary monetary policy shocks, and these declines are statistically significant.

Next, we evaluate the impact of contractionary monetary policy shocks on the mean, standard deviation, and skewness of EID (Figure 8). For one-year ahead inflation expectations (top row), contractionary monetary policy shocks do not affect the average and standard deviation with statistical significance. Skewness for one-year ahead inflation expectations, however, declines, implying a disproportionate increase in the share of households with lower or negative inflation expectations. For medium-run inflation expectations (bottom row), monetary policy shocks do not affect the selected sample moments.

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

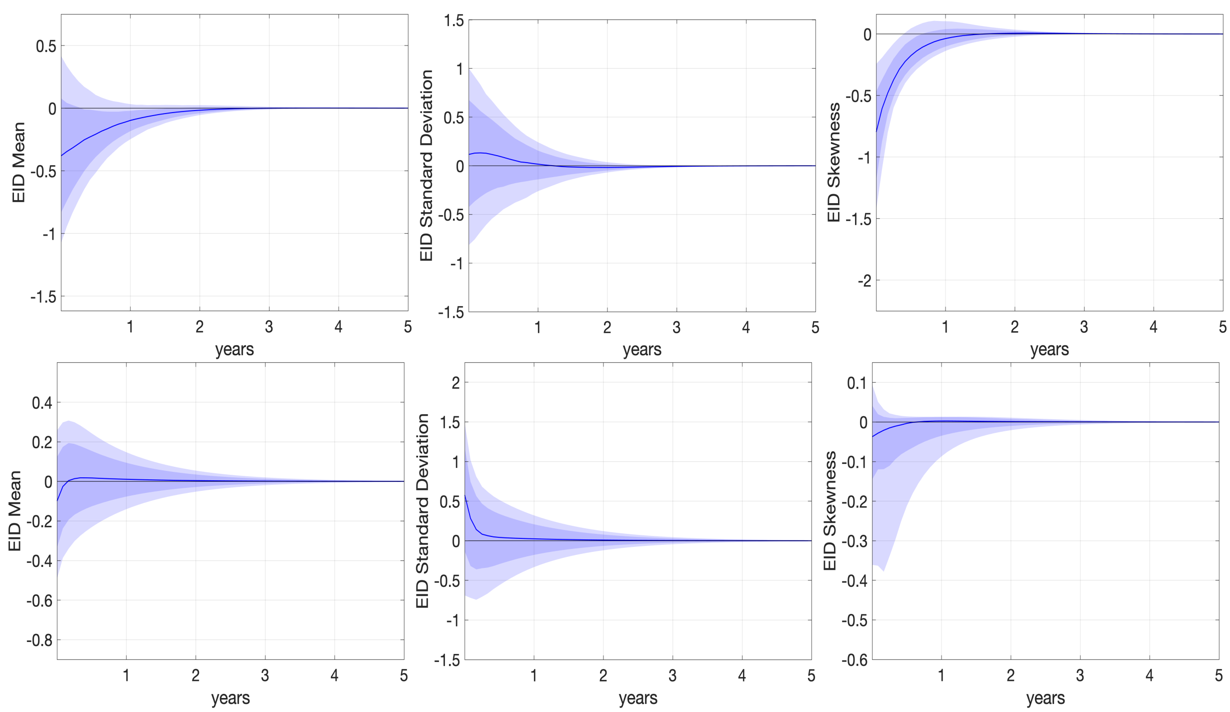

Notes: Each column shows the impact of contractionary monetary policy shocks of 1 percent on the mean (left), standard deviation (center), and skewness (right) of expected inflation distributions (EIDs). The top (bottom) row shows the impact on one-year ahead (medium-run) inflation expectations. The sample runs from 1991Q1 to 2009Q4. The IRFs include two confidence bands (68 and 90 percent) estimated using bootstrap methods.

5.3.2 Government Spending Shocks

Next, we explore the effects of government spending shock on EIDs based on the shocks from Auerbach and Gorodnichenko (2012). The top chart of Figure 9 presents the at-impact response for one-year ahead inflation expectations (left) and medium-run inflation expectations (right). For one-year ahead inflation expectations (top left chart), government spending shocks decrease the densities of inflation expectations below about 1 percent, while increasing those with positive and relatively higher inflation expectations compared to the moderate level observed for monetary policy shocks. The at-impact response of medium-run inflation expectations is more muted in comparison (top right chart), except for an increase in the frequency of inflation expectations at around 10 percent.

The bottom charts represent the results from the same distributional analysis conducted for monetary policy shocks, by dividing the range of EIDs into ten equal-probability intervals based on the deciles using the initial levels of household inflation expectations prior to government spending shocks. For one-year ahead inflation expectations (bottom left chart), we find that government spending shocks significantly increase the share of households with inflation expectations above 4.3 percent. This shows that government spending shocks have different distributional effects on households' inflation expectations than monetary policy shocks. In particular, government spending shocks affect the tail behaviors of inflation expectations, with the share of households with high levels of inflation expectations increasing, particularly the households in the 9th and 10th deciles prior to the shock. Monetary policy shocks, in contrast, did not affect high-levels of inflation expectations. For medium-run inflation expectations (bottom right chart), government spending shocks affect the share of households in the 10th decile only, by increasing the probability by 0.3 percentage points from 10 percent to 10.3 percent.

__Figure 9.__ One-year Ahead EID and Medium-Run EID Responses to Government Spending Shocks

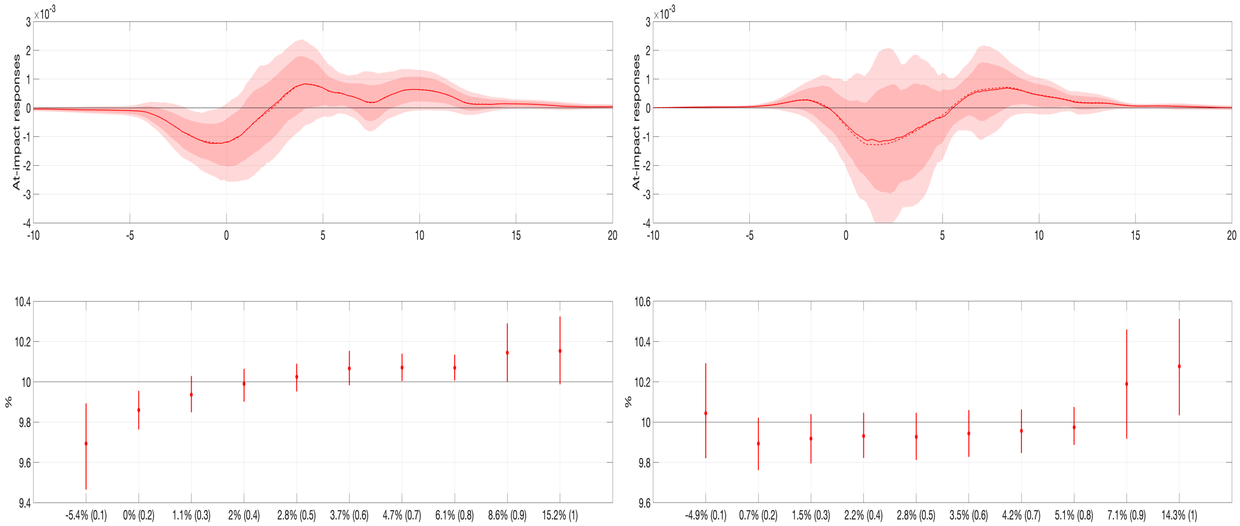

Notes: The top charts show the at-impact responses of one-year ahead (left) and medium-run inflation expectations (right) to government spending shocks using the sample from 1983Q1 (1991Q1 respectively) to 2008Q4. The bottom charts shows the estimated post-shock probabilities of decile bins for one-year 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.

Figure 10 compares how government spending shocks affect the mean (left column), standard deviations (center column), and skewness (right column). The mean of one-year ahead inflation expectations (top row) increases, confirming the observation in Figure 9. The average impact on medium-inflation expectations (bottom row) is not statistically significant. Government spending shocks do not affect standard deviations or skewness of households' inflation expectations for both forecast horizons.

5.3.3 Personal Income Tax Shocks

Next, we examine the effects of personal income tax shocks on EIDs. Top charts in Figure 11 exhibit the at-impact responses for one-year ahead (left) and medium-run inflation expectations (right) following personal income tax shocks. The effects of personal income tax shocks on one-year ahead inflation expectations are visibly different from those on

Figure 10. *Impact of Government Spending Shocks on Specific Moments of One-Year Ahead and Medium-Run EIDs*

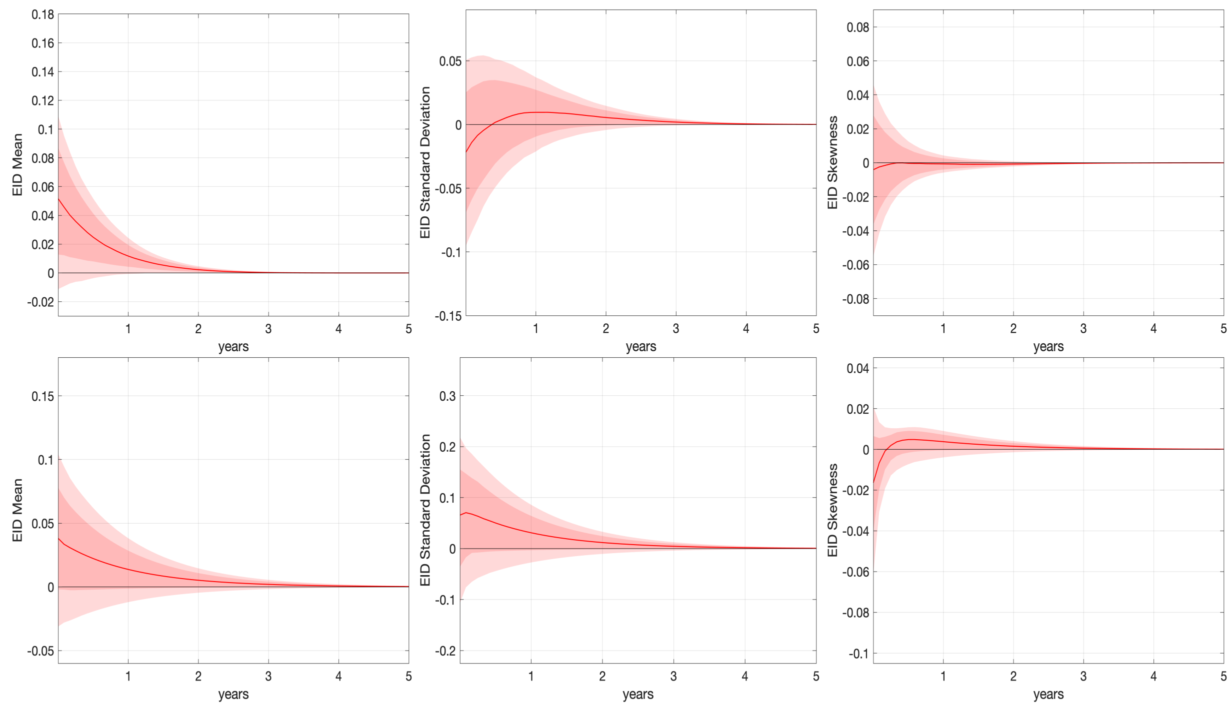

Notes: Each column shows the impact of government spending shocks on the mean (left), standard deviation (center), and skewness (right). The top row reports for the results for one-year ahead expected inflation distributions (EIDs), using the sample from 1983Q1 to 2008Q4. The bottom row reports the results for medium-run inflation expectations, using the sample from 1991Q1 to 2008Q4.

medium-run inflation expectations, as the results are statistically significant only for medium-run inflation expectations. For medium-run inflation expectations, personal income tax shocks tend to increase the average inflation expectations, with a decline in the share of households with low levels of inflation expectations and an increase in the share of households with higher levels of inflation expectations above 3 percent.

The bottom charts examining the change of frequencies by deciles confirm these observations. Compared to the muted at-impact response of one-year ahead inflation expectations, personal income tax shocks increase medium-run inflation expectations, by lowering the probabilities of inflation expectations below 3 percent, while increasing the probabilities of inflation expectations higher than 3 percent. The bottom right chart shows a general increase in the share of households with positive medium-run inflation expectations greater than 3.2 percent, accompanied by a general decline in the share of households with inflation expectations below 1.8 percent, with a greater decline observed for lower initial levels of medium-run inflation expectations prior to the shock.

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

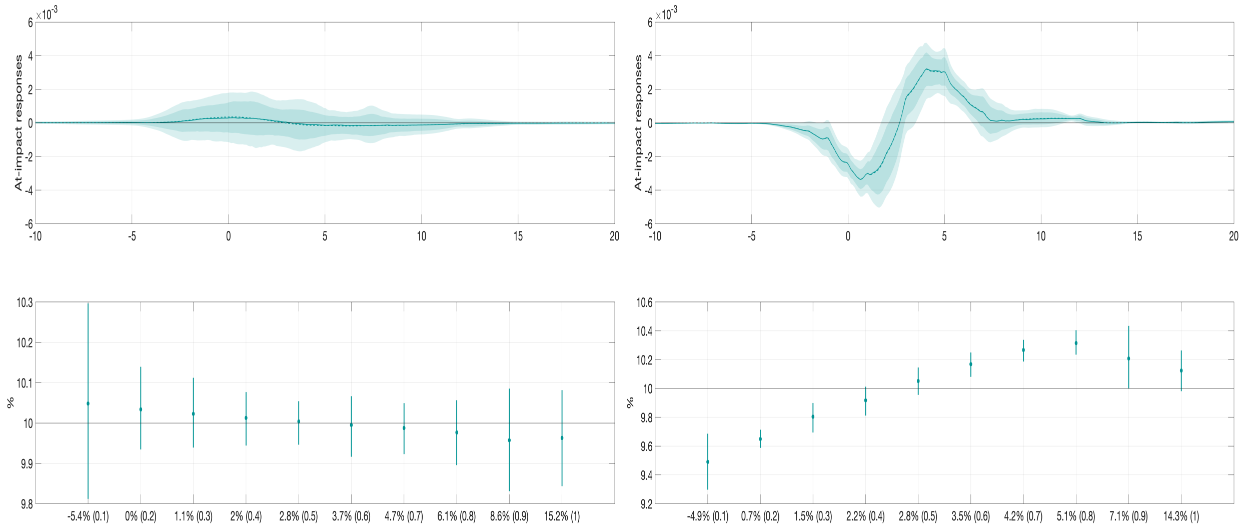

Notes: The top charts show the at-impact response of one-year ahead (left) and medium-run (right) inflation expectations to personal income tax shock using the sample from 1983Q1 (1991Q1 respectively) to 2006Q4. The bottom charts shows the estimated post-shock probabilities of decile bins for one-year 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.

Figure 12 presents the impulse responses of personal income tax shocks on the mean (left), standard deviations (center), and skewness (right). The sample statistics of one-year ahead inflation expectations are not affected by personal income tax shocks. Personal income tax shocks increase the average medium-run inflation expectations at-impact with statistical significance, but do not affect the other two statistics.

5.3.4 Gasoline Price Shocks

Finally, we examine the impact of gasoline price shocks on EIDs. The top charts of Figure 13 compare the at-impact responses of EIDs to gasoline price shocks for one-year ahead (left)

Figure 12. *Impact of Personal Income Tax Shocks on Specific Moments of One-Year Ahead and Medium-Run EIDs*

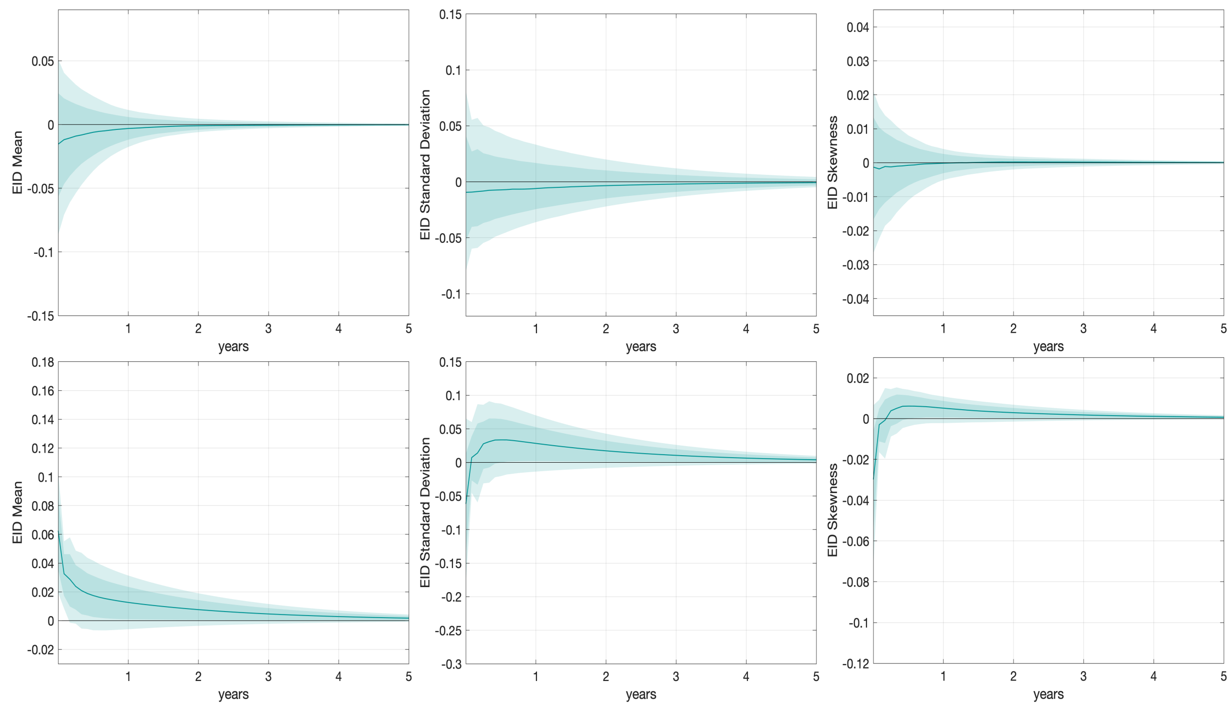

Notes: Each column shows the impact of personal income tax shock on the mean (left), standard deviation (center), and skewness (right). The top (bottom) row reports the results for one-year ahead (medium-run) expected inflation distributions, using the sample from 1983Q1 to 2006Q4.

and medium-run inflation expectations (right). These confirm the findings in the literature that gasoline price shocks are inflationary (Harris et al., 2009; Coibion and Gorodnichenko, 2015). Furthermore, We find that, from distributional perspectives, gasoline price shocks affect one-year inflation expectations differently from medium-run inflation expectations, which is a novel finding in the literature.

For one-year ahead EIDs (top left), a surprise hike in gasoline prices decreases the frequency of inflation expectations below 3 percent, while increasing the frequency of inflation expectations higher than 3 percent at the same time. This confirms the findings in the literature that gasoline price shocks lead to an increase in the average level of one-year ahead inflation expectations. In contrast, for medium-run inflation expectations, the same shock not only increases the frequency of positive inflation expectations, but also that of negative inflation expectations. In other words, the disagreement among households on the impact of the current gasoline price shocks on the future price level widens.

The bottom chart examines the distributional changes following the shock. For one-year ahead inflation expectations (bottom left), the share of households with pre-shock inflation expectations lower than 1.5 percent, declines in a linear way, in the sense that the magnitude of the decline is larger for households with lower levels of pre-shock inflation expectations. This decline is accompanied by an increase in the share of households with pre-shock inflation expectations above 2.5 percent. The bottom right chart shows the distributional impact of the gasoline price shocks on medium-run inflation expectations. Contrary to the one-year ahead inflation expectations, the share of households with low medium-run inflation expectations, below 0.1 percent, increases with statistical significance, pointing to the increase in disagreement following gasoline price shocks.

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

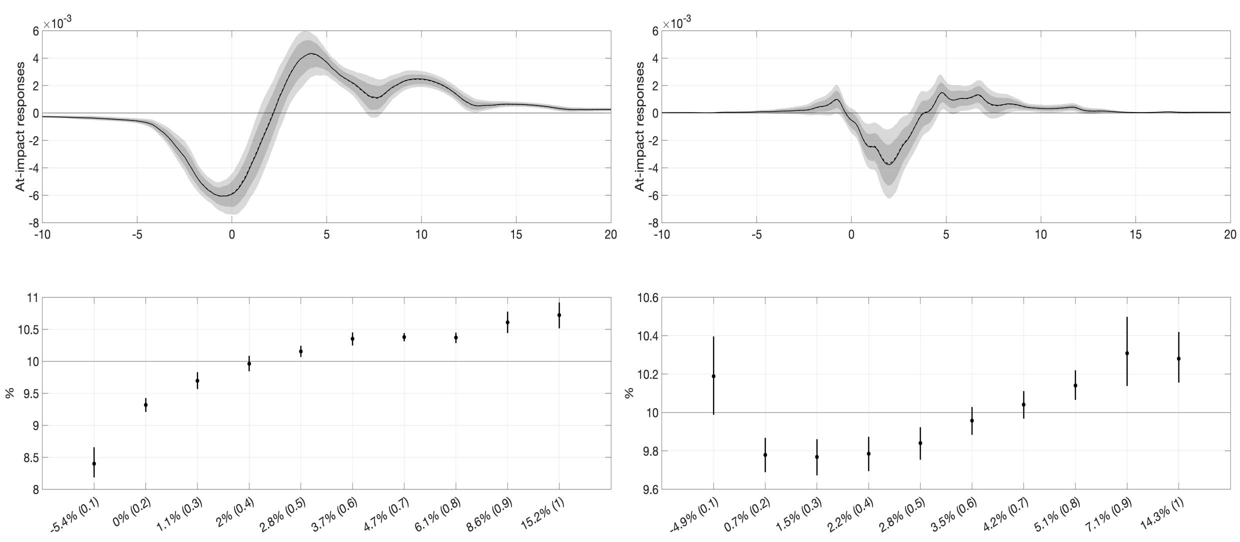

Notes: The top charts show the at-impact response of one-year ahead (left) and medium-run (right) expected inflation distributions (EIDs) to gasoline price shocks using the sample from 1983Q1 to 2021Q4. The bottom charts shows the estimated post-shock probabilities of decile bins for one-year ahead (left) 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.

The examination of the impact of gasoline prices shocks on selected moments reaffirms such observations (Figure 14). A surprise increase in gasoline prices leads to higher average inflation expectations for both horizons, with a larger increase for one-year ahead inflation expectations (left column). In fact, the response of one-year ahead inflation expectations at-impact is nearly four times larger than that of medium-run inflation expectations. Standard deviation decreases for one-year ahead inflation expectations (top center), implying that the disagreement among households on the level of inflation one year from the shock is reduced following the gasoline price shock. On the contrary, standard deviation for medium-run inflation expectations (bottom center) increases in response to gasoline price shocks, and the shock has a significant impact even 5 years after the shock. Finally, skewness declines for medium-run inflation expectations (bottom right), but increases for one-year ahead inflation expectations (top right) for several months after the shock.

5.4 Interpretation and Plausibility of Findings

The precise transmission channels through which economic shocks affect the distribution of inflation expectations are beyond the scope of our paper. Nonetheless, in this section, we attempt to provide plausible explanations for our findings in the established transmission mechanisms of each shock in the literature, while leaving a more rigorous analysis for future research. But, we also note that there are channels that would predict a different outcome than ours. The lack of empirical evidence on the transmission of specific channels of shocks on the distribution of inflation expectations is a significant barrier to testing different predictions, emphasizing the need for additional research in this area.

First, why would contractionary monetary policy shocks lead to more frequent responses of negative inflation expectations?[14] *A priori*, contractionary monetary policy

[14] In fact, empirical studies have shown that contractionary monetary policy shocks lead to an increase in inflation, through neo-fisherian effects (Uribe, 2022), or price puzzles (Sims, 1992; Hanson, 2004; Rusnák et al., 2013).

Figure 14. *Impact of Gasoline Price Shocks on Specific Moments of One-Year Ahead and Medium-Run EIDs*

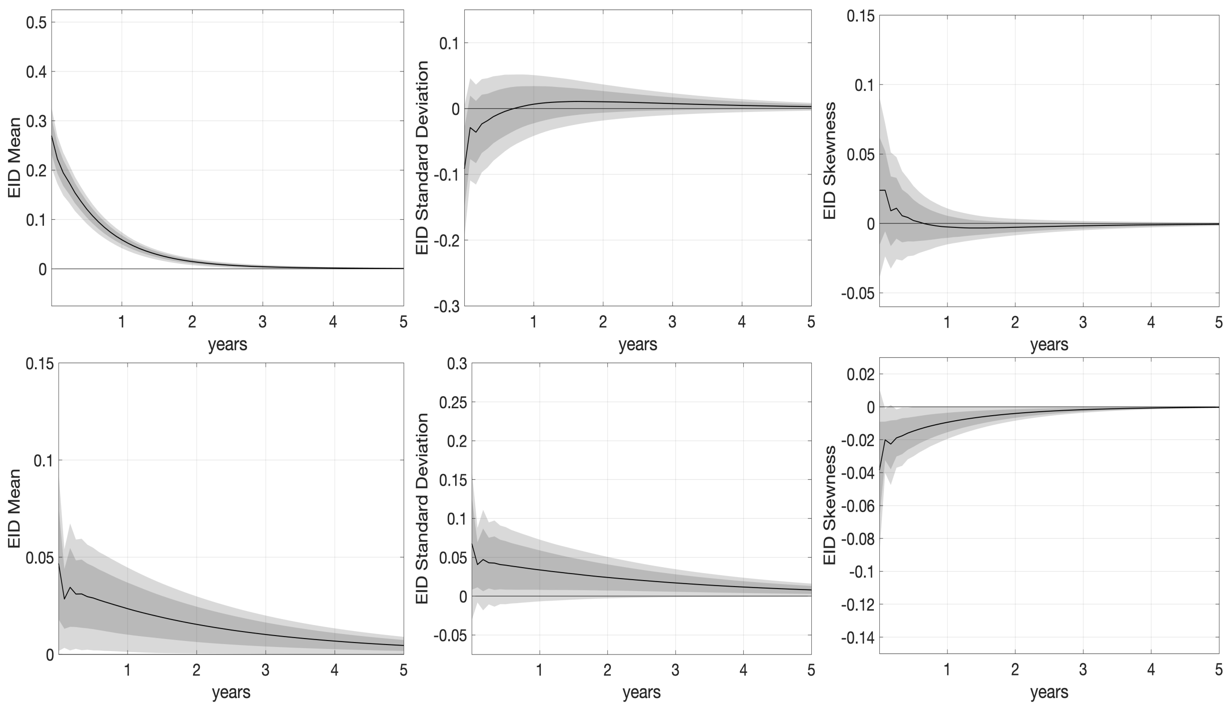

Notes: Each column shows the impact of gasoline price shock on the mean (left), standard deviation (center), and skewness (right). The top (bottom) row reports the results for one-year ahead (medium-run) expected inflation distributions (EIDs), using the sample from 1983Q1 to 2021Q4.

shocks can either increase or decrease household inflation expectations. Empirical evidence is also mixed. While some studies show that policy news, including monetary policy, does not affect household inflation expectations insofar as inflation expectations are well-anchored (Coibion et al., 2022; De Fiore et al., 2022), others document that monetary policy tightening in the United States during the high inflation period of 2021-22 has resulted in a share of households with deflationary inflation expectations, particularly for households with more optimistic economic outlook (Armantier et al., 2022). Our findings may appear consistent with the conventional Phillips Curve channel, whereby agents expect slower economic activity following contractionary monetary policy, leading to lower inflationary pressures, and, thus, lower inflation expectations. On the other hand, this interpretation may be at odds with the findings in Coibion and Gorodnichenko (2015) that show households tend to associate inflation with negative news to future economic conditions. If so, households may perceive contractionary monetary policy shocks as negative shocks to the economy, which would lead to an increase in inflation expectations.

The findings that government spending shocks increase inflation expectations in the short-run but not for medium-run suggest myopic consumption behaviors and may not bode well with the predictions from well-established models using forward-looking consumers. For instance, from the perspective of an inter-temporal government budget constraint where fiscal decisions in the current period affect the level of future prices (Leeper and Nason, 2010), current government spending shocks may affect households' beliefs about future prices. Conversely, our finding that shows no meaningful impact of government spending shocks on medium-run inflation expectations may be at odds with the literature that discusses a delay in public spending implementation, as suggested in Leeper et al. (2010), where the economic impact of current spending shocks span several periods.

Finally, to our knowledge, few studies have looked into the channels through which personal income tax shocks affect household inflation expectations. Personal income tax shocks associated with an increase in inflation expectations may be consistent with the studies of perfect foresight, as households may perceive the current increase in personal income tax as a higher future government spending, leading to an upward revision in medium-run inflation expectations. However, a handful of studies caution against the restrictiveness of perfect foresight as an assumption, as the perfect foresight assumption does not allow agents to be uncertain about the beliefs and the response of others as documented in Angeletos and Lian (2018) and García-Schmidt and Woodford (2019).

6 Conclusion

In this paper, we document several novel findings on how economic shocks affect the distribution of household inflation expectations for one-year ahead and medium-run inflation expectations using a functional approach. Starting with monetary policy shocks, there are more households who report negative inflation expectations following contractionary monetary policy shocks. Fiscal policy shocks also affect the distribution of inflation expectations. We find that inflation expectations increase for one-year ahead inflation expectations following government spending shocks, but not for medium-run. Conversely, personal income tax shocks affect medium-run inflation expectations only. Finally, we confirm that the levels of household inflation expectations for both one-year ahead and medium-run inflation expectations respond to changes in gasoline price shocks. However, gasoline price shocks also have more persistent effects and increase disagreement among households' medium-run inflation expectations.

Our findings contribute to the policy discussion on how to anchor households' inflation expectations. So far, there has been limited evidence and discussion on the channels through which economic shocks affect inflation expectations, let alone the distributional impact of these shocks. Our findings highlight that both fiscal and monetary policy shocks exert their influence on the distribution of household inflation expectations in their own ways. A fruitful avenue for future research is to better understand the exact channel through which these policy measures affect household inflation expectations, individually and jointly. Another promising avenue for future research would be to explore to what extent the varying distributional aspects of inflation expectations may matter for the movements of aggregate macroeconomic variables that are of first-order importance to policymakers, such as inflation, unemployment rate, 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. (2023). Oil and the stock market revisited: A mixed functional var approach. *CAMP Working Paper Series 03/2023*, BI Norwegian Business School.

Bosq, D. (2000). *Linear Processes in Function Spaces*. Springer-Verlag.

Chang, Y., Gómez-Rodríguez, F., and Matthes, C. (2020a). How do the u.s. government's decisions affect its borrowing costs? Mimeo, Indiana University.

Chang, Y., Hu, B., and Park, J. Y. (2020b). Econometric analysis of functional dynamics in the presence of persistence. Mimeo, Indiana University.

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. (2023). How do macroaggregates and income distribution interact dynamically? A novel structural mixed autoregression with aggregate and functional variables. Mimeo, 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., and Kamdar, R. (2018). The formation of expectations, inflation, and the phillips curve. *Journal of Economic Literature*, 56(4):1447–91.

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.

Coibion, O., Gorodnichenko, Y., and Weber, M. (2022). Monetary policy communications and their effects on household inflation expectations. *Journal of Political Economy*, 130(6):1537–1584.

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.

De Fiore, F., Lombardi, M. J., and Schuffels, J. (2022). Are households indifferent to monetary policy announcements? *CEPR Discussion Paper No. DP17041*.

Favero, C. and Giavazzi, F. (2004). Inflation targeting and debt: lessons from brazil. *NBER Working Paper*, 10390.

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.

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.

Hu, B., Park, J. Y., and Qian, J. (2016). Analysis of distributional dynamics for repeated cross-sectional and intra-period observations. Mimeo, Indiana University.

Inoue, A. and Rossi, B. (2019). The Effects of Conventional and Unconventional Monetary Policy: A New Approach. Working Papers 1082, Barcelona Graduate School of Economics.

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.

Nakamura, E. and Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: the information effect. *The Quarterly Journal of Economics*, 133(3):1283–1330.

Park, J. Y. and Qian, J. (2012). Functional regression of continuous state distributions. *Journal of Econometrics*, 167:397–412.

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

Ramsay, J. O. and Silverman, B. W. (2005). *Functional Data Analysis*. Springer.

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.

Appendices

For Online Publication

A Choice of Basis

Here, we explain the approach adopted in this paper in order to select the basis. In fact, our approach based on the VAR representation in (6) may be implemented with any orthonormal basis (v_i) of H. However, the finite sample performance of the approach is critically dependent upon the choice of basis. In what follows, we denote by (v_i^*) the functional principal component basis used in the paper, and by V^* the subspace of H spanned by the sub-basis $(v_i^*)_{i=1}^m$ assuming $m < T$ and P^* to be the Hilbert space projection on V^*. In contrast, we let $(v_i)_{i=1}^m$ be an arbitrary sub-basis spanning the subspace V of H, and P be the Hilbert space projection on V in H.

As shown in Chang et al. (2021), the π's in (4) and (5) are *isometries*, not just one-to-one mappings between V and R^m. For π in (4), we have

$$\|f\|^2 = \|(f)\|^2$$

for any f in V, where we use the same notation $\|\cdot\|$ for the Hilbert space norm of f in V and the Euclidean norm of (f) in R^m. Similarly, for π in (5), we may show

$$|||A|||^2 = \text{trace}\,(A'A) = \text{trace}\,\big((A)'(A)\big) = |||(A)|||^2,$$

where A' is the adjoint of A, $(A)'$ is the transpose of (A), and $\|\cdot\|$ denotes both the Hilbert-Schmidt norm for a linear operator A on V whose trace$(A'A)$ is finite and the Frobenius norm for the $m \times m$ matrix (A).

Let

$$\mathrm{FR}^2 = \frac{\sum_{t=1}^T \|Pf_t\|^2}{\sum_{t=1}^T \|f_t\|^2}$$

be the functional R-squared (FR-squared) of an arbitrary sub-basis $(v_i)_{i=1}^m$, which represents the proportion of the total variation of (f_t) explained by its projection (Pf_t) on the subspace V spanned by $(v_i)_{i=1}^m$, with FR_*^2 denoting the FR-squared of $(v_i^*)_{i=1}^m$. Then we have

$$\mathrm{FR}_*^2 \geqslant \mathrm{FR}^2,$$

which implies that $(P^* f_t)$ has the maximum temporal variation. For both $k = 1$ and 2, the approximation of A_k given by $P^* A_k P^*$ thus restricts A_k to the subspace of V^* of H, where (f_t) has the largest variation and thus A_k is most strongly identified. In this sense, the basis (v_i^*) provides the most effective approximation of A_k for $k = 1$ and 2 by its restriction on an m-dimensional subspace V^* of H spanned by $(v_i^*)_{i=1}^m$.

Given any basis (v_i), we may make the FR-squared as large as we want simply by increasing the truncation number m. However, this does not come without a cost. As m gets large, the variances of the estimators $\widehat{A}_k$ for the autoregressive operator A_k for $k = 1$ and 2 are expected to increase. They increase often very sharply in many practical applications we have observed so far, and therefore, we also need to examine how fast the variance of $\widehat{A}_k$ increases as m gets large.

Let $\widehat{A}_k$ for $k = 1$ or 2 be the estimator obtained from $\widehat{(A_k)}$ by $\widehat{A}_k = \pi^{-1}\big(\widehat{(A_k)}\big)$, which we may regard more explicitly as the estimator of $\overline{A}_k = \pi^{-1}\big((A_k)\big) = PA_kP$, and let

$$\widehat{A} = \begin{pmatrix} \widehat{A}_1 \\ \widehat{A}_2 \end{pmatrix} \quad \text{and} \quad \overline{A} = \begin{pmatrix} \overline{A}_1 \\ \overline{A}_2 \end{pmatrix},$$

which are operators from H to $H \times H$. Furthermore, we define

$$Q = \begin{pmatrix} P & 0 \\ 0 & P \end{pmatrix} \quad \text{and} \quad \Delta = \frac{1}{T}\sum_{t=1}^{T}\left[\begin{pmatrix} f_t \\ f_{t-1} \end{pmatrix} \otimes \begin{pmatrix} f_t \\ f_{t-1} \end{pmatrix}\right].$$

Then the mean-squared-error (MSE) of $\widehat{A}$ is given by

$$\mathbb{E}\big\|\widehat{A} - \overline{A}\big\|^2 = \mathbb{E}\big\|\widehat{A} - \mathbb{E}\widehat{A}\big\|^2 + \big\|\mathbb{E}\widehat{A} - \overline{A}\big\|^2,$$

where we decompose it into its variance and squared bias terms. The variance term of $\widehat{A}$ is approximately given by

$$(\operatorname{trace}\Sigma)\big(\operatorname{trace}(Q\Delta Q)^+\big),$$

where Σ is the variance operator of (ε_t) as defined earlier, and $(Q\Delta Q)^+$ is the inverse of the bounded linear operator $Q\Delta Q$ restricted to the subspace $V \times V$ of $H \times H$. The squared bias term of $\widehat{A}$ is approximately given by

$$\big\|A(1-Q)\Delta Q(Q\Delta Q)^+\big\|^2,$$

where A is defined from A_1 and A_2 similarly as $\widehat{A}$ and $\overline{A}$.

In our subsequent discussions, we denote by $\widehat{A}^*$ and $\overline{A}^*$ the versions of $\widehat{A}$ and $\overline{A}$ obtained using our functional principal component basis (v_i^*), and show that the variance term $\mathbb{E}\big\|\widehat{A}^* - \mathbb{E}\widehat{A}^*\big\|^2$ of $\widehat{A}^*$ is significantly smaller than that of the estimator $\widehat{A}$ based on other bases in our empirical analysis. Although we cannot explicitly compute and compare them, the bias term $\left\|\mathbb{E}\widehat{A}^* - \overline{A}^*\right\|^2$ of $\widehat{A}^*$ is generally expected to be smaller than that of the other estimator $\widehat{A}$ since

$$P^*\Gamma(1-P^*) = 0,$$

whereas $P\Gamma(1-P) \neq 0$ for any other choice of basis (v_i). In fact, our methodology yields an unbiased estimator $\widehat{A}^*$ if the first order FAR, viz., $f_t = Af_{t-1} + \varepsilon_t$, is used in place of the second order FAR in (1).

To demonstrate the importance of the choice of basis in explaining the variance of EID, we compare the FR-squared's and the variance terms of the estimators of A based on our basis (v_i^*) and other bases. As an alternative to our basis (v_i^*), we consider three other bases given by the orthonormalized moments, histograms and quantiles, which will be referred to as the *moment basis*, *histogram basis* and *quantile basis*, respectively. 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 $\big(\langle u_i, f_t \rangle\big)$ represents the i-th moments of the EID given by the densities (f_t) for $i \geqslant 1$.

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). As before, we let $p = -0.5$ and $q = 0.5$ and obtain the $(m+1)$-number of sub-intervals $\big([p_i, q_i)\big)$ of equal length, from which we take only m indicators as a basis, ignoring the first sub-interval. This is because the $(m+1)$ indicators over the $(m+1)$-number of sub-intervals $\big([p_i, q_i)\big)$ are linearly dependent. The quantile basis $(v_i)_{i=1}^m$ is defined similarly as indicators over a

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 FR2 are reported for four different choices of basis including the functional principal component (FPC) basis, histogram basis, quantile basis and moment basis. The FR2 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 FR2's.

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 $(Q\Lambda Q)^+$, which are asymptotically proportional to the variances of the autoregressive operator estimators in the Hilbert-Schmidt norm relying on 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.

different set of partition $\big([p_i, q_i)\big)$. The $(m+1)$-sub-intervals $\big([p_i, q_i)\big)$ in the partition are obtained with (q_i) defined as the $i/(m+1)$-th sample quantiles of entire observations for $i = 1, \ldots, m+1$. Similar to the histogram basis, we only include m indicators in a quantile basis.

Tables 4 and 5 report the FR-squared's and the variances of $\widehat{A}$ in the Hilbert-Schmidt norm based on four different choices of basis, including the functional principal component (FPC) basis, histogram basis, quantile basis and moment basis, computed using the time series of one-year ahead expected inflation distributions from January 1978 to December 2021. We may clearly see that the FPC basis effectively represents the temporal variation of the EID even with the trunction number m as small as $m = 1$ or 2. For $m = 3$, it captures more than 95 percent of the total temporal variation of the EID. Moreover, the variance of the autoregressive operator estimator based on the FPC basis increases as m gets large, but only at a moderate rate. In sharp contrast, all other bases obviously do not represent the temporal variation of the EID adequately. The moment basis is especially ineffective. It captures only 8.8 percent of the total variation of the EID over time even for $m = 10$. What is worse, the variances of the autoregressive operator estimator based on three other bases increase very rapidly as m increases. In particular, the use of the histogram basis or moment basis yields exploding variances of the autoregressive operator estimator even for a moderately large value of m.

As discussed, we use $m = 3$ in our empirical analysis of the EID using the FPC basis. Our choice of m with the use of the FPC basis explains 95.1 percent of the total EID variation and yields our variance measure 4.7 for the autoregressive operator estimator. If the moment basis is used, only 1.7 percent of the total EID variation is explained while we have the corresponding variance measure 3916 for the autoregressive operator estimator. The choice of basis is therefore critically important for our functional method in the paper.

B 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 15 shows the change in these four moments with respect to loadings for each FPC.

Figure 15. *Range of variation of the mean, standard deviation, skewness and kurtosis due to the functional principal components*

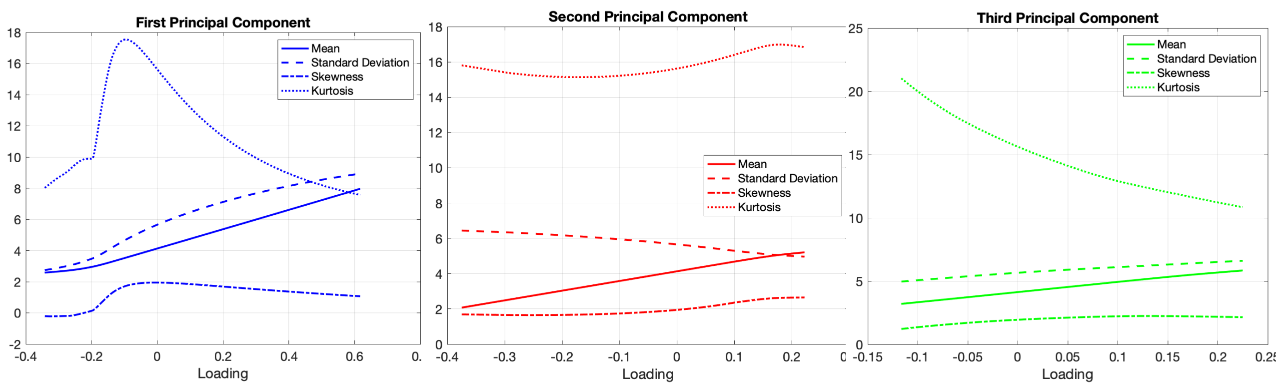

Notes: Each figure reports the results for each FPC. The blue/red/green colors signify the first/second/third FPCs. The ordered loadings are on the horizontal axis and the values of the four standardized moments are shown on the vertical axis.

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.

C Three-Dimensional Impulse Response Function Surfaces of EIDs to Economic Shocks

Figure 16 shows three-dimensional IRF surfaces of one-year ahead (top row) and medium-run EIDs (bottom row) to monetary (blue), government spending (red), personal income tax (green), and gasoline price (black) shocks.

Figure 16. *Three-dimensional Impulse Response Function Surfaces of EIDs to Economic Shocks*

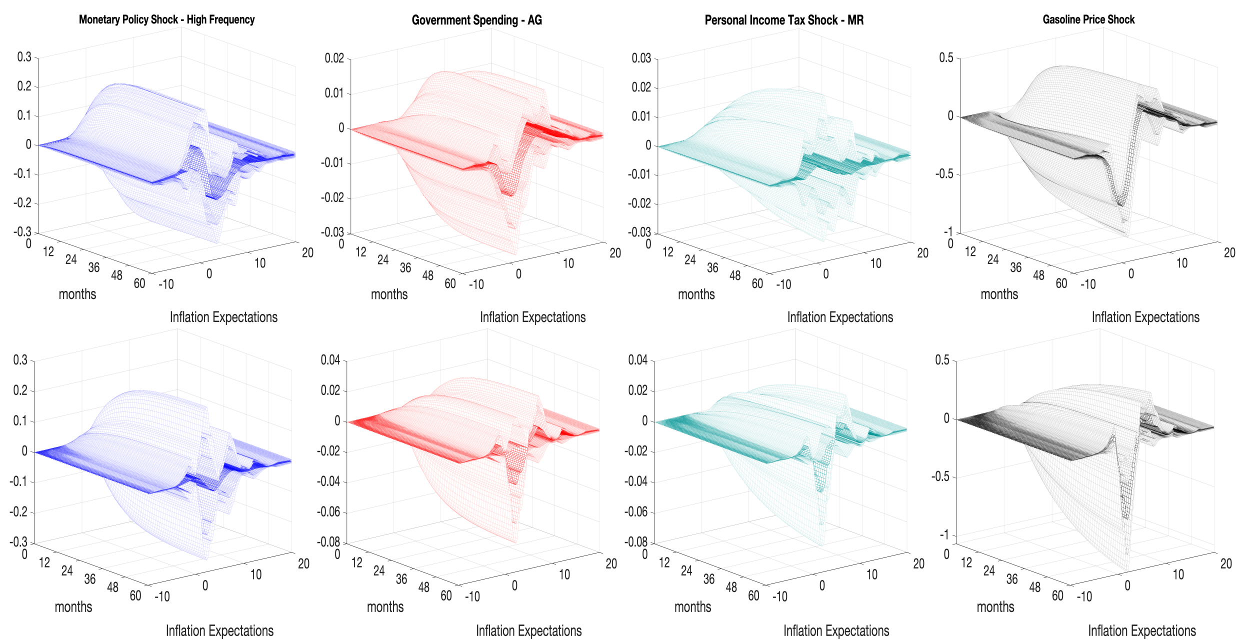

Notes: Top panel shows cumulative IRFs of EIDs of one-year inflation expectations. Bottom panel shows IRFs of EIDs of medium-run inflation expectations. IRFs to monetary policy shocks, government spending shocks, personal income tax shocks, and gasoline price shocks are represented in blue, red, green, and black, respectively. Horizon is in months.