CES ifo
Working Papers

13047
2026

October 2026

Energy Supply Shocks and Inflation Expectations: A Functional Var Approach

Maximilian Boeck, Anja Bauer, Jonas Dovern

CES ifo

Imprint:

CESifo Working Papers

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740

Email office@cesifo.de
https://www.cesifo.org

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge
· from the CESifo website: www.ifo.de/en/cesifo/publications/cesifo-working-papers
· from the SSRN website: www.ssrn.com/index.cfm/en/cesifo/
· from the RePEc website: https://ideas.repec.org/s/ces/ceswps.html

Energy supply shocks and inflation expectations: a functional VAR approach*

Maximilian Boeck
Vienna University of Economics and Business

Anja Bauer
Nuremberg Institute for Market Decisions

Jonas Dovern
Friedrich-Alexander-Universität Erlangen-Nürnberg

October, 2026

Abstract

This paper studies how energy supply shocks affect the entire cross-sectional distribution of household inflation expectations. Using monthly German survey data from 2003 to 2025, we estimate a functional vector autoregression that combines macroeconomic aggregates with the expectation distribution and identify oil supply news shocks and natural gas supply shocks using external instruments. Oil shocks induce a hump-shaped rightward shift of the distribution, increase disagreement, and disproportionately move probability mass toward very high expectations. Natural gas shocks generate faster but smaller responses. In both cases, households overestimate the inflationary effects of the shocks. We further document substantial heterogeneity across household types by splitting the survey sample based on age, gender, and other characteristics. Given evidence that higher moments of the expectation distribution matter for macroeconomic dynamics, monitoring only average expectations is likely to understate the business-cycle consequences of energy supply shocks.

Keywords: inflation expectations; functional VAR; energy; supply shocks.

JEL classification: E31, D84, C32, Q43, E52.

*Contact: Maximilian Boeck, Vienna University of Economics and Business. E-mail: maximilian.boeck@wu.ac.at. Anja Bauer, Nuremberg Institute for Market Decisions. E-mail: anja.bauer@nim.org. Jonas Dovern, Friedrich-Alexander-Universität Erlangen-Nürnberg. E-mail: jonas.dovern@fau.de. The authors thank Pia Heckl for helpful comments. The study has been partly funded by the Nuremberg Institute for Market Decisions as part of a project on "Forecasting and Analysis of Macroeconomic Expectations and Consumer Behavior".

1. Introduction

Macroeconomic expectations, especially those of private households, are heterogeneous, and this heterogeneity matters. A growing body of evidence shows that differences in beliefs across households shape differences in consumption, saving, and portfolio decisions (Roth and Wohlfart, 2020; Coibion et al., 2024; Chopra et al., 2025). At the aggregate level, the distribution of expectations feeds back into macroeconomic dynamics. Recent work demonstrates that accounting for the heterogeneity of inflation expectations substantially alters standard macroeconomic relationships, such as the Phillips curve (Meeks and Monti, 2023), and that variations in the higher moments of the distribution of inflation expectations influence aggregate prices, real activity, and interest rates (Ascari et al., 2025). Yet, little is known about how the entire distribution of expectations responds dynamically to macroeconomic shocks.

Recently, arguably the most important macroeconomic shocks in many advanced economies, particularly in Europe, have been energy supply shocks. In 2022, the Russian war against Ukraine substantially disrupted the global market for natural gas, because Russian exports to Europe shrank dramatically over a short period and were substituted by imports from other regions. European wholesale gas prices (Dutch TTF) increased from approximately €80 per MWh immediately before the Russian invasion to €340 per MWh in late August 2022, roughly seventeen times their pre-2021 average of about €20 per MWh. Likewise, the oil market experienced price increases from around $74 per barrel (Brent oil) in December 2021 to $117 per barrel in March 2022. More recently, the war between the U.S. and Iran halted shipping through the Strait of Hormuz, through which roughly a fifth of global oil and a comparable share of global LNG pass, causing monthly oil prices to rise from $67 per barrel in January 2026 to $117 per barrel in April, marking the largest quarterly increase in real oil prices in decades. Given that Qatar is one of the largest exporters of LNG globally, the war also caused another natural gas supply shock. Both episodes posed enormous challenges for monetary policymakers, who had to navigate the trade-off between stimulating sluggish economic growth and taming inflation and inflation expectations (Reis, 2023; Gagliardone and Gertler, 2026).

This paper studies how the distribution of inflation expectations of German households reacts to oil supply news shocks and gas supply shocks. We move beyond average expectations or specific moments, analyzing the full distributional response. Our contribution is twofold: First, we provide new empirical evidence on how different types of energy supply shocks propagate through the distribution of household inflation expectations. Second, we uncover systematic heterogeneity in this propagation across different types

of households. Our empirical results are informative for the assessment of models of expectation formation and relate to recent findings about the relevance of higher-moment movements in expectations (Meeks and Monti, 2023; Ascari et al., 2025).

Our analysis combines micro-level survey data with aggregate macroeconomic information within a unified framework. We use household-level expectation data to estimate the cross-sectional density of inflation expectations for each period. We summarize this high-dimensional object using functional data techniques and jointly model expectations with macroeconomic variables in a functional vector autoregression (FVAR) that we estimate using Bayesian methods. The information in the densities is compressed via functional principal component analysis, which we apply to the log-quantile density transformation of the densities. For identification, we rely on established external instruments to identify oil supply news shocks (Känzig, 2021; Mori and Peersman, 2026) and gas supply shocks (Alessandri and Gazzani, 2025). We then compute impulse response functions to trace the dynamic effects of these shocks on the entire expectation distribution. Finally, we split the cross-section into various subsamples across different types of households to examine potentially differential responses across groups.

Our descriptive analysis of the inflation expectation data shows that their distribution is very wide and that both its location and shape change considerably over time. Mean expectations were low between 2010 and 2019, rose rapidly after 2020, and fell sharply in 2023. Disagreement among households fluctuates substantially, with pronounced increases around 2008 and 2021. Moreover, the distribution became markedly right-skewed and fat-tailed following the global financial crisis and the European debt crisis. Despite this complex variation, seven functional principal components are sufficient to capture more than 95 percent of the variation in the log-quantile density transformation of the densities.

Our main results regarding the expectation effects of energy supply shocks can be summarized as follows: First, we document substantial effects on the higher moments of the expectation distribution for all supply shocks that we consider. Second, a contractionary oil supply news shock shifts the entire expectation distribution to the right in a hump-shaped manner, increases disagreement, and disproportionately moves probability mass toward very high expectations. Mean expectations rise by nearly 0.4 percentage points, while headline inflation increases by at most 0.2 percentage points. This implies that households, on average, overestimate the inflationary impact of the shock for two to three years. Third, after natural gas supply shocks, households adjust their inflation expectations more quickly than after oil supply shocks. However, given an equally large rise in energy prices, the effect is only about half as strong. Given the exceptional size of the

gas supply disruptions in 2022, gas supply shocks shifted mean expectations more than oil supply news shocks between January and August of that year (0.6 versus 0.4 percentage points). Fourth, the responses differ systematically across household types: younger households react more strongly to both shocks, and the peak response of rural households' mean expectations to an oil supply news shock is more than twice that of urban households. In contrast, the well-documented gender gap in the unconditional level of expectations does not carry over to gender differences in the dynamic expectation responses.

Our work contributes to the literature on the macroeconomic expectations of private households; Weber et al. (2022) and D'Acunto et al. (2023) provide recent surveys. This literature documents that household expectations deviate systematically from the full-information rational-expectations benchmark and from the expectations of professional forecasters.

More specifically, our paper relates to the literature on heterogeneity in households' inflation expectations. A first strand documents persistent differences in the *level* of expectations across demographic and socioeconomic groups, linking them to income and socioeconomic status (Madeira and Zafar, 2015), exposure to specific prices (D'Acunto et al., 2021a), gender (D'Acunto et al., 2021b), and lifetime inflation experiences (Malmendier and Nagel, 2016). A second strand treats the resulting disagreement – among professional forecasters and households alike – as an object of interest in its own right, documenting its variation over time and the information frictions that generate it (Patton and Timmermann, 2010; Dovern et al., 2012; Coibion and Gorodnichenko, 2015a; Tsiaplias, 2020). A third strand uses information-provision experiments embedded in expectation surveys to show that households update their beliefs heterogeneously, even when confronted with identical signals (Roth and Wohlfart, 2020; Coibion et al., 2024; Piccolo and Gorodnichenko, 2025), a finding commonly attributed to differences in attention (Maćkowiak and Wiederholt, 2026). We contribute to this literature in two ways. First, we estimate how the entire expectation distribution responds dynamically to cleanly identified energy supply shocks. Second, we show that the dynamic responses differ systematically across household types. This is complementary to the experimental evidence on heterogeneous updating and shifts the focus from differences in what households believe on average to differences in how they dynamically adjust their beliefs.[1]

[1] In two related papers, Ahn et al. (2024) investigate how renters' and homeowners' inflation expectations and labor market outlooks react differently to mortgage rate shocks, and Doh et al. (2025) investigate how quantiles of the inflation expectations of different socioeconomic groups react to monetary policy shocks.

Heterogeneity and its dynamics over time—in our case of inflation expectations—connect our paper to the literature on heterogeneous agent New Keynesian (HANK) models (Kaplan et al., 2018; Auclert et al., 2025). While this class of models stresses heterogeneity in income, wealth, and marginal propensities to consume, it typically retains the assumption of full-information rational expectations. Recent work relaxes this assumption, albeit mostly in ways that imply identical, if biased, beliefs across all households and, thus, no disagreement (Farhi and Werning, 2019). So far, only a few models generate truly heterogeneous expectations, for instance, through sticky information (Auclert et al., 2020) or noisy information in a HA model without price rigidities (Overhage, 2026). Our results on how the entire distribution of inflation expectations responds to aggregate supply shocks provide empirical moments that may prove useful for calibrating the next generation of HANK models featuring heterogeneous expectations.

Naturally, the paper also contributes to the literature on the effects of energy supply shocks on inflation expectations. Thus far, the literature has found mixed results regarding the pass-through of such shocks to inflation expectations in the cases of oil (Peersman and Van Robays, 2009; Wong, 2015; Kilian and Zhou, 2022a; Kilian and Zhou, 2022b; Aastveit et al., 2023) and natural gas (Boeck and Zörner, 2025; Alessandri and Gazzani, 2025; Adolfsen et al., 2026), or both in a joint framework (Casoli et al., 2024); generally, it finds only moderate effect sizes.[2] However, Coibion and Gorodnichenko (2015b) show that the effect of oil prices on private households' inflation expectations is important for aggregate inflation dynamics. What all studies have in common, however, is that they characterize the pass-through in terms of a single summary of inflation expectations, typically a market-based break-even rate or the mean of a survey, even though the anchoring aspect concerns the entire distribution of expectations rather than only its center.[3] Our contribution is to estimate how the entire cross-sectional density of household inflation expectations responds to identified energy supply shocks. This allows us to assess directly whether moderate average responses mask a substantial reallocation of probability mass toward extreme beliefs.

Regarding our modeling approach, we draw heavily on the statistical literature on functional data. Horváth and Kokoszka (2012) provide an excellent overview of this literature. In the macroeconomic literature, tools for functional data are relatively new. An early example is Tsay (2016), who models the variation in the distribution of stock returns and

[2] Related papers document the effects of carbon price shocks on the inflation expectations of households (Frank, 2026), firms (Hensel et al., 2024), and financial market participants (Bauer et al., 2026), which work through higher energy prices.

[3] Chang et al. (2022) estimate the effects of gasoline price changes on the distribution of inflation expectations; however, they do so without properly identifying oil supply (or demand) shocks.

the income distribution using functional autoregressive models. FVARs are developed by Chang et al. (2024) to trace aggregate shocks on cross-sectional distributions. Their connection to HANK models is discussed in Huber et al. (2026a). For the finite-dimensional repeated cross-sectional data that we use, it is important that the model accounts for the fact that only a sample of observations from the distribution to be modeled is available for each time period. Whereas Chang et al. (2024) use cubic splines to approximate cross-sectional densities, we follow Huber et al. (2026b) in using functional principal components (FPCs) as proposed in Kneip and Utikal (2001). FVARs have been used to study the impact of uncertainty shocks (Huber et al., 2026b), technology shocks (Chang et al., 2024), and monetary policy shocks (Chang and Schorfheide, 2026) on the earnings and consumption distributions in the U.S., as well as how the distribution of inflation expectations reacts to monetary policy, government spending, taxation, and gasoline price shocks (Chang et al., 2022); FVARs are also used to assess the effects of oil demand and supply shocks on the distribution of U.S. stock market returns (Bjørnland et al., 2026). While all of these papers investigate the impact of conventionally identified macroeconomic shocks *on* distributions, Inoue and Rossi (2021) identify monetary policy shocks *from* exogenous shifts in the distribution of government bond yields across different maturities. An alternative to FVARs has recently been proposed by Koop et al. (2026) in the form of large-scale pseudo VARs that include a vector of individual expectations for "pseudo individuals" derived from the repeated-cross-section data instead of an approximation of the expectation density.

Our paper proceeds as follows. In Section 2, we document the data used for our analysis. This includes the household-level expectation data and macroeconomic data used to identify the aggregate supply shocks. Section 3 outlines the methodological approaches we use to estimate the distribution of inflation expectations and evaluate how this distribution responds to aggregate supply shocks. Section 4 documents our empirical results. It provides descriptive results about the dynamics of the distribution of inflation expectations in Germany, the core findings regarding the dynamic response of this distribution to oil supply news shocks and gas supply shocks, an account of how much the two shocks contributed to the inflation surge and expectation dynamics in 2022, and it describes the extent to which expectations of different types of households respond differently to these aggregate shocks. Finally, Section 5 concludes.

2. Data

2.1 Inflation Expectation Data

Our data on inflation expectations of private households in Germany come from the "NIM Consumer Climate *powered by GfK*", which is co-funded by the European Commission as part of the European Business and Consumer Surveys (BCS). It is a monthly repeated cross-sectional survey of approximately 2,000 households, published by the Nuremberg Institute for Market Decisions (NIM). The sample is selected using a quota procedure that ensures a distribution representative of the German-speaking population in Germany by considering age, gender, region, household size, and other socioeconomic characteristics. For each monthly wave, the survey period covers the first two weeks of the month. Until April 2019, the survey was conducted through face-to-face, computer-assisted personal interviews. Starting in May 2019, the mode of interviewing was converted to an online survey to achieve greater flexibility and efficiency; it has since been conducted via the GfK online panel "GfK eBUS". This change brings several methodological improvements, including enhanced adaptability, the integration of complex questionnaires, faster response times, and increased sample diversity. However, it introduces a potential structural break in our data. Therefore, we include suitable dummy variables in our empirical analyses to capture these potential breaks.

The survey questions most relevant to our analysis elicit twelve-month-ahead inflation expectations. First, there is a qualitative question regarding the future development of inflation. In particular, participants are asked:

> How, in your opinion, will consumer prices develop over the next 12 months compared with the last 12 months?
> They will ...
> 1: increase more strongly than before,
> 2: increase by about the same amount as before,
> 3: increase less strongly than before,
> 4: remain about the same,
> 5: decrease,
> 9: don't know

If participants select 1, 2, 3, or 5, they receive a follow-up question regarding their quantitative inflation expectations:

> By what percentage do you think consumer prices will rise or fall over the next 12 months?

The survey design accepts answers with up to 3 digits before the decimal point and one decimal place. There is an "I-do-not-know" option.

For our analysis, we assume that households indicating prices will remain about the same have inflation expectations of 0 percent. All of our results are robust to the exclusion of these households from the sample. We exclude all households that selected the I-do-not-know option for either of the two questions. The dataset covers May 2003 to December 2025 and contains, on average, 2,028 observations per period (SD: 87.27). In total, we have around 551,000 observations. For the months from August to October 2007, inflation expectation data are missing. For each of these three months, we simulate observations from a non-normal distribution by using interpolated moments for the mean, variance, skewness, and kurtosis.[4] We make two adjustments to the sample. First, we remove any missing values. Second, we trim the dataset and remove any inflation expectations below $L = -25$ and above $U = 55$. This second step only removes about 4,000 observations from the total sample. Overall, this results in an average of 1,636 observations per wave (SD: 107.49) over the sample period from May 2003 to December 2025. The smallest sample contains 854 observations in January 2006. The largest sample contains 1,933 observations in March 2008. This large cross-sectional dimension of our dataset allows for a good approximation of the expectation distribution. To align our sample with the German-speaking population, we apply the survey weights provided by the NIM, which adjust for quota deviations and align the sample distributions of gender, age, household-size, city size, and education with the corresponding population distributions.

Figure 1 shows how the distribution of inflation expectations evolves over time.[5] Two observations stand out: First, the center of the distribution moves substantially, indicating pronounced shifts in average expectations over time. Second, the distributions are very wide, and the variance changes substantially over time, especially widening around 2008 and 2021. Since we transform the data according to equation (7) before we estimate the FPCs to be included in the FVAR, we show the same data in terms of the variation of the log-quantile density (LQD) over time in Figure A1 in the appendix.

Figure 2 shows the evolution of the mean, variance, skewness, and kurtosis of the distribution of inflation expectations across all waves of our sample. It confirms the pronounced shifts in average expectations, particularly the low expectations between 2010

[4] In particular, we perform linear interpolation for the mean, variance, skewness, kurtosis, and the standard error of the mean. Using the interpolated variance and standard error, we can compute the implied sample size via $\hat{N}_t = \left(\sqrt{Var(x_t)}/se(x_t)\right)^2$. Given these objects, we simulate a non-normal continuous variable for each missing month using the method by Fleishman (1978).

[5] Based on the cross-sections of survey-based inflation expectations, we use a kernel density estimator to approximate the distribution of inflation expectations for each month in our sample. Details are presented in Section 3 below.

Figure 1: Distribution of inflation expectations over time.

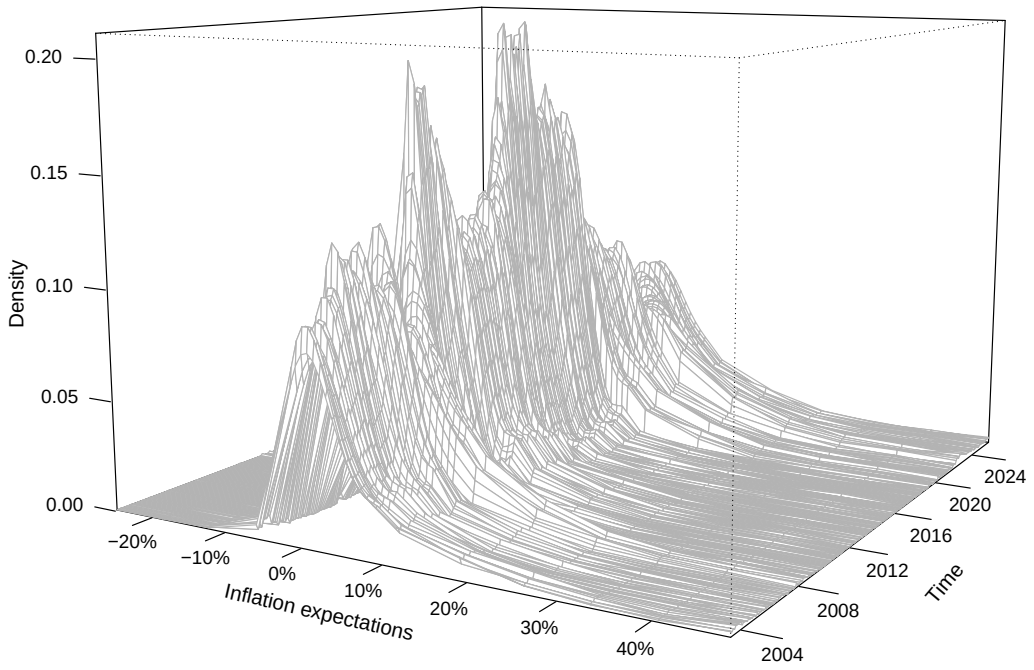

Notes: The figure shows estimates of the distribution of inflation expectations for each survey wave. The densities are computed month-by-month using a Gaussian kernel density estimator with bandwidth size determined according to Silverman's rule of thumb. Boundary correction is applied by augmenting the data with their reflections near the boundaries of $\mathscr{X}$. For August to October 2007 data are interpolated as described in the text.

and 2019, along with the rapid increase in expectations after 2020, followed by a sharp decline in 2023. Changes in the width of the expectation distribution are substantial, as indicated by the large fluctuations in variance. This variation in disagreement over time confirms previous evidence (Dovern et al., 2012). The changes in skewness and kurtosis reveal higher-order shifts in the distribution of inflation expectations that are often overlooked. The distribution became particularly right-skewed and fat-tailed after the financial crisis of 2008 and the European debt crisis a few years later. Note that a substantial part of the shifts observed in 2019 is likely due to the change in survey methodology highlighted above. We account for this possibility in our analysis below by including suitable dummy variables in our time-series models.

Figure 2: Moments of the distribution of inflation expectations.

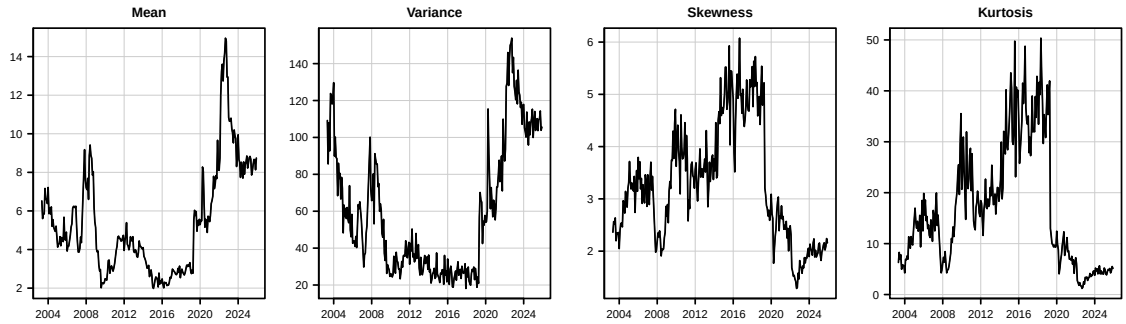

Notes: The figure shows the first four moments of the distribution of inflation expectations for each survey wave. They are computed based on density estimates obtained using a Gaussian kernel density estimator with bandwidth size determined according to Silverman's rule of thumb. Boundary correction is applied by augmenting the data with their reflections near the boundaries of $\mathscr{X}$. For August to October 2007 data are interpolated as described in the text.

In addition to the expectation questions, the questionnaire contains a section on the personal characteristics of the respondent and the household in which they live. In particular, the "NIM Consumer Climate *powered by GfK*" provides information on the respondents' age, gender, family status, household composition, income, location of residence, living conditions, and education. Below, we make use of these data to analyze whether there are any heterogeneities in the response of expectations to macroeconomic supply shocks across household types.

2.2 Macroeconomic Data

In addition to the microeconomic expectation data, we require a set of macroeconomic variables to identify the aggregate shocks and estimate the dynamic response of expectations to these shocks. In particular, we use the real oil and natural gas prices, which we deflate using the harmonized consumer price index in the euro area. For the German macroeconomy, we use data on German industrial production, consumer price inflation (π_t), energy prices, and the stock market. We list the data sources and their transformations in the VAR in Appendix A.

3. Methodological approach

In this section, we first show how we jointly model the relationship between the cross-sectional density of inflation expectations and macroeconomic aggregates. Our econometric framework builds upon Chang et al. (2024) and Huber et al. (2026b). After outlining this general framework, we explain how we approximate the cross-sectional dynam-

ics using a functional principal component approach. Lastly, we discuss the estimation of our object of interest: distributional impulse response functions.

3.1 A functional vector autoregression

We model the joint dynamics of an $n_y \times 1$ vector of macroeconomic aggregates $\boldsymbol{y}_t$ and a cross-sectional continuous density function $p_t(x)$, with $x \in \mathscr{X}$ in the domain $\mathscr{X}$ of the expectation distribution and a corresponding cumulative distribution function $F_t(x)$. We do not directly observe $p_t(x)$ over time. In each period, we observe a quota-selected survey sample of size N_t, with associated survey weights, which we collect in the $N_t \times 1$ vector $\boldsymbol{x}_t$. We estimate the target cross-sectional distribution using the weighted kernel density estimator described below. We denote the time-series dimension of our sample as T. Let $f_t(z) = g\{p_t(\cdot)\}(z)$ denote an invertible transformation of the density function on the support $z \in [0,1]$. Since we are interested in fluctuations around the steady state distribution, we express $\tilde{f}_t(\cdot) = f_t(\cdot) - \bar{f}$ as deviations from the sample mean of $f_t(\cdot)$, defined as $\bar{f} = T^{-1}\sum_{t=1}^{T} f_t$.

We assume a functional vector autoregression (FVAR) with p lags as

$$\begin{pmatrix} \boldsymbol{y}_t \\ \tilde{f}_t(z) \end{pmatrix} = \begin{pmatrix} \boldsymbol{c}_y \\ c_f(z) \end{pmatrix} + \sum_{\ell=1}^{p} \begin{pmatrix} \boldsymbol{B}_{\ell,yy}\boldsymbol{y}_{t-\ell} + \int \boldsymbol{B}_{\ell,yf}(\hat{z})\tilde{f}_{t-\ell}(\hat{z})\, d\hat{z} \\ \boldsymbol{B}_{\ell,fy}(z)\boldsymbol{y}_{t-\ell} + \int B_{\ell,ff}(z,\hat{z})\tilde{f}_{t-\ell}(\hat{z})\, d\hat{z} \end{pmatrix} + \begin{pmatrix} \boldsymbol{u}_{y,t} \\ u_{f,t}(z) \end{pmatrix}, \tag{1}$$

where $z, \hat{z} \in [0,1]$. $\boldsymbol{c}_y$ denotes a vector of intercepts for the macroeconomic aggregates, and $c_f(z)$ is a function-valued intercept for the density equation. The reduced-form errors $\boldsymbol{u_t} = [\boldsymbol{u}'_{y,t}, u_{f,t}(z)]'$ are innovations with zero mean and a covariance matrix given by

$$\boldsymbol{\Sigma}(z,\hat{z}) = \begin{pmatrix} \boldsymbol{\Sigma}_{yy} & \boldsymbol{\Sigma}_{yf}(\hat{z}) \\ \boldsymbol{\Sigma}_{fy}(z) & \Sigma_{ff}(z,\hat{z}) \end{pmatrix} \tag{2}$$

that accommodates non-zero contemporaneous correlations between macroeconomic variables and the distribution of expectations. This specification assumes that macroeconomic aggregates, summarized by $\boldsymbol{y}_t$, depend on the cross-sectional density both contemporaneously and dynamically. The same holds for the cross-sectional density, which depends on macroeconomic aggregates both contemporaneously and dynamically. Structural macroeconomic shocks can impact the cross-sectional distribution both directly and through dynamic inter-dependencies.

Since the functions $\tilde{f}_t(\cdot)$ and $u_{f,t}(\cdot)$ are continuous, the dimension of the model is infinite. In order to estimate the model, we resort to a finite-dimensional representation.

The Karhunen-Loéve theorem (see, e.g., Bosq (2000), Theorem 1.5) expresses $\tilde{f}_t(\cdot)$ as the infinite-order expansion

$$\tilde{f}_t(z) = \sum_{k=1}^{\infty} \zeta_k(z)\alpha_{k,t}, \tag{3}$$

with $\zeta_k(z)$ denoting time-invariant orthogonal functional bases and $\alpha_{k,t}$ time-varying factors. Conditional on choosing an appropriate truncation level q, we can approximate equation (3) through

$$\tilde{f}_t(z) \approx \sum_{k=1}^{q} \zeta_k(z)\alpha_{k,t} = \boldsymbol{\zeta}(z)'\boldsymbol{\alpha}_t, \tag{4}$$

where $\boldsymbol{\zeta}(z) = (\zeta_1(z), \ldots, \zeta_q(z))'$ and $\boldsymbol{\alpha}_t = (\alpha_{1,t}, \ldots, \alpha_{q,t})'$ are two $q \times 1$ vectors. In a similar fashion, we can approximate the remaining objects with infinite dimensions in equation (1) as $c_f(z) \approx \boldsymbol{\zeta}(z)'\boldsymbol{c}_\alpha$, $u_{f,t}(z) \approx \boldsymbol{\zeta}(z)'\boldsymbol{u}_{\alpha,t}$, $\boldsymbol{B}_{\ell,fy}(z) \approx \boldsymbol{\zeta}(z)'\boldsymbol{B}_{\ell,\alpha y}$, $\boldsymbol{B}_{\ell,yf}(\hat{z}) \approx \boldsymbol{B}_{\ell,y\alpha}\boldsymbol{\delta}(\hat{z})$, and $B_{\ell,ff}(z,\hat{z}) \approx \boldsymbol{\zeta}(z)'\boldsymbol{B}_{\ell,\alpha\alpha}\boldsymbol{\delta}(\hat{z})$. Both $\boldsymbol{\zeta}(z)$ and $\boldsymbol{\delta}(\hat{z})$ denote a q-dimensional vector of time-invariant orthogonal functional bases.

By using this finite-dimensional approximation, we can re-write the infinite-dimensional FVAR in equation (1) as an $n = n_y + q$ finite-dimensional FVAR

$$\begin{pmatrix} \boldsymbol{y}_t \\ \boldsymbol{\alpha}_t \end{pmatrix} = \begin{pmatrix} \boldsymbol{c}_y \\ \boldsymbol{c}_\alpha \end{pmatrix} + \sum_{\ell=1}^{p} \begin{pmatrix} \boldsymbol{B}_{\ell,yy} & \boldsymbol{B}_{\ell,y\alpha}\boldsymbol{C}_\alpha \\ \boldsymbol{B}_{\ell,\alpha y} & \boldsymbol{B}_{\ell,\alpha\alpha}\boldsymbol{C}_\alpha \end{pmatrix} \begin{pmatrix} \boldsymbol{y}_{t-\ell} \\ \boldsymbol{\alpha}_{t-\ell} \end{pmatrix} + \boldsymbol{\Pi}\boldsymbol{d}_t + \begin{pmatrix} \boldsymbol{u}_{y,t} \\ \boldsymbol{u}_{\alpha,t} \end{pmatrix}, \tag{5}$$

where $\boldsymbol{C}_\alpha = \int \boldsymbol{\delta}(\hat{z})\boldsymbol{\zeta}(\hat{z})' d\hat{z}$ and $\boldsymbol{u}_t = (\boldsymbol{u}'_{y,t}, \boldsymbol{u}'_{\alpha,t})'$ has zero mean and an $n \times n$ covariance matrix $\boldsymbol{\Sigma}$. This covariance matrix is a full matrix. Conditional on knowing $\boldsymbol{\alpha}_t$, the model in equation (5) is a standard linear VAR, resembling a factor-augmented VAR proposed by Bernanke et al. (2005). Note that we add $n_e = 6$ exogenous dummies for the Covid months (one dummy per month from March to July 2020) and for the change in the survey methodology (one dummy from May 2019 onward) in the $n_e \times 1$ vector $\boldsymbol{d}_t$ with the corresponding $n \times n_e$ coefficient matrix $\boldsymbol{\Pi}$. We estimate the model using Bayesian techniques with an independent Gaussian and inverse-Wishart prior for the VAR coefficients and innovation variances, respectively. We calibrate the Gaussian variances using a Minnesota specification (Sims, 1980; Doan et al., 1984; Karlsson, 2013). In the empirical analysis, we use $p = 12$ lags. We discuss further details on the prior specification and the Bayesian estimation of the linear VAR in Appendix B.

Structural identification. To transform this model into a structural FVAR, we exploit that the covariance matrix can be decomposed using $\boldsymbol{\Sigma} = \boldsymbol{S}\boldsymbol{S}'$, where $\boldsymbol{S}$ denotes the structural impact matrix. This implies the usual relationship between the reduced-form errors and the structural innovations as $\boldsymbol{u}_t = \boldsymbol{S}\boldsymbol{\varepsilon}_t$, where the structural error term $\boldsymbol{\varepsilon}_t$ is normally distributed with zero mean and unit variance: $\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_n)$. For the identification of energy supply shocks, we rely on external instruments. Hence, we partition the vector

of structural shocks as $\boldsymbol{\varepsilon}_t = (\varepsilon_t^{\text{ES}}, \boldsymbol{\varepsilon}_t^{-\text{ES}\prime})'$ (w.l.o.g.), where $\varepsilon_t^{\text{ES}} \in \mathbb{R}$ denotes the respective energy supply shock, and $\boldsymbol{\varepsilon}_t^{-\text{ES}} \in \mathbb{R}^{n-1}$ denotes all other unidentified shocks.

We identify two shocks: oil supply news shocks and natural gas supply shocks. To identify oil supply news shocks, we rely on the external instrument by Känzig (2021) and the refinement in Mori and Peersman (2026).[6] The approach utilizes high-frequency variation in oil price futures around production announcements by the Organization of the Petroleum Exporting Countries (OPEC). This isolates the impact of those announcements, as the pre-announcement prices incorporate the prevailing state of the global economy. The interpretation as a supply shock is based on the plausible assumption that OPEC's announcements do not contain new information about demand conditions. The upper panel of Figure C2 in the appendix shows the obtained sequence of posterior densities for the oil supply news shocks, where the external instrument is used to identify this structural shock.

In order to identify natural gas supply shocks, we use the instrument provided by Alessandri and Gazzani (2025).[7] Their instrument relies on selecting days with major changes in the price of gas futures and a careful analysis of news to identify daily price changes that are clearly driven by supply factors, such as current or future gas imports and changes in the risks of future supply disruptions. The monthly instrument z_t is then the average of all daily price changes related to supply-side news in that month, while it is set to zero if there are no days with exclusively supply-side news in a particular month. The lower panel of Figure C2 in the appendix shows the obtained sequence of posterior densities for the natural gas supply shocks, where the external instrument is used to identify this structural shock

Distributional impulse response functions. Following Chang et al. (2024), we assume that the system is at its steady state before the shock hits. At steady state, the LQD function corresponds to $f_{ss}(z) = \boldsymbol{\zeta}(z)'\boldsymbol{\alpha}_{ss} + \bar{f}$, where $\boldsymbol{\alpha}_{ss}$ corresponds to the unconditional mean implied by the model parameters. The steady state distribution can be found by applying the inverse transformation $p_{ss}(x) = g^{-1}\{f_{ss}(\cdot)\}(x)$.

To compute the distributional impulse response functions, we use the impulse responses of $\boldsymbol{\alpha}_t$ to a structural shock of interest by standard methods. We denote the impulse response of the factors as $\tilde{\boldsymbol{\alpha}}_{j,h} = \mathbb{E}[\alpha_{t+h}|I_{t-1}, \varepsilon_{t,j} = 1, \boldsymbol{\varepsilon}_{t,-j} = 0] - \mathbb{E}[\alpha_{t+h}|I_{t-1}, \boldsymbol{\varepsilon}_t = \mathbf{0}]$, defined as the difference between expected factor values conditional on a unit shock and on no

[6] The argument made by Mori and Peersman (2026) is that the instrument suffers from an invertibility issue, as it is predictable using stock market information. By additionally controlling for this forward-looking information, the instrument yields stable and reliable results over time. We use the series provided on the authors' website.

[7] We use the shock series extending to December 2025, as provided by the authors on their website.

shock, holding the pre-shock information set I_{t-1} fixed. These can be used to express the LQD function h periods after the shock, which we define as $f_{ss+h}(z) = \boldsymbol{\zeta}(z)'(\boldsymbol{\alpha}_{ss} + \tilde{\boldsymbol{\alpha}}_{j,h}) + \bar{f}$. Using the inverse transformation, we back out the corresponding distribution of interest, $p_{ss+h}(x) = g^{-1}\{f_{ss+h}(\cdot)\}(x)$. The distributional impulse response functions we report are then computed as the difference between $p_{ss+h}(x)$ and $p_{ss}(x)$. Similarly, when we report how moments or quantiles change due to a structural shock, we report the difference between the moment/quantile of the distribution after h periods and the steady state distribution.

3.2 A finite-dimensional representation of a distribution

We have been silent on three issues: (i) how to approximate the cross-sectional density, $p_t(x)$, (ii) which invertible transformation of the density function, $f_t(z) = g\{p_t(\cdot)\}(z)$, we use, and (iii) what kind of orthogonal basis functions are contained in $\boldsymbol{\zeta}(\cdot)$ and $\boldsymbol{\delta}(\cdot)$. We approximate the cross-sectional density using a kernel density estimator (KDE). For the transformation, we use the log-quantile density (LQD) transformation proposed by Petersen and Müller (2016). Regarding the choice of the functional basis, we resort to functional principal components (FPCs) as described in Tsay (2016) and Huber et al. (2026b).[8]

Our approach offers several advantages. Generally, we introduce relatively few parametric restrictions. The kernel density approach allows us to estimate the cross-sectional density. Probability density functions, however, are not unconstrained random functions in the Hilbert space L^2, because they are constrained by non-negativity and unit-integration. As they constitute only a subset of L^2, this leads to problems when applying common techniques from functional data analysis (Ramsay and Silverman, 2005). We use the LQD mapping associated with the density function. The LQD representation enables functional principal component analysis without directly imposing the non-negativity and unit-integration constraints on densities. After approximation, a normalized reconstruction enforces the prescribed support and yields valid densities, allowing us to compute distributional impulse response functions.

We evaluate the LQD functions on a sufficiently fine grid to approximate their functional principal components numerically. We then use the estimated component scores directly in the VAR, treating them as observed in the second-stage estimation. This follows the two-step principal-components approach discussed by Stock and Watson (2016). In our functional setting, grid refinement improves numerical accuracy but does not by itself

[8] An alternative is to consider splines to approximate the cross-sectional distribution, as done by Chang et al. (2024). While their approach is flexible and non-parametric, the number of truncation points (i.e., the number of spline knots) needed can easily become very large, which undermines inference in the VAR.

eliminate uncertainty arising from density and component estimation. The approach is computationally convenient, although the resulting posterior inference is conditional on the estimated densities and component scores and does not propagate their estimation uncertainty.

Kernel density estimation. We observe the cross-sectional distribution as a sample $x_{it} \in \mathscr{X}$ $(i = 1, \dots, N_t)$ together with survey weights ω_{it}, which we use to estimate the population distribution $p_t(x)$. We normalize the weights to sum to unity, i.e., $\tilde{\omega}_{it} = \omega_{it} / \sum_{i=1}^{N_t} \omega_{it}$. We specify $n_x = 1000$ grid points $g_j^x \in [L, U]$, where L and U denote the lower and upper bounds of the support. We use a Gaussian kernel density estimator, specifying the bandwidth size h_t according to Silverman's rule of thumb.[9] We apply the boundary correction to the kernel density estimator by augmenting the data with their reflections near the boundaries. Specifically, our estimator is

$$\hat{p}_t(x) = \frac{1}{h_t \sum_{i=1}^{N_t} \tilde{\omega}_{it}} \sum_{i=1}^{N_t} \tilde{\omega}_{it} \Big[K\Big(\frac{x - x_{it}}{h_t}\Big) + K\Big(\frac{x - x_{it}^L}{h_t}\Big) + K\Big(\frac{x - x_{it}^U}{h_t}\Big) \Big], \tag{6}$$

where $x_{it}^L = 2L - x_{it}$ and $x_{it}^U = 2U - x_{it}$. $K(w) = (2\pi)^{-1/2} \exp(-w^2/2)$ denotes the Gaussian kernel. We evaluate the density on the fixed grid, $p_{jt} = \hat{p}_t(g_j^x)$, which leaves us with a balanced panel of size $T \times n_x$.

Log-quantile density transformation. The transformation is defined as the logarithm of the first derivative of the quantile function $Q_t(z) = F_t^{-1}(z)$:

$$f_t(z) = g\{p_t(\cdot)\}(z) = \log\Big(\frac{d}{dz} Q_t(z)\Big) = \log Q_t'(z) = -\log p_t(Q_t(z)), \qquad z \in [0,1]. \tag{7}$$

Regarding the inverse mapping, which will be necessary to obtain the IRFs of $p_t(x)$, we note that $Q_t'(z) = \exp(f_t(z))$; therefore, we simply need to find $Q_t(z) = Q_t(0) + \int_0^z Q_t'(s)\,ds$. For an exact LQD function, the boundary conditions imply $\int_0^1 \exp(f_t(z))\,dz = U - L$. Approximation and model-implied changes in the transformed function need not preserve this constraint. We therefore use a normalized reconstruction that ensures the reconstructed quantile function satisfies $Q_t(0) = L$ and $Q_t(1) = U$, restoring the prescribed support $[L, U]$. The corresponding reconstructed density is

$$p_t(Q_t(z)) = \frac{1}{Q_t'(z)} = \frac{\int_0^1 \exp(f_t(u))\,du}{(U - L)\exp(f_t(z))} \tag{8}$$

For the numerical evaluation, we construct a grid $g_j^z \in [0,1]$ $(j = 1, \dots, n_z)$. We evaluate the LQD as $f_{jt} = f_t(g_j^z)$, resulting in a balanced panel of size $T \times n_z$. To match the specification of the FVAR above, we can compute $\tilde{f}_{jt} = f_{jt} - \bar{f}_j$ as deviations from the

[9] We use an adjusted weighted Silverman bandwidth, which we define as $h_t = 1.06 \cdot \sqrt{\sum_{i=1}^{N_t} \tilde{\omega}_{it}(x_{it} - \mu_{w,t})^2} \cdot N_{\text{eff,t}}^{-1/5}$. Here, $N_{\text{eff},t} = \frac{(\sum_{i=1}^{N_t} \tilde{\omega}_{it})^2}{\sum_{i=1}^{N_t} \tilde{\omega}_{it}^2}$ and $\mu_{w,t} = \sum_{i=1}^{N_t} \tilde{\omega}_{it} x_{it}$.

time-mean, defined as $\bar{f}_j = T^{-1}\sum_{t=1}^{T} f_{jt}$. We refer to Appendix B for more details on the LQD transformation.

Functional principal component analysis. For the orthogonal basis functions contained in $\boldsymbol{\zeta}(\cdot)$ and $\boldsymbol{\delta}(\cdot)$, we rely on FPC analysis described in Kneip and Utikal (2001) and Tsay (2016). We use FPC analysis to summarize a large part of the variation in $\tilde{f}_t(\cdot)$ through a small number of factors $\boldsymbol{\alpha}_t$. Let $\boldsymbol{F}$ be a $T \times n_z$ matrix with elements $\tilde{f}_{jt}$. Applying singular value decomposition (SVD) yields

$$\boldsymbol{F} = \boldsymbol{U}\boldsymbol{D}\boldsymbol{V}' + \boldsymbol{E},$$

where, $\boldsymbol{U}$ is a $T \times q$ real-valued matrix with orthonormal columns containing the first q left singular vectors, $\boldsymbol{D}$ is a $q \times q$ diagonal matrix with the first q largest singular values on its main diagonal, and $\boldsymbol{V}$ is a $n_z \times q$ real-valued matrix with orthonormal columns containing the first q right singular vectors. Lastly, the $T \times n_z$ matrix $\boldsymbol{E}$ measures the approximation error arising from using only the first q singular vectors. This allows for the computation of the first q FPCs as $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{D}$, with factors $\boldsymbol{\alpha}_t = \boldsymbol{A}_t'$ being represented by the rows of $\boldsymbol{A}$. The associated orthogonal basis functions are the q columns of $\boldsymbol{V}$.

The key takeaway from the FPC analysis is that it allows us to summarize cross-sectional information in a few $q << n_z$ factors. If the distribution we wish to approximate exhibits complex variation, the number of factors needed to approximate $\tilde{f}_t(\cdot)$ increases. In practice, we resort to the following strategy to decide on the number of factors: we truncate the number of factors when the cumulative share of explained variation in the observed cross-section exceeds a desired threshold.[10]

4. Empirical results

In this section, we summarize and discuss our empirical results. We first focus on the approximation of the distribution of inflation expectations using FPC analysis. We then present the full-sample results for the oil supply news shock and the corresponding results for the natural gas supply shock, followed by a discussion of the economic implications of our results. Then we look into the quantitative relevance of energy supply shocks in 2022. At the end of the section, we present results that highlight the heterogeneity in the dynamic responses of expectations to the two shocks across different types of households and discuss the connections of these results to the existing literature.

[10] Alternative strategies include the information criteria of Bai and Ng (2002), performing cross-validation as suggested by Tsay (2016), or minimizing the Kullback-Leibler divergence between the estimate obtained by KDE and the approximate distribution described by the FPCs.

Figure 3: Cumulative explained share of the variation by FPCs.

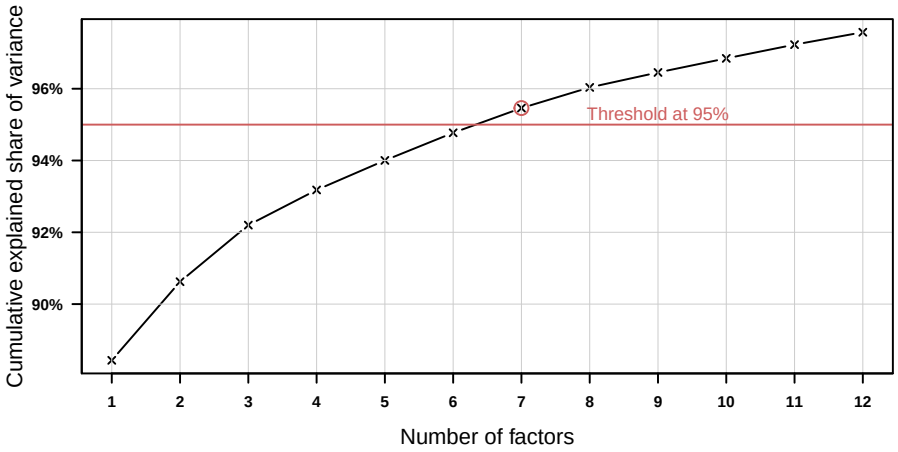

Notes: The figure shows how much of the cumulative share of the variation in the log-quantile density transformation of the distribution of inflation expectations is explained by a particular number of functional principal components (FPCs). We apply the threshold of 95 % to determine how many FPCs we include in the FVAR.

4.1 Approximation of the distribution of inflation expectations

We start by describing the approximation through the finite expansion in equation (4), which involves the choice of the number of functional principal components q. Choosing q too small results in an inaccurate approximation of the density and its movements over time, potentially leading to misleading estimates of its response to aggregate shocks. Choosing q too large results in an FVAR with more parameters than necessary, which may cause excessively wide error bands for IRFs. To decide upon this number, we use the following criterion: the optimal number of factors has to explain at least a certain threshold of the cumulative explained share of the variation. We set this threshold to 95 % in the baseline.

Figure 3 shows the cumulative share of variation in $\boldsymbol{F}$ (consisting of elements $\tilde{f}_{jt}$) explained by a specific number of FPCs. Based on these results, we choose $q = 7$ for our baseline estimations because it is the smallest number of factors that explains at least 95 percent of the variation. We note that none of our findings are particularly sensitive to this choice. In Appendix C, Figure C1 shows the first twelve FPCs of the baseline model. The approximate distribution based on the first seven FPCs captures the time dynamics of different features of the distribution well. The correlations of the first four moments with the implied corresponding approximate moments are very high ($\rho_{\text{mean}} = 0.913$, $\rho_{\text{var}} = 0.997$, $\rho_{\text{skew}} = 0.972$, and $\rho_{\text{kurtosis}} = 0.973$, with p-values all below 0.001).

4.2 Oil supply news shock

In this section, we examine how the inflation expectation distribution responds to the oil supply news shock. We begin by assessing the plausibility of the identified oil supply news shock through the responses of the macroeconomic aggregates, as shown in the first row of Panel (a) in Figure 4. A contractionary supply shock triggers a sharp, on-impact increase in the real oil price that gradually reverts over the five-year horizon. German energy price inflation rises by almost 1.5 percentage points, with the response building over the first six months and then slowly dissipating, consistent with the gradual pass-through of crude oil costs to retail energy prices. Broader consumer prices respond positively, albeit more modestly; the effect on headline inflation peaks at approximately 0.2 percentage points and remains elevated for two years, consistent with the standard cost-push channel. Industrial production contracts slightly, with the peak decline of almost 0.5 percent reached after two to three years, reflecting the negative supply-side effect on economic activity. Following Mori and Peersman (2026), we include stock prices in the set of endogenous variables as well to avoid possible identification issues. Stock prices decline persistently by more than 2 percent, consistent with a downward revision of firm valuations in response to higher energy costs and weaker growth prospects. Overall, these responses are qualitatively and quantitatively consistent with the findings of Känzig (2021) and confirm that the identification strategy successfully isolates a plausible contractionary oil supply news shock.

The second row of Panel (a) shows the responses of the first four moments of the expectation distribution, as implied by the IRFs of the FPCs included in the FVAR, as well as the dynamics of the forecast error implied by the reaction of inflation expectations and inflation itself.[11] The distribution undergoes a hump-shaped shift to the right, becomes more dispersed, is less skewed to the right, and exhibits thinner tails—all of which occur over a period of two to three years. The maximum mean shift of almost 0.4 percentage points is larger than the maximum impact of the oil supply news shock on general inflation of about 0.2 percentage points and lags the actual inflation dynamics. Consequently, we observe a significant movement of inflation expectation errors (defined by $FE_t = \pi_{t+12} - \bar{x}_t$, where $\bar{x}_t$ denotes the mean inflation expectations at time t) following the oil supply news shock, implying a temporary overestimation (for two to three years) of the inflation response.[12] The increase in the dispersion of inflation expectations following the oil supply news shock is consistent with the patterns that Reis (2021) documents

[11] We show the IRFs of the FPCs in Figure C3 in the appendix.

[12] The strong reaction of households' inflation expectations is consistent with the evidence in Känzig (2021), which shows that their expectations react more strongly to an oil supply news shock than those of professional forecasters in the U.S.

Figure 4: Impulse responses of macroeconomic aggregates and moments of the inflation expectations distribution.

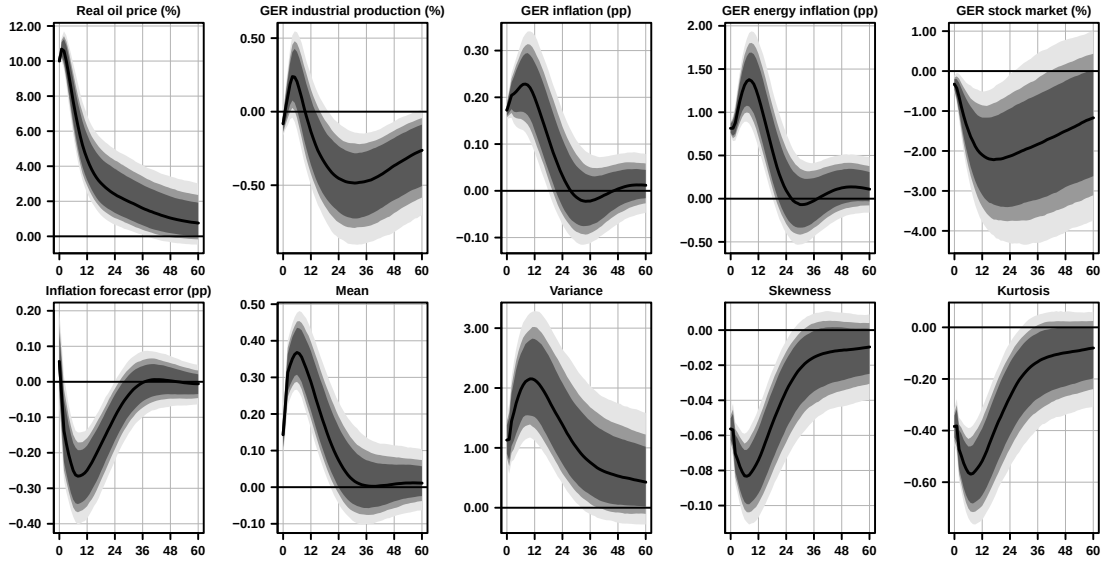

(a) Oil supply news shock.

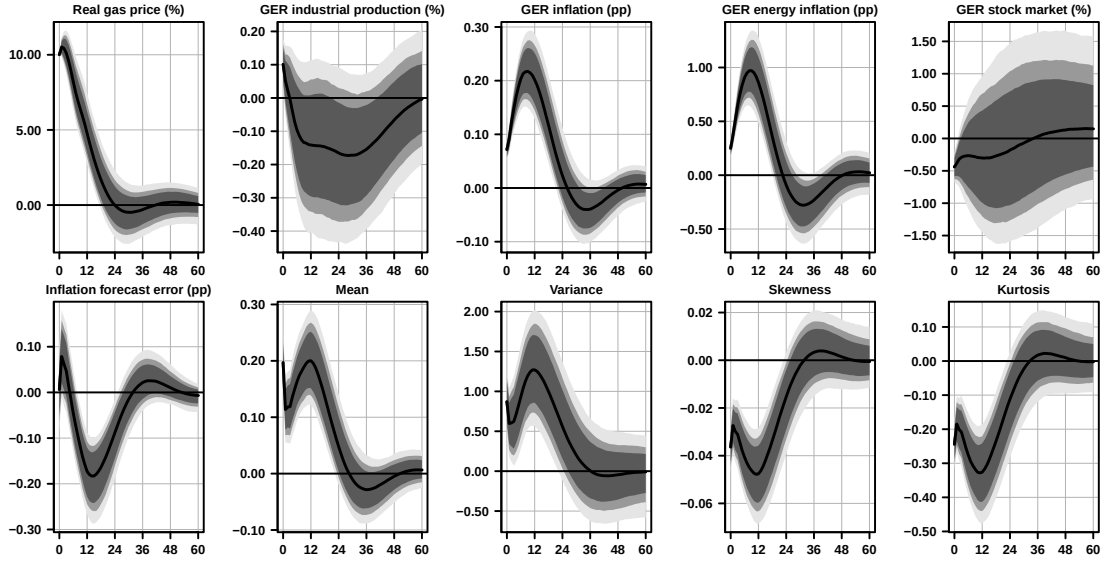

(b) Natural gas supply shock.

Notes: The figure shows the impulse responses of macroeconomic aggregates and the first four moments of the inflation expectation distribution to an oil supply news shock (panel a) and natural gas supply shock (panel b). The shock is normalized to induce a 10% increase in the real price of oil (panel a) or natural gas (panel b). The estimation covers the period 2003M5 to 2025M12. The black solid lines refer to the median estimates. The grey areas refer to the 68%/80%/90% credible sets. The horizontal axis measures the impulse horizon in months. The vertical axis denotes deviations from trend/steady state in absolute values (skewness, kurtosis), percent (real oil price, real gas price, industrial production, stock market), percentage points (inflation, energy inflation, mean) or squared percentage points (variance). The forecast error is defined as $FE_t = \pi_{t+12} - \bar{x}_t$, where $\bar{x}_t$ denotes the mean inflation expectations at time t.

for inflation expectations in the U.S. in 1974, during the aftermath of the oil shock. The decline in skewness reflects that the core of the distribution moves to the right, shrinking the gap to the extreme expectations in the long upper tail.

To display the effect of the shock on the entire distribution of inflation expectations, we must focus on specific horizons separately. Figure 5 shows the difference between the inflation expectation distribution following an oil supply news shock and the one prevailing in the steady state for $h = \{0, 3, 6, 12, 24, 60\}$. The response is non-monotonic, with the share of households holding inflation expectations close to two percent decreasing on impact, while the share of households reporting higher inflation expectations increases. Both movements are statistically significant and strengthen over time, peaking after one year. Consistent with the impulse response functions (IRFs) of the first moments shown in the previous figure, the distribution returns to the steady state after five years. The shift in probability mass to the right, visible in the first two years following the shock, explains the decrease in skewness and kurtosis; a substantial number of households adjust their expectations toward those of households that consistently hold extremely high inflation expectations.

Are these expectation shifts sizable from a monetary policy perspective? We believe so. For demonstration, we cut the support of inflation expectations into four intervals that are meaningful from the perspective of the ECB and analyze how the share of expectations falling into any of these intervals changes dynamically after an oil supply news shock.[13] The first interval includes all expectations of deflation and stagnant prices. The second interval collects all positive expectations up to three percent and constitutes the range of relatively well anchored expectations. The third interval ranges from above three to ten percent, capturing high yet reasonable expectations. The final interval includes all very high inflation expectations above 10 percent. The IRFs in Figure C4 in the appendix show substantial probability mass shifts. An oil supply news shock that raises real oil prices by an initial ten percent leads to a decline in the share of deflation or unchanged-price expectations by one percentage point and a decline in the share of moderate inflation expectations by almost two percentage points. In contrast, it increases the share of expectations between three and ten percent by 1 percentage point and the share of very high inflation expectations by more than 1.5 percentage points. For comparison: the total standard deviation of the latter share across all waves in our sample is 8.7 percentage points.

To investigate whether our results are sensitive to the specific approach chosen for identifying the supply shock, we re-compute the main IRFs for an oil supply shock that we

[13] As for the other IRFs, we treat the unconditional sample average of the distribution as our baseline.

Figure 5: Responses of the inflation expectation distribution to an oil supply news shock.

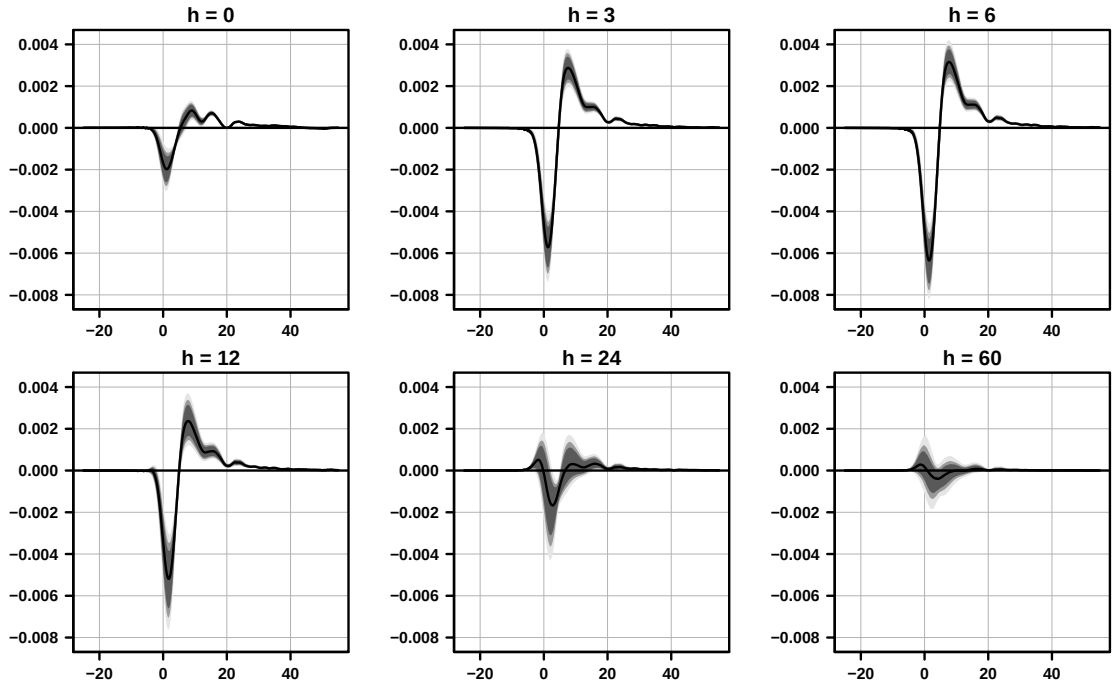

Notes: The figure shows the difference between the inflation expectation distribution after an oil supply news shock and the one prevailing in the steady state. The estimation covers the period 2003M5 to 2025M12. The black solid lines refer to the median estimates. The grey areas refer to the 68%/80%/90% credible sets. The horizontal axis measures inflation expectations in percent. The vertical axis denotes the effect on the density $p_t(x)$.

identify as in Baumeister and Hamilton (2019). The results are shown in Figures C5 and C6 in the appendix. Its estimated dynamic effects on the distribution of inflation expectations are very similar to those obtained for the oil supply news shock, both in terms of maximum magnitude and persistence. What is different is the more muted response of the macroeconomic variables, which aligns with Känzig's (2021) insight that the oil market is forward looking and that not accounting for anticipated shocks misses important oil supply effects. All important results, including the lagged overestimation of the inflation effect, stand.

4.3 Natural gas supply shock

We now turn to the IRFs for the natural gas supply shock, which we report in panel (b) of Figure 4. We notice that the shock is less persistent than the oil supply news shock and has no significant effect on industrial production (although point estimates are negative for most horizons). The response of the stock market index is also much more muted and very short-lived. The effect on headline inflation is comparable in size to that of the oil supply news shock, peaking at about 0.2 percentage points after roughly one year.

However, it is more hump-shaped, with a more moderate initial response. The effect on energy price inflation exhibits the same shape but is markedly weaker than that in response to the oil supply news shock. The peak effect after one year aligns with the pattern of the inflation effect of a gas flow supply shock estimated by Güntner et al. (2024).

Interestingly, the response of the distribution of inflation expectations, as shown by the IRFs for the first four moments in the second row of Panel (b), exhibits less build-up over time. All four moment measures respond on impact and exhibit further adjustment during the first year before returning toward their steady-state values.[14] This is confirmed by the responses of the entire distribution, as shown in Figure 6. The IRFs of the four moments also demonstrate that the maximum impacts of the natural gas supply shock on these moments are approximately half as large as those from the oil supply news shock described in the previous section. A major difference in the expectation response is that, initially, we do not observe an overshooting of expectations for about six months; if anything, households, on average, underestimate the impact that the natural gas supply shock will have twelve months ahead. Until approximately two years after the shock, they do not anticipate the eventual decline in the inflation response and, consequently, overestimate future inflation just as they almost immediately do following an oil supply news shock.

Due to the more muted response of the expectation distribution, we naturally obtain smaller shifts in the interval probabilities following the natural gas supply shock than those observed for the oil supply news shock. Figure C8 in the appendix shows that all dynamic responses remain below one percentage point throughout. Overall, just over one percentage point of probability mass is shifted from the two lower intervals to the two higher intervals.

4.4 Economic implications of the baseline results

Our findings relate to a number of economic models and carry important economic implications. The hump-shaped dynamics of the expectation distribution following the oil supply news shock are a hallmark of informational frictions in the expectation formation process. Under the sticky-information model of Mankiw and Reis (2002), only a fraction of agents update their information set in each period; the remainder continue to base their decisions on outdated beliefs. This predicts precisely the pattern we observe: a gradual build-up in the mean response that peaks well after the shock and then slowly dissipates, rather than an immediate jump and reversion. The concurrent increase in cross-sectional variance is equally diagnostic. As long as non-updaters remain anchored

[14] Interestingly, some of the FPCs' IRFs shown in Figure C7 in the appendix exhibit a hump shape.

Figure 6: Responses of the inflation expectation distribution to a natural gas supply shock.

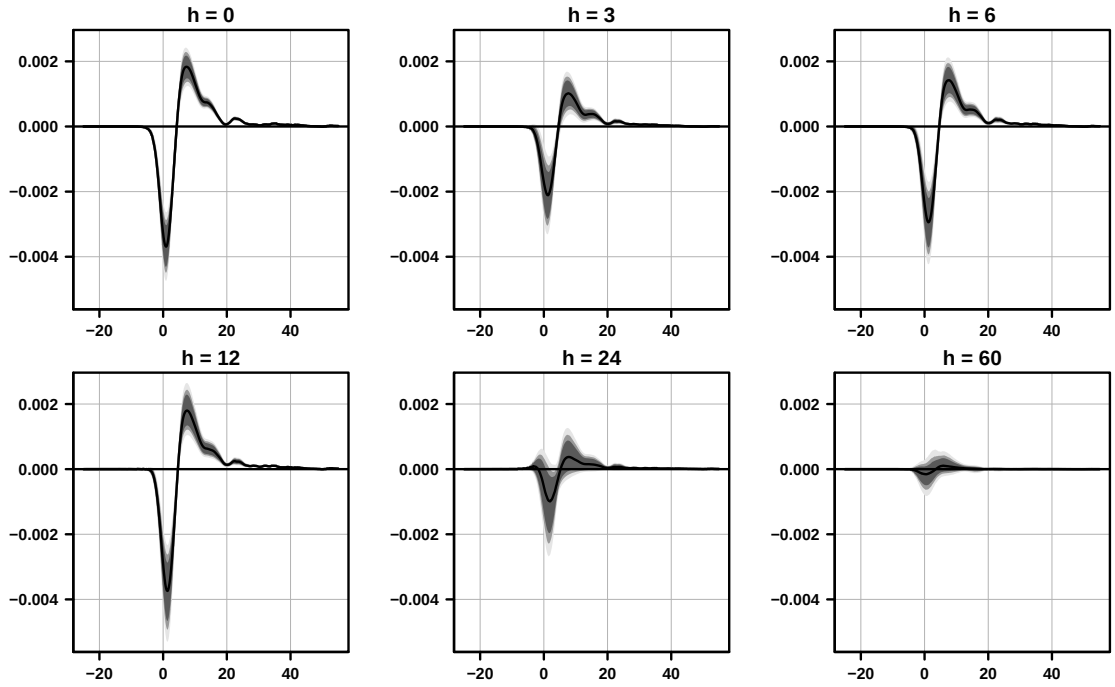

Notes: The figure shows the difference between the inflation expectation distribution after a natural gas supply shock and the one prevailing in the steady state. The estimation covers the period 2003M5 to 2025M12. The black solid lines refer to the median estimates. The grey areas refer to the 68%/80%/90% credible sets. The horizontal axis measures inflation expectations in percent. The vertical axis denotes the effect on the density $p_t(x)$.

at their prior while updaters revise upward, the distribution becomes a mixture of the two groups, mechanically widening the dispersion of beliefs. The model of rational inattention (Sims, 2003; Maćkowiak and Wiederholt, 2009) offers a complementary perspective. Because both shocks have direct effects on raw material prices but only indirectly influence the prices of goods and services that are part of the consumer price index, households face a non-trivial signal-extraction problem, which raises the cognitive cost of processing information and updating their inflation expectations. As a consequence, the response of mean expectations is sluggish, and the dispersion of expectations increases. The more immediate effect on mean expectations of the natural gas supply shock is instructive in this regard. Although, unlike gasoline prices at the pump, retail natural gas prices are not, in normal times, a particularly salient signal for German households,[15] much of the identifying variation for the natural gas shock stems from the period following Russia's invasion of Ukraine, when natural gas supply moved to the very center

[15] The prevalence of long-term fixed-price supply contracts means that wholesale natural gas price movements pass through to household bills only with considerable delay and substantial attenuation, weakening the direct price-experience channel emphasized by Coibion and Gorodnichenko (2015b) and D'Acunto et al. (2021a).

of public and media attention in Germany. During this episode, salience was generated not by households observing higher prices on their own bills, but by the extraordinary prominence of the supply situation in the national discourse. This interpretation is consistent with the response we document: all four distributional moments react more or less immediately to the natural gas shock, with little subsequent build-up, suggesting that households updated swiftly once the topic became unavoidably salient.

While sluggish adjustments and rising disagreements are consistent with information frictions, the shocks also lead to a disproportionate reweighting toward high-inflation states of the distribution. This pattern is consistent with a diagnostic-expectations mechanism in the sense of Bordalo et al. (2022): adverse inflation news makes high-inflation outcomes particularly representative, leading households to overweight those states relative to a Bayesian benchmark. The peak increase in mean expectations of almost 0.4 percentage points exceeds the peak effect of the shock on realized headline consumer price inflation of approximately 0.2 percentage points, reflecting an overshooting of beliefs relative to the actual inflationary impact of the shock. In the case of the natural gas supply shock, there is no overshooting regarding the peak increase. However, the reaction of expectations is too persistent, leading to an overestimation of the shock's medium-run inflationary impact. Similarly, the substantial increases in the share of expectations above ten percent (which exceed the increase in the share of expectations in the interval from three to ten percent in both cases) suggest a strong impact on the salience of high-inflation states.

These shifts in the full distribution of inflation expectations have direct implications for the dynamics of inflation itself through the heterogeneous-beliefs Phillips curve of Meeks and Monti (2023). In their framework, aggregate inflation depends not only on the mean of the expectation distribution and the output gap, but also on a cross-sectional disagreement term that captures the dispersion of beliefs around the consensus. Empirically, they show that including up to three FPCs that approximate the variation in the distribution of inflation expectations improves the explanatory power of the Phillips curve. Because our results document significant and persistent responses of multiple higher moments—not merely the mean—to the energy supply shocks, there is an additional mechanism through which these supply shocks affect future inflation beyond what their effect on consensus expectations alone would imply. Crucially, this additional channel is invisible to a central bank that monitors only the mean forecast. Even if policymakers successfully stabilized average expected inflation at the pre-shock level, the widening of the distribu-

tion would still push inflation upwards.[16] The persistence of our responses implies that this additional inflationary impetus constitutes sustained inflationary pressure. A key contribution of our analysis is to identify the *source* of these distributional shifts. While Meeks and Monti document that higher-moment variation in beliefs adds significant explanatory power to the Phillips curve, they treat the dynamics of the expectation distribution as an empirical regularity. We show that energy supply shocks are a structural driver of precisely those distributional shifts, providing the upstream causal chain that their framework takes as given.

The stagflationary consequences of these distributional dynamics are analyzed by Ascari et al. (2025). They show, using a Bayesian VAR augmented with the distribution of inflation expectations held by U.S. households, that exogenous shocks to the mean and dispersion of inflation expectations are stagflationary and that accounting for higher moments of the expectation distribution is important to avoid underestimating the macroeconomic implications of these stagflationary shocks.[17] We show that energy supply shocks simultaneously generate a mean shift and an increase in dispersion, which are the two types of perturbations that Ascari et al. show to be independently stagflationary. Where they treat distributional perturbations as primitive shocks and ask what macro effects follow, our paper tackles one further step by identifying an important class of real-world events—energy supply disruptions—that generate such perturbations in the first place. The implication is that the stagflationary character of energy supply shocks is amplified beyond the standard cost-push channel.

4.5 Relevance of energy supply shocks in 2022

The period of rising energy prices following the Russian war against Ukraine in 2022 was very challenging for European economies, particularly Germany, which relied heavily on energy imports from Russia. As discussed above, both oil and natural gas prices increased substantially during this period. To assess the supply effect on inflation and inflation expectations, we analyze how much oil supply news shocks and natural gas supply shocks contributed jointly to the rise in these two target variables after January 2022.

To achieve this, we rely on a joint, larger structural FVAR that includes both price series. We use the same external instruments and disentangle the two energy supply shocks by additionally imposing a recursive structure that assumes that the gas supply shock instrument does not have a contemporaneous effect on the instrument used to identify oil

[16] Note that this is an argument about the distribution of mean expectations rather than the uncertainty of individual expectations, which is arguably another expectation moment that is relevant in the context of expectation anchoring (Dovern and Kenny, 2020).

[17] Interestingly, they find that the expectations themselves are not driven by the (broadly defined) structural macroeconomic shock identified through sign-restrictions.

supply news shocks.[18] This is a common assumption, and our results confirm that real gas prices react to oil shocks, while the reverse is not true.[19] The resulting IRFs corresponding to the two shocks are very similar to those generated by the two shock-specific smaller FVARs, with the exception of the inflation and expectations responses to the natural gas supply shock, which are smaller in the large FVAR, and the response of industrial production to the same shock, which reverts back to zero much quicker in the large FVAR.[20]

First, we focus on the effects on inflation. The left part of Figure 7 displays the actual inflation rate alongside a hypothetical inflation rate derived from the joint large FVAR by excluding all estimated effects of oil supply news shocks or natural gas supply shocks, respectively, which we obtain from a historical decomposition of the structural shocks. In particular, the gas supply shocks increased inflation by almost two percentage points in 2022. This estimate is considerably larger than the inflationary impact attributed to gas flow supply shocks for the year 2022 by Güntner et al. (2024). The contribution of oil supply news shocks is more muted but not negligible, reaching almost 0.85 percentage points on average.

Next, we examine the effects on the distribution of inflation expectations. In particular, we compute the extent to which the two types of energy supply shocks shifted the distribution of inflation expectations between January 2022 and August 2022 (when energy prices peaked), starting from the observed distribution at the beginning of 2022 (black solid line with dots). We obtain the hypothetical expectation distribution for August by first accumulating the effects of the estimated shocks on the PCs in our FVAR and then deriving the implied shift of the distribution that the PCs describe. The results are shown in the right part of Figure 7. The blue solid curves represent the distributions obtained without the energy supply shocks. The red dashed curves represent the distributions obtained when taking these supply shocks into account, too. Comparing these two shows the shift to the right, accompanied by rising dispersion and a decrease in skewness, all of which are consistent with the IRFs shown above and a sequence of negative supply shocks. The sequence of natural gas supply shocks increases the mean expectation by roughly 0.55 percentage points, while the oil supply news shocks shift mean expectations upward by 0.45 percentage points. Hence, just as with actual inflation, the gas supply shocks were more relevant for the shift in inflation expectations during this period. In-

[18] Note that this does not imply that the real oil price does not move instantaneously in response to a natural gas supply shock.

[19] Results are not sensitive to this assumption. A re-ordering yields the same results, as the correlation between the external instruments is close to zero and statistically indistinguishable from zero.

[20] We reproduce the full set of IRFs for the joint, large FVAR in Figure C9 in the appendix.

Figure 7: Effects of energy supply shocks in 2022.

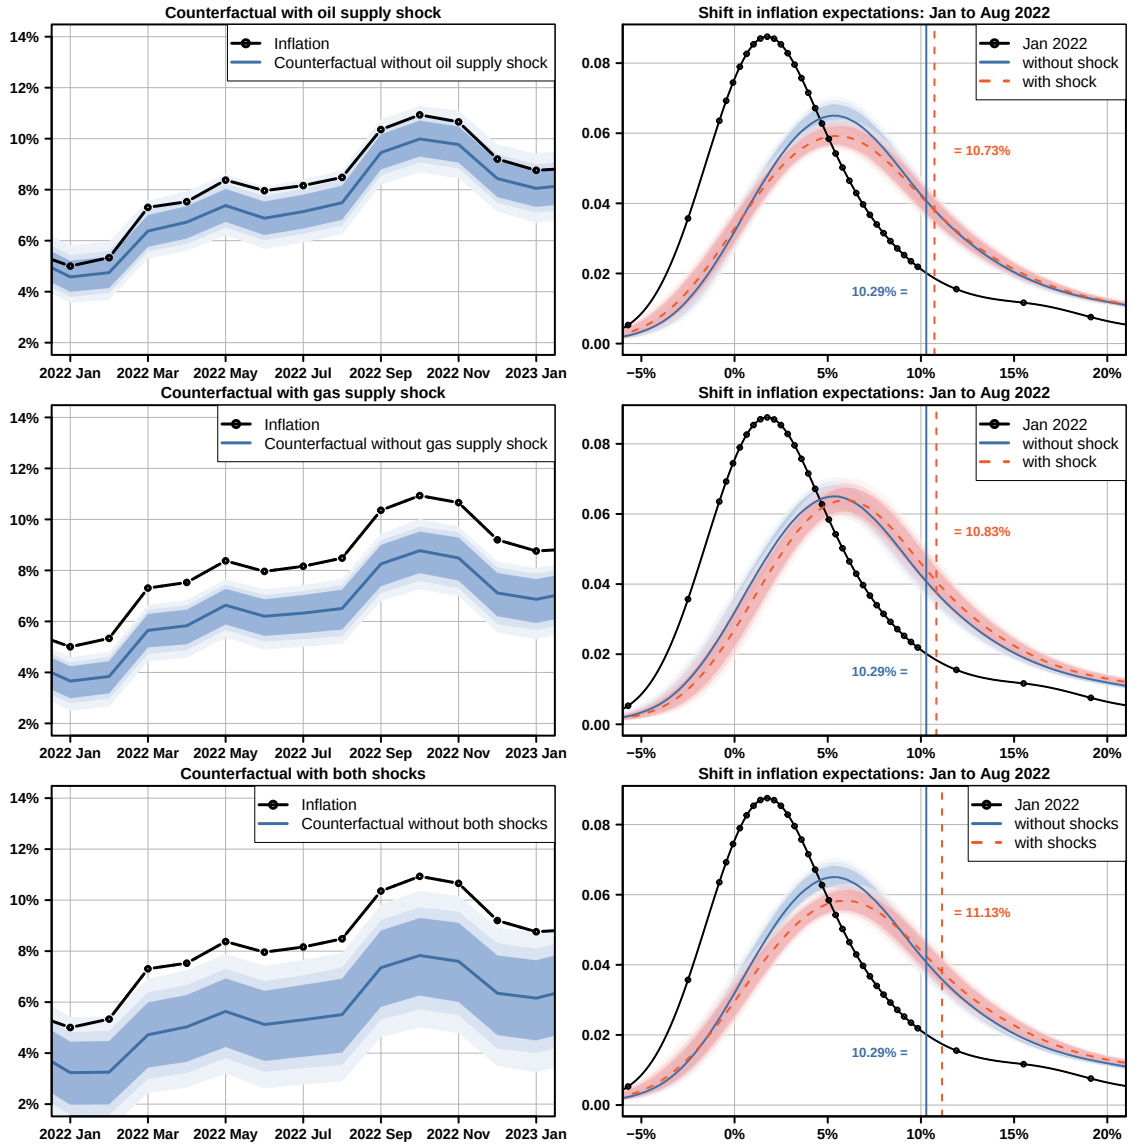

Notes: The figure shows the model-based effects of energy supply shocks on inflation and the distribution of inflation expectations in 2022. Results are based on the joint, large FVAR, estimated over the period 2003M5 to 2025M12. In the left plots, the black solid lines with dots show the actual inflation rate. The blue solid lines refer to the median estimates for the hypothetical inflation rate obtained by subtracting the effects of the respective supply shocks. The shaded areas refer to the 68%/80%/90% credible sets. In the right plots, the black solid lines with dots show the actual distribution of inflation expectations in January 2022. The blue solid lines show the median hypothetical distribution in August 2022 under the assumption that no energy supply shocks occur after January 2022. The red dashed lines show the corresponding distribution taking energy supply shocks into account. The shaded areas refer to the 68%/80%/90% credible sets. The two vertical lines indicate the respective means of the two distributions.

terestingly, the joint effect is almost additive, even though the transformations involved are not linear.

In summary, energy supply shocks, particularly those related to natural gas, substantially contributed to both the increase in inflation and the upward shift in the expectation distribution following the Russian attack on Ukraine in 2022.

4.6 Heterogeneity across household types

Since the literature has identified several socioeconomic variables that cause important heterogeneities in inflation expectations across different types of households (D'Acunto et al., 2023), we now turn to an analysis of such differences in terms of household types and their dynamic responses to energy supply shocks. To do so, we apply the same approach that we use for the full sample (incl. the data interpolation described in Section 2) to the subsamples corresponding to the different household types. It is not clear a priori that heterogeneities in the unconditional means of expectations (or the implied biases, for that matter) translate into the same heterogeneities in the dynamic responses to structural shocks.

Age cohorts. Figure 8 reports the moment responses separately for the young (under 30) and the old (aged 30 and above). The younger cohorts react more strongly to both shocks than the older cohorts regarding the first two moments. Following the oil supply news shock, the mean expectation of the young rises by 0.4 percentage points at its peak, compared to roughly 0.3 percentage points for the old. The age gradient is even sharper for the natural gas supply shock: the mean expectation of the young jumps on impact and remains nearly 0.3 percentage points above trend, while that of the older cohorts increases by only 0.15 percentage points at maximum. The cross-sectional variance of the young increases by more than twice as much as that of the older group after both shocks. While the response of skewness is very similar in the case of the oil supply news shock, skewness declines more strongly for younger households than for older ones. The old display more negative responses in kurtosis in the case of the oil supply news shock. Overall, younger households revise their beliefs more strongly and heterogeneously than older households.

The age differences in dynamic responses to energy supply shocks are naturally connected to the experience-based theory of expectation formation. In the model of Malmendier and Nagel (2016), households overweight the inflation they have personally experienced, with a gain that declines as the length of their accumulated experience increases; younger households, whose experience is shorter, attach a larger weight to any given inflation surprise and revise their beliefs more strongly than older households. This

aligns precisely with the ordering we find for the mean response to both energy supply shocks. Combining this insight with the sluggish revision of expectations by some households, due to the information frictions discussed above, explains the larger increase in cross-sectional variance for the young, as a more elastic response of the updaters widens the dispersion of beliefs more strongly than that observed for the old. Our findings stand in apparent tension with the memory-based mechanism of Gennaioli et al. (2024), which implies that it is the *old* who de-anchor most strongly. In their framework, however, a broad and salient inflation surge acts as a similarity cue that retrieves earlier high-inflation episodes, an effect that is strongest for those who have stored the most such episodes in memory. We interpret the two sets of results as regime-dependent rather than contradictory. Gennaioli et al. (2024)'s mechanism explains the de-anchoring of expectations during a generalized, highly salient inflation surge, where the retrieval of distant memories dominates. In contrast, our estimates isolate the marginal response to cleanly identified energy supply shocks. Although our sample period overlaps with the post-2021 inflation surge, it covers other periods when energy supply shocks were not followed by such a pronounced increase in inflation. Furthermore, our identification isolates the component of expectation revisions specifically attributable to the energy shock. Apparently, along this margin, it is the experience channel that dominates the selective-recall channel. The findings suggest that the age profile of expectation updating is not a fixed structural feature. Rather, it depends on the nature of the inflationary impulse.

Homeownership. Figure 8 also contrasts the effects of supply shocks on the expectations of both homeowners and renters. Following an oil supply news shock, the mean response of homeowners is more rapid and slightly stronger (by roughly 0.1 percentage points) compared to that of renters. The stronger mean response aligns with substantially greater declines in skewness and kurtosis, again reflecting the fact that for homeowners, a greater probability mass is shifted more strongly toward the right tail of the pre-shock inflation expectation distribution. Interestingly, we observe only minor differences between these two household types in the response of all four moments in the case of the natural gas supply shock.

The stronger and more rapid response of homeowners to the oil supply news shock is consistent with rational inattention. In models of optimally allocated attention, agents devote more of their limited attention to states of the world that matter more for their payoffs (Sims, 2003; Maćkowiak and Wiederholt, 2009; Maćkowiak et al., 2023). Homeowners—especially those carrying long-term nominal mortgage debt, as is common in Germany—hold considerably larger inflation-sensitive nominal positions than

Figure 8: Heterogeneity analysis: age and homeownership.

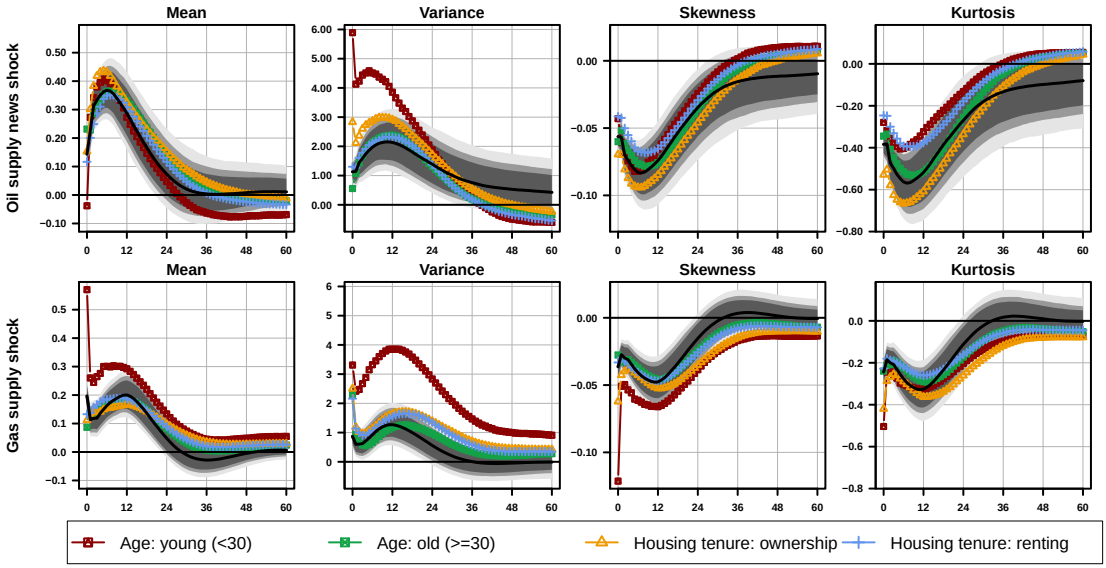

Notes: The figure shows the impulse responses of the first four moments of the inflation expectation distribution to an oil supply news shock (upper row) and a natural gas supply shock (lower row). Both shocks are normalized to induce a 10% increase in the real price of oil or natural gas, respectively. The estimation covers the period 2003M5 to 2025M12. The black solid lines refer to the full-sample median estimates. The grey areas refer to the corresponding 68%/80%/90% credible sets. The colored lines with markers refer to the median estimates obtained for subsamples of particular household types. The horizontal axis measures the impulse horizon in months. The vertical axis denotes deviations from steady state.

renters, whose nominal exposures are small and short-lived (Doepke and Schneider, 2006); inflation is therefore more relevant for them, giving them a stronger incentive to attend to inflation news and to revise their expectations more decisively when such news arrives. This is consistent with direct evidence that homeowners are more attentive to inflation at baseline (Piccolo and Gorodnichenko, 2025), an attentiveness that, in their information-provision experiments, leaves owners *less* influenced by being provided with inflation data. However, in our setting, it plausibly underlies a sharper reaction to a genuine inflationary shock. It is also consistent with the finding that the sensitivity of economic expectations to macroeconomic news is concentrated among households with greater exposure to macroeconomic risk (Roth and Wohlfart, 2020). Furthermore, our finding is in line with the evidence presented by Ahn et al. (2024), who show that homeowners in the U.S. adjust their inflation expectations in response to monetary policy shocks that move mortgage rates, while renters do not do so to the same extent. Two qualifications are noteworthy. First, because homeowners are on average older, whereas our age results show that *younger* households react more strongly, the unconditional homeowner-renter gap most likely understates the true difference: the stronger reaction of homeowners emerges despite an age composition that works against it. Second, the

gap disappears in the case of the natural gas shock. One potential explanation is that the confounding age effect completely overrides the homeowner-renter gap.

Gender. Differences across genders, as shown in Figure 9, are more muted, with male and female responses tracking each other closely throughout. There is a small difference in mean responses, but this difference is smaller than 0.1 percentage points for both shocks. The other notable gender difference is a stronger decline in kurtosis for men in response to the oil supply news shock. On balance, women react a bit more strongly compared to men, but the gaps are small relative to the widths of the credible sets.

The muted gender differences are noteworthy in light of the existing literature. One of the most robust findings in the household-expectations literature is that women report substantially higher inflation expectations than men (Bruine de Bruin et al., 2010), a gap that D'Acunto et al. (2021b) trace to gender roles and the associated differences in exposure to grocery prices rather than being an intrinsic gender-related bias. Our results show that this well-documented gap in the *level* of expectations does not carry over to the *dynamic response* to energy supply shocks: men and women revise their expectations by very similar amounts. This distinction is instructive. The gender gap appears to operate as a roughly parallel shift in the level of expectations that leaves the reaction to aggregate energy shocks largely unaffected. Our results align with evidence in Jo and Klopack (2025), which shows that men and women in the U.S. react very similarly regarding their inflation expectations in response to shifts in gasoline prices.

Education. Education seems to be a relevant factor influencing the distributional response of inflation expectations to energy supply shocks. In particular, the mean response is stronger and more persistent for households with lower levels of education compared to those with higher education levels (Figure 9).[21] Interestingly, the variance response is stronger for the highly educated following the natural gas supply shock, suggesting that the interpretation of the shock's impact on future inflation is more heterogeneous within this group compared to the dispersion of views among the less educated.

These education results can be discussed in terms of the debate over whether household expectations are better described by information rigidity or by salience and exposure. At first glance, the finding that the mean response is stronger and more persistent among less-educated households is in contrast to rational-inattention models (Sims, 2003; Coibion and Gorodnichenko, 2015a), in which less-informed agents update infrequently and sluggishly and should, therefore, react *less*. The tension dissolves once the nature of the signal is taken into account. A large body of evidence shows that less-

[21] We classify households with at least upper secondary school, technical college, or commercial college as high-education households.

Figure 9: Heterogeneity analysis: education and gender.

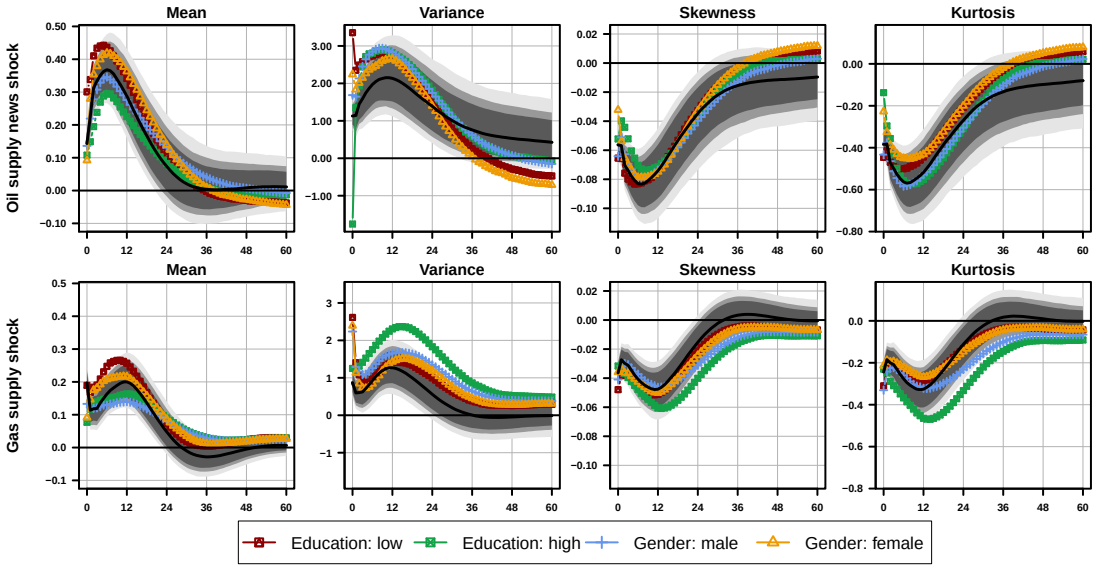

Notes: The figure shows the impulse responses of the first four moments of the inflation expectation distribution to an oil supply news shock (upper row) and a natural gas supply shock (lower row). Both shocks are normalized to induce a 10% increase in the real price of oil or natural gas, respectively. The estimation covers the period 2003M5 to 2025M12. The black solid lines refer to the full-sample median estimates. The grey areas refer to the corresponding 68%/80%/90% credible sets. The colored lines with markers refer to the median estimates obtained for subsamples of particular household types. The horizontal axis measures the impulse horizon in months. The vertical axis denotes deviations from steady state.

educated households disproportionately form their inflation expectations based on concrete, frequently observed prices, such as fuel prices, rather than on aggregate statistics or professional forecasts (D'Acunto et al., 2021a; Weber et al., 2022). An energy supply shock is precisely the kind of salient, visible price change to which this group is most exposed. This potentially explains why they react more strongly and why the response persists as they continue to observe elevated energy prices. It is harder to explain why dispersion rises more strongly among households with higher education following the natural gas shock. One potential explanation is offered by models in which such dispersion is driven by heterogeneous priors and forecasting models rather than by differences in information sets (Patton and Timmermann, 2010). Because the pass-through of a natural gas price increase to headline inflation is complex (including second-round effects, substitution, and the monetary-policy response), highly educated households that actively attempt to map this shock into an aggregate inflation forecast using different models arrive at more divergent conclusions. In contrast, the dispersion among less educated households, who think less in terms of complex transmission channels when translating the shock into the aggregate, increases to a lesser extent.

Hand-to-mouth status. Figure 10 contrasts the responses of hand-to-mouth (HtM) households, i.e., households that do not save from their current income, with those of all other households.[22] Following the oil supply news shock, the initial reactions of the mean expectations are very similar in both groups, but they are slightly more persistent in the case of HtM households. Following the natural gas supply shock, the mean response of HtM households is weaker on impact but builds up to approximately 0.2 percentage points during the first year, whereas that of non-HtM households remains at roughly 0.15 percentage points and declines more rapidly. These differences in mean responses are moderate relative to the estimation uncertainty. The higher moments, however, reveal that this masks differences in the shifts of the entire distribution of expectations. The variance of HtM households' expectations increases more strongly at its peak (by 2.8 compared to 2.4 squared percentage points after the oil supply news shock and by 2.5 compared to 1.9 squared percentage points after the natural gas supply shock), whereas their skewness and, in particular, their kurtosis decline less.

At first sight, the more persistent and, if anything, stronger response of HtM households is at odds with models in which households choose how much information about the aggregate economy to acquire. In the heterogeneous-agent model of Broer et al. (2026), acquiring such information has no value for households at the borrowing constraint. Thus, HtM households should be less attentive to aggregate conditions. Similar to our interpretation of the education results, the two perspectives can be reconciled by distinguishing between information that households must acquire at a cost and information that they receive as a by-product of their own purchases. HtM households tend to have lower incomes, and low-income households spend a larger share of their income on energy. This makes their personal cost of living particularly sensitive to energy prices (Känzig, 2026). Having little to no scope for savings, these households are forced to cut spending on other items when their energy bills rise, making energy supply shocks hard to ignore. There is ample evidence that such personal experiences shape households' expectations about aggregate outcomes (Kuchler and Zafar, 2019; D'Acunto et al., 2021a). Most closely related to our findings, Wehrhöfer (2026) shows that German households raise their inflation expectations when they personally experience an increase in electricity prices; this effect is driven by low-income households. Similarly, Kim and Sung (2026) find that the link between expected gasoline prices and expected inflation is more than three times as strong for low-income U.S. households as for high-income households. Unlike in the

[22] We classify households as HtM if they report in the "NIM Consumer Climate *powered by GfK*" that they are just managing to make ends meet on their income, are drawing on their savings, or are running into debt, and as non-HtM if they report that they are saving a little or a lot. While we deviate from the balance-sheet-based HtM definition of Kaplan et al. (2014) due to limited data availability, our filter identifies households that lack a financial buffer from current income.

Figure 10: Heterogeneity analysis: hand-to-mouth status and urbanicity.

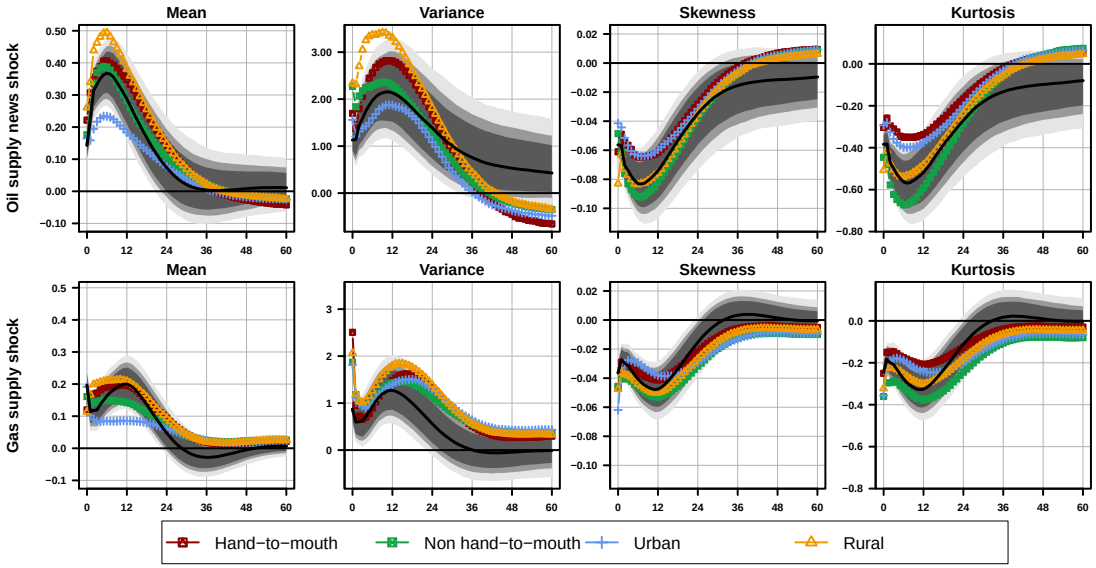

Notes: The figure shows the impulse responses of the first four moments of the inflation expectation distribution to an oil supply news shock (upper row) and a natural gas supply shock (lower row). Both shocks are normalized to induce a 10% increase in the real price of oil or natural gas, respectively. The estimation covers the period 2003M5 to 2025M12. The black solid lines refer to the full-sample median estimates. The grey areas refer to the corresponding 68%/80%/90% credible sets. The colored lines with markers refer to the median estimates obtained for subsamples of particular household types. The horizontal axis measures the impulse horizon in months. The vertical axis denotes deviations from steady state.

case of gender, the higher *level* of inflation expectations of HtM households goes hand in hand with a more persistent *response* to energy supply shocks, suggesting that exposure to energy prices shapes both.[23] One qualification is noteworthy: HtM households are more often less educated (79 compared to 68 percent) but less often homeowners (38 compared to 57 percent) than non-HtM households. According to our results above, these two composition effects work in opposite directions, at least for the oil supply news shock. Finally, from the perspective of HANK models with heterogeneous expectations, our results imply that the inflation expectations of those households that typically have the highest marginal propensities to consume (Kaplan et al., 2014) respond most persistently to energy supply shocks.

Urbanicity. Figure 10 distinguishes between households living in rural areas and those living in cities.[24] The differences between the two groups are large. Following the oil supply news shock, the mean expectation of rural households rises by almost 0.5 percentage

[23] Over our sample period, the average (median) inflation expectation equals 6.2 (3.0) percent for HtM households and 4.7 (2.5) percent for non-HtM households.

[24] We classify households as rural if they live in a municipality with fewer than 20,000 inhabitants, and as urban otherwise.

points at its peak after six months, which is more than twice the peak response of urban households (0.2 percentage points). The variance of rural households' expectations also increases considerably more (by up to 3.4 compared to 1.9 squared percentage points), and their skewness and kurtosis decline more strongly, reflecting a more pronounced shift of probability mass toward the upper part of the distribution. The natural gas supply shock produces different dynamics. Initially, urban households raise their mean expectation by nearly 0.2 percentage points, approximately twice as much as rural households (0.1 percentage points). However, the response of urban households drops below 0.1 percentage points within three months and remains at this level, whereas that of rural households builds up to about 0.2 percentage points and stays there for the remainder of the first year. One year after the shock, the mean response of rural households is about two and a half times as large as that of urban households. The responses of skewness and kurtosis mirror this pattern, with stronger declines for urban households on impact and for rural households at longer horizons.

The large rural–urban gap is consistent with differences in exposure to energy prices. Rural households depend much more on cars, and many rural communities still rely more on heating oil than urban areas. According to Többen et al. (2023), these differences in car dependency and heating infrastructure, together with larger floor space, imply that the average rural household in Germany is burdened more than twice as much by carbon pricing of heating and transport fuels as the average urban household with the same income. In the U.S., inflation differentials between rural and urban households are driven to a large extent by gasoline prices (Hobijn and Lagakos, 2005). Our results suggest that this higher exposure translates into stronger revisions of inflation expectations, in line with evidence that fuel prices are an important driver of household inflation expectations (Coibion and Gorodnichenko, 2015b; Kilian and Zhou, 2022a; Jo and Klopack, 2025) and that individuals extrapolate from personal experiences to aggregate outcomes (Kuchler and Zafar, 2019). Combined with the information frictions discussed above, the stronger response of updating households also explains the larger increase in the variance of rural households' expectations. The different dynamics after the natural gas supply shock are harder to explain. One potential explanation is that, in line with models of rational inattention (Sims, 2003; Maćkowiak and Wiederholt, 2009), news about the gas supply situation attracts more attention among urban households because natural gas dominates heating in cities (Többen et al., 2023), while long-term supply contracts delay and dampen the pass-through of higher gas prices to their bills. Urban households would then react more strongly to the salient news about the shock, in line with the salience-based interpretation offered above. This reaction, however, would not be reinforced by

subsequent personal price experiences. In contrast, the gradual build-up of the response of rural households mirrors the gradual increase in energy prices following the shock.

5. Conclusion

This paper studies how oil supply news shocks and natural gas supply shocks influence the entire distribution of households' inflation expectations in Germany. We estimate an FVAR using monthly survey data from 2003 to 2025 and identify the shocks with external instruments.

We find that these energy supply shocks change the entire shape of the expectation distribution. Following shocks that increase the prices of oil and gas, the distribution shifts to the right, dispersion rises, and mass moves disproportionately into the upper tail above 10 percent. Households overestimate the inflationary impact of both shocks: After oil supply news shocks, the mean response overestimates the inflation impact for two to three years and by up to 0.25 percentage points. After gas supply shocks, mean expectations react immediately without overshooting, but remain elevated for too long. For a given price increase, gas supply shocks move mean expectations about half as much as oil supply news shocks. Yet, during the energy crisis of 2022, gas was the more important driver of both inflation and expectation dynamics.

Three results from our analysis of heterogeneities across household types stand out. First, young households react much more strongly than older households. This is consistent with experience-based expectation formation (Malmendier and Nagel, 2016). Second, rural households' mean expectations react more than twice as strongly to oil supply news shocks as those of urban households, which is consistent with their greater direct exposure to energy prices (Többen et al., 2023). Third, the well-documented gender gap in the level of inflation expectations does not translate into substantial differences in the dynamic response to energy supply shocks.

We draw three conclusions from our results. First, the evidence suggests that at least three features are relevant for the dynamic response of expectations to the shocks. The sluggish response and rising dispersion are consistent with information frictions, as in models of rational inattention; the distributional shift toward high-inflation outcomes points to the diagnostic-expectations mechanism; many of the heterogeneity results suggest a role for salience. Second, the heterogeneous responses across household types imply that models with information frictions, but otherwise identical expectation formation for all households, miss important features of expectation dynamics. The strong reaction of hand-to-mouth households is particularly relevant in the context of HANK mod-

els, where these households with high marginal propensities to consume play a crucial role in aggregate dynamics.[25] Finally, because energy supply shocks raise expectation dispersion and the share expecting inflation above 10%, they potentially add inflationary pressure beyond the normal cost-push channel, as suggested by Meeks and Monti (2023). Thus, monitoring higher moments of the expectation distribution is relevant for policymakers to accurately estimate the inflation effects of supply shocks.

Looking ahead, we believe there are several issues that would be interesting to study. Given our results and the evidence in Meeks and Monti (2023) and Ascari et al. (2025), one could study optimal monetary policy in a standard New Keynesian model with information frictions when supply shocks move the entire expectation distribution. Empirically, one could study whether demand shocks also have effects on the higher moments of the expectation distribution.

References

Aastveit K A., Bjørnland H C., and Cross J L. (2023) "Inflation Expectations and the Pass-Through of Oil Prices". *Review of Economics and Statistics*, Vol. 105(3), pp. 733–743.

Adolfsen J F., Lappe M.-S., Manu A.-S., Rößler D., Schupp F., and Stalla-Bourdillon A. (2026) "Gas Market Shocks: Tracing the Effect on Euro Area Inflation Expectations". *Journal of Macroeconomics*, Vol. 88, p. 103760.

Ahn H J., Xie S., and Yang C. (2024) "Effects of Monetary Policy on Household Expectations: The Role of Homeownership". *Journal of Monetary Economics*, Vol. 147, p. 103599.

Alessandri P. and Gazzani A. (2025) "Natural Gas and the Macroeconomy: Not All Energy Shocks Are Alike". *Journal of Monetary Economics*, Vol. 151, p. 103749.

Ascari G., Bonomolo P., and Celani A. (2025) "The Macroeconomic Effects of Inflation Expectations: The Distribution Matters". Working paper.

Auclert A., Rognlie M., and Straub L. (2020) *Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model*. NBER Working Paper 26647. National Bureau of Economic Research.

Auclert A., Rognlie M., and Straub L. (2025) "Fiscal and Monetary Policy with Heterogeneous Agents". *Annual Review of Economics*, Vol. 17, pp. 539–562.

Bai J. and Ng S. (2002) "Determining the Number of Factors in Approximate Factor Models". *Econometrica*, Vol. 70(1), pp. 191–221.

Bauer M D., Känzig D R., and Rudebusch G D. (2026) "Carbon Pricing and Inflation Expectations". *The Econometrics Journal* (forthcoming).

[25] As a caveat, we need to remind readers that our measure of HtM status is approximate due to limited data availability.

Baumeister C. and Hamilton J D. (2019) "Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks". *American Economic Review*, Vol. 109(5), pp. 1873–1910.

Bernanke B S., Boivin J., and Eliasz P. (2005) "Measuring the Effects of Monetary Policy: a Factor-Augmented Vector Autoregressive (FAVAR) Approach". *The Quarterly Journal of Economics*, Vol. 120(1), pp. 387–422.

Bjørnland H C., Chang Y., and Cross J L. (2026) "Oil and the Stock Market Revisited: A Mixed Functional VAR Approach". *Quantitative Economics*, Vol. 17(2), pp. 541–589.

Boeck M. and Zörner T O. (2025) "Natural Gas Prices, Inflation Expectations, and the Pass-Through to Euro Area Inflation". *Energy Economics*, Vol. 141, p. 108061.

Bordalo P., Gennaioli N., and Shleifer A. (2022) "Overreaction and Diagnostic Expectations in Macroeconomics". *Journal of Economic Perspectives*, Vol. 36(3), pp. 223–244.

Bosq D. (2000) *Linear Processes in Function Spaces: Theory and Applications*. Vol. 149. Springer Science & Business Media.

Broer T., Kohlhas A., Mitman K., and Schlafmann K. (2026) "Expectation and Wealth Heterogeneity in the Macroeconomy". *Journal of Political Economy* (forthcoming).

Bruine de Bruin W., Klaauw van der W., Downs J S., Fischhoff B., Topa G., and Armantier O. (2010) "Expectations of Inflation: The Role of Demographic Variables, Expectation Formation, and Financial Literacy". *Journal of Consumer Affairs*, Vol. 44, pp. 381–402.

Casoli C., Manera M., and Valenti D. (2024) "Energy Shocks in the Euro Area: Disentangling the Pass-Through from Oil and Gas Prices to Inflation". *Journal of International Money and Finance*, Vol. 147, p. 103154.

Chang M., Chen X., and Schorfheide F. (2024) "Heterogeneity and Aggregate Fluctuations". *Journal of Political Economy*, Vol. 132(12), pp. 4021–4067.

Chang M. and Schorfheide F. (2026) "On the Effects of Monetary Policy Shocks on Income and Consumption Heterogeneity". *American Economic Journal: Macroeconomics* (forthcoming).

Chang Y., Gómez-Rodríguez F., and Hong G H. (2022) *The Effects of Economic Shocks on Heterogeneous Inflation Expectations*. IMF Working Paper WP/22/132. International Monetary Fund.

Chopra F., Roth C., and Wohlfart J. (2025) "Home Price Expectations and Spending: Evidence from a Field Experiment". *American Economic Review*, Vol. 115(7), pp. 2267–2305.

Coibion O., Georgarakos D., Gorodnichenko Y., Kenny G., and Weber M. (2024) "The Effect of Macroeconomic Uncertainty on Household Spending". *American Economic Review*, Vol. 114(3), pp. 645–677.

Coibion O. and Gorodnichenko Y. (2015a) "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts". *American Economic Review*, Vol. 105(8), pp. 2644–2678.

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

D'Acunto F., Malmendier U., Ospina J., and Weber M. (2021a) "Exposure to Grocery Prices and Inflation Expectations". *Journal of Political Economy*, Vol. 129(5), pp. 1615–1639.

D'Acunto F., Malmendier U., and Weber M. (2021b) "Gender Roles Produce Divergent Economic Expectations". *Proceedings of the National Academy of Sciences*, Vol. 118(21), e2008534118.

D'Acunto F., Malmendier U., and Weber M. (2023) "What Do the Data Tell Us About Inflation Expectations?" *Handbook of Economic Expectations*. Ed. by R Bachmann, G Topa, and W van der Klaauw. Academic Press, pp. 133–161.

Doan T., Litterman R., and Sims C. (1984) "Forecasting and Conditional Projection Using Realistic Prior Distributions". *Econometric Reviews*, Vol. 3(1), pp. 1–100.

Doepke M. and Schneider M. (2006) "Inflation and the Redistribution of Nominal Wealth". *Journal of Political Economy*, Vol. 114(6), pp. 1069–1097.

Doh T., Lee J H., and Park W Y. (2025) "Heterogeneity In Household Inflation Expectations And Monetary Policy". *Journal of Financial Econometrics*, Vol. 23(1), nbae034.

Dovern J., Fritsche U., and Slacalek J. (2012) "Disagreement Among Forecasters in G7 Countries". *The Review of Economics and Statistics*, Vol. 94(4), pp. 1081–1096.

Dovern J. and Kenny G. (2020) "Anchoring Inflation Expectations in Unconventional Times: Micro Evidence for the Euro Area". *International Journal of Central Banking*, Vol. 16(5), pp. 309–347.

Farhi E. and Werning I. (2019) "Monetary Policy, Bounded Rationality, and Incomplete Markets". *American Economic Review*, Vol. 109(11), pp. 3887–3928.

Fleishman A I. (1978) "A Method for Simulating Non-Normal Distributions". *Psychometrika*, Vol. 43(4), pp. 521–532.

Frank P. (2026) "Carbon Pricing and Heterogeneity in Household Inflation Expectations: Evidence from the Euro Area". *Macroeconomic Dynamics*, Vol. 30, e58.

Gagliardone L. and Gertler M. (2026) "Oil Prices, Monetary Policy and Inflation Surges". *American Economic Journal: Macroeconomics*, Vol. 18(4), pp. 468–501.

Gennaioli N., Leva M., Schoenle R., and Shleifer A. (2024) *How Inflation Expectations De-Anchor: The Role of Selective Memory Cues*. NBER Working Paper 32633. National Bureau of Economic Research.

Güntner J., Reif M., and Wolters M. (2024) "Sudden Stop: Supply and Demand Shocks in the German Natural Gas Market". *Journal of Applied Econometrics*, Vol. 39(7), pp. 1282–1300.

Hensel J., Mangiante G., and Moretti L. (2024) "Carbon Pricing and Inflation Expectations: Evidence from France". *Journal of Monetary Economics*, Vol. 147, p. 103593.

Hobijn B. and Lagakos D. (2005) "Inflation Inequality in the United States". *Review of Income and Wealth*, Vol. 51(4), pp. 581–606.

Horváth L. and Kokoszka P. (2012) *Inference for Functional Data with Applications*. Vol. 200. Springer Science & Business Media.

Huber F., Koop G., and Matthes C. (2026a) *Beyond Aggregate VARs: A Bayesian Benchmark for HANK Models*. arXiv preprint arXiv:2609.06827.

Huber F., Marcellino M., and Tornese T. (2026b) "The Distributional Effects of Economic Uncertainty". *International Economic Review* (forthcoming).

Inoue A. and Rossi B. (2021) "A New Approach to Measuring Economic Policy Shocks, with an Application to Conventional and Unconventional Monetary Policy". *Quantitative Economics*, Vol. 12(4), pp. 1085–1138.

Jo Y J. and Klopack B. (2025) "Fueling Expectations: The Causal Impact of Gas Prices on Inflation Expectations". *Journal of Monetary Economics*, Vol. 153, p. 103797.

Känzig D R. (2021) "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements". *American Economic Review*, Vol. 111(4), pp. 1092–1125.

Känzig D R. (2026) "The Unequal Economic Consequences of Carbon Pricing". *American Economic Review* (forthcoming).

Kaplan G., Moll B., and Violante G L. (2018) "Monetary Policy According to HANK". *American Economic Review*, Vol. 108(3), pp. 697–743.

Kaplan G., Violante G L., and Weidner J. (2014) "The Wealthy Hand-to-Mouth". *Brookings Papers on Economic Activity*, Vol. 45(1), pp. 77–153.

Karlsson S. (2013) "Forecasting with Bayesian Vector Autoregression". *Handbook of economic forecasting*, Vol. 2, pp. 791–897.

Kilian L. and Zhou X. (2022a) "Oil Prices, Gasoline Prices, and Inflation Expectations". *Journal of Applied Econometrics*, Vol. 37(5), pp. 867–881.

Kilian L. and Zhou X. (2022b) "The Impact of Rising Oil Prices on US Inflation and Inflation Expectations in 2020–23". *Energy Economics*, Vol. 113, p. 106228.

Kim T. and Sung Y. (2026) *Consumer Vibe Check: Gas Prices and Inflation Expectations*. FRBSF Economic Letter 2026-24. Federal Reserve Bank of San Francisco.

Kneip A. and Utikal K J. (2001) "Inference for Density Families Using Functional Principal Component Analysis". *Journal of the American Statistical Association*, Vol. 96(454), pp. 519–542.

Koop G., McIntyre S., Mitchell J., and Wu P. (2026) *Incorporating Micro Data into Macro Models Using Pseudo VARs*. Working Paper 26-04. Federal Reserve Bank of Cleveland.

Kuchler T. and Zafar B. (2019) "Personal Experiences and Expectations about Aggregate Outcomes". *The Journal of Finance*, Vol. 74(5), pp. 2491–2542.

Maćkowiak B., Matějka F., and Wiederholt M. (2023) "Rational Inattention: A Review". *Journal of Economic Literature*, Vol. 61(1), pp. 226–273.

Maćkowiak B. and Wiederholt M. (2009) "Optimal Sticky Prices under Rational Inattention". *American Economic Review*, Vol. 99(3), pp. 769–803.

Maćkowiak B. and Wiederholt M. (2026) "Rational Inattention during a Randomized Controlled Trial". *AEA Insights*, Vol. 8(1), pp. 19–35.

Madeira C. and Zafar B. (2015) "Heterogeneous Inflation Expectations and Learning". *Journal of Money, Credit and Banking*, Vol. 47(5), pp. 867–896.

Malmendier U. and Nagel S. (2016) "Learning from Inflation Experiences". *The Quarterly Journal of Economics*, Vol. 131(1), pp. 53–87.

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

Meeks R. and Monti F. (2023) "Heterogeneous Beliefs and the Phillips Curve". *Journal of Monetary Economics*, Vol. 139, pp. 41–54.

Mori L. and Peersman G. (2026) "Estimating the Macroeconomic Effects of Oil Supply News". *Review of Economics and Statistics* (forthcoming).

Overhage J. (2026) *Inferring the Aggregate: Information Frictions and Fiscal Policy*. IIES Working Paper. Institute for International Economic Studies.

Patton A J. and Timmermann A. (2010) "Why Do Forecasters Disagree? Lessons from the Term Structure of Cross-Sectional Dispersion". *Journal of Monetary Economics*, Vol. 57(7), pp. 803–820.

Peersman G. and Van Robays I. (2009) "Oil and the Euro Area Economy". *Economic Policy*, Vol. 24(60), pp. 603–651.

Petersen A. and Müller H.-G. (2016) "Functional Data Analysis for Density Functions by Transformation to a Hilbert Space". *The Annals of Statistics*, Vol. 44(1), pp. 183–218.

Piccolo J. and Gorodnichenko Y. (2025) *Homeownership and Attention to Inflation: Evidence from Information Treatments*. NBER Working Paper 33595. National Bureau of Economic Research.

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

Reis R. (2021) "Losing the Inflation Anchor". *Brookings Papers on Economic Activity*, Vol. Fall 2021, pp. 307–379.

Reis R. (2023) "What Can Keep Euro Area Inflation High?" *Economic Policy*, Vol. 38(115), pp. 495–517.

Roth C. and Wohlfart J. (2020) "How Do Expectations about the Macroeconomy Affect Personal Expectations and Behavior?" *The Review of Economics and Statistics*, Vol. 102(4), pp. 731–748.

Sims C A. (1980) "Macroeconomics and Reality". *Econometrica*, Vol. 48 (1), pp. 1–48.

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

Stock J H. and Watson M W. (2016) "Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics". *Handbook of Macroeconomics*. Vol. 2. Elsevier, pp. 415–525.

Többen J., Pichler P.-P., Jaccard I S., Kratena K., Moran D., Zheng H., and Weisz H. (2023) "Unequal Carbon Tax Impacts on 38 Million German Households: Assessing Spatial and Socio-Economic Hotspots". *Environmental Research: Climate*, Vol. 2(4), p. 045006.

Tsay R S. (2016) "Some Methods for Analyzing Big Dependent Data". *Journal of Business & Economic Statistics*, Vol. 34(4), pp. 673–688.

Tsiaplias S. (2020) "Time-Varying Consumer Disagreement and Future Inflation". *Journal of Economic Dynamics and Control*, Vol. 116, p. 103903.

Weber M., D'Acunto F., Gorodnichenko Y., and Coibion O. (2022) "The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications". *Journal of Economic Perspectives*, Vol. 36(3), pp. 157–184.

Wehrhöfer N. (2026) "Energy Prices and Inflation Expectations: Evidence from Households and Firms". *Journal of Political Economy Macroeconomics*, Vol. 4(1), pp. 135–163.

Wong B. (2015) "Do Inflation Expectations Propagate the Inflationary Impact of Real Oil Price Shocks? Evidence from the Michigan Survey". *Journal of Money, Credit and Banking*, Vol. 47(8), pp. 1673–1689.

Online Appendix for
Energy supply shocks and inflation expectations: a functional VAR approach

Maximilian Boeck Anja Bauer Jonas Dovern

Contents

A Additional data information **A-2**

B Further details on the methodological approach **A-5**

C Additional empirical results **A-11**

A. Additional data information

All macroeconomic data were gathered from the sources listed below, which include the ECB Statistical Data Warehouse (ECB), the Federal Statistical Office of Germany (DeStatis), the World Bank Pink Sheets, and Yahoo Finance. In Table A1, we provide an overview of the macroeconomic variables, their transformations, and sources.

Table A1: Definitions, transformations, and sources of the macroeconomic data.

Variable (y_t)	Transformation	Source
Raw data:		
Crude oil, Brent ($/bbl)	x_t	World Bank Pink Sheet
Natural gas, Europe ($/mmbtu)	x_t	World Bank Pink Sheet
Harmonized consumer price index, Overall index	p_t	ECB
Harmonized consumer price index, Overall index, Germany	x_t	ECB
Harmonized consumer price index, electricity, gas, and other fuels, Germany	x_t	ECB
German industrial production, manufacturing	x_t	DeStatis
German stock market index, DAX	x_t	Yahoo Finance
VAR variables:		
Real price of oil	$100 \times \log(x_t / p_t)$	
Real price of natural gas	$100 \times \log(x_t / p_t)$	
GER industrial production	$100 \times \log x_t$	
GER inflation, yoy	$100 \times \log(x_t / x_{t-12})$	
GER energy price inflation, yoy	$100 \times \log(x_t / x_{t-12})$	
GER stock market	$100 \times \log x_t$	

Notes: yoy is short for year-on-year.

The data on households' inflation expectations and socioeconomic backgrounds are from the "NIM Consumer Climate *powered by GfK*" survey. This is a survey of 2,000 households in Germany collected via the GFK eBus (now NIQ eBus) and published by the Nuremberg Institute for Market Decisions (NIM). It is co-funded by the European Commission as part of the European Business and Consumer Surveys (BCS). The sample is a monthly repeated cross-section that is selected using a quota procedure, which considers age, gender, region, size of community, size of the household, and the education of the main income earner, and provides sampling weights. The respondents are sampled from a pool of 20,000 potential respondents aged 18-74 years, and after responding to the survey, they are excluded for at least 3 months, making resampling the same respondent in each year rare. The survey period covers the first two weeks of the

Figure A1: LQD of inflation expectation distribution over time.

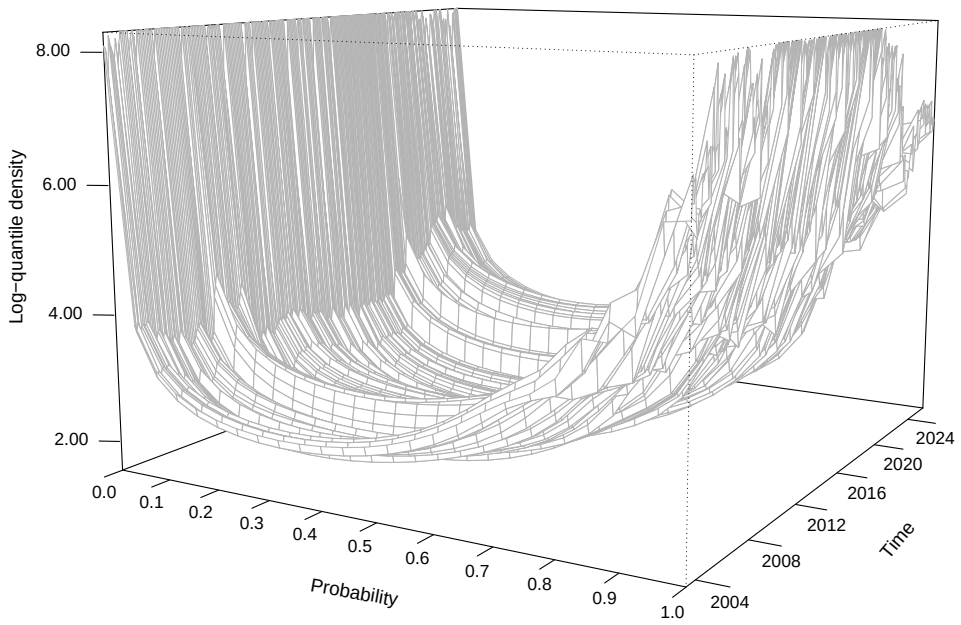

Notes: The figure shows the log-quantile density (LQD) transformation over time in our sample.

month and has been conducted via Computer-Assisted Web Interviewing (CAWI) since May 2019. More information on metadata and the questionnaire can be found on the European Commission's website (Metadata, Questionnaire). The question wording that is used to elicit inflation expectations is displayed in the main text in Section 2. Table A2 documents the questions on socioeconomic characteristics of households that we use to analyze the heterogeneity in expectation responses in Section 4.6.

Table A2: Sociodemographic variables used from "NIM Consumer Climate *powered by GfK*"

Variable	Question	Possible Answers
age	How old are you?	open answer, two digits, range 18-74 years
homeownership	Do you or your family live ...	1 in your own house 2 in your own flat 3 in a rented accommodation
gender	What is your gender (according to birth register)?	1 male 2 female
education	What school-leaving qualification do you have?	1 secondary school without a completed apprenticeship 2 secondary school with a completed apprenticeship 3 upper secondary school / technical college / commercial college 4 A-levels /university entrance qualification 5 university degree 99 N-A
HtM status	What is your household's current financial situation?	1 We save a lot 2 We save a little 3 We just about make ends meet 4 We're dipping into our savings a bit 5 We're getting into debt 9 Don't know
Urbanicity	Please indicate the population of your place of residence:	1 up to 1,999 inhabitants 2 2,000 to 2,999 inhabitants 3 3,000 to 4,999 inhabitants 4 5,000 to 9,999 inhabitants 5 10,000 to 19,999 inhabitants 6 20,000 to 49,999 inhabitants 7 50,000 to 99,999 inhabitants 8 100,000 to 199,999 inhabitants 9 200,000 to 499,999 inhabitants 10 500,000 inhabitants and over

Notes: The question regarding urbanicity has only five options from May 2019 onward (change to NIQ eBus). This does not affect our classification regarding urban/rural.

B. Further details on the methodological approach

In this appendix section, we provide further details on the Bayesian estimation of the vector autoregression. We describe the prior setup in detail. We also provide more details on the identification via external instruments. In the second subsection, we provide further details on the log-quantile density transformation.

B.1 Bayesian estimation of the vector autoregression

As detailed in Section 3, the finite-dimensional approximation of the functional vector autoregression yields a standard linear vector autoregressive model. In this appendix, we provide additional information on the model and its estimation. The time series process $\{\boldsymbol{Y}_t\}_{t=1}^T$ $(n \times 1)$ with $\boldsymbol{Y}_t = (\boldsymbol{y}_t', \boldsymbol{\alpha}_t')'$ follows

$$\boldsymbol{Y}_t = \boldsymbol{c} + \sum_{\ell=1}^{p} \boldsymbol{B}_\ell \boldsymbol{Y}_{t-\ell} + \boldsymbol{\Pi}\boldsymbol{d}_t + \boldsymbol{u}_t, \qquad \boldsymbol{u}_t \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma}), \tag{B.1}$$

where $\boldsymbol{B}_\ell$ are $n \times n$ coefficient matrices for lag $\ell \in \{1, 2, \dots, p\}$, and the $n \times 1$ vector $\boldsymbol{c}$ denotes the constant. The $n_e \times 1$ vector $\boldsymbol{d}_t$ denotes the set of exogenous variables with the corresponding $n \times n_e$ coefficient matrix $\boldsymbol{\Pi}$. The $n \times 1$ vector $\boldsymbol{u}_t$ denotes the reduced-form innovations and is normally distributed, $\boldsymbol{u}_t \sim \mathcal{N}(0, \boldsymbol{\Sigma})$. Note that

$$\boldsymbol{B}_\ell = \begin{pmatrix} \boldsymbol{B}_{\ell,yy} & \boldsymbol{B}_{\ell,y\alpha}\boldsymbol{C}_\alpha \\ \boldsymbol{B}_{\ell,\alpha y} & \boldsymbol{B}_{\ell,\alpha\alpha}\boldsymbol{C}_\alpha \end{pmatrix} \qquad \text{and} \qquad \boldsymbol{c} = \begin{pmatrix} \boldsymbol{c}_y \\ \boldsymbol{c}_\alpha \end{pmatrix}. \tag{B.2}$$

The model estimation is done in a Bayesian fashion, and thus we discuss our prior choices. We collect all parameters of the VAR coefficient matrices in the $K \times 1$ vector $\boldsymbol{b} = \text{vec}(\boldsymbol{B}_1, \dots, \boldsymbol{B}_p, \boldsymbol{c}, \boldsymbol{\Pi})$ with $K = n^2 p + n + n n_e$. Generally, we propose a normally distributed prior

$$\boldsymbol{b} \sim \mathcal{N}(\underline{\boldsymbol{b}}, \underline{\boldsymbol{V}}), \tag{B.3}$$

where $\underline{\boldsymbol{b}} = \text{vec}(\boldsymbol{I}, \boldsymbol{0}, \dots, \boldsymbol{0})'$ translates to the prior mean for first own-lag coefficients being centered at unity, with all other coefficients at zero. This corresponds to a random-walk specification a priori, as proposed by Litterman (1986). The corresponding prior variances are collected in $\underline{\boldsymbol{V}} = \text{diag}(\underline{v_1}, \dots, \underline{v_K})$. The prior variance of coefficient r $(r = 1, \dots, K)$

depends on the equation i, the lagged variable j, and the lag order ℓ associated with that coefficient. Under the Minnesota specification,

$$\underline{v}_r = \begin{cases} \left(\frac{\lambda_1}{\ell^{\lambda_3}}\right)^2 & \text{for an own lag coefficient } j = i, \\ \left(\frac{s_i^2}{s_j^2}\right)\left(\frac{\lambda_1\lambda_2}{\ell^{\lambda_3}}\right)^2 & \text{for a cross-lag coefficient } j \neq i, \\ s_i^2(\lambda_1\lambda_4)^2 & \text{for a deterministic or exogenous coefficient in equation } i, \end{cases} \tag{B.4}$$

Here $i, j \in \{1, \ldots, n\}$ denote the equation and lagged-variable indices, respectively, and $\ell \in \{1, \ldots, p\}$ denotes the lag order. The deterministic and exogenous terms include the constant and dummy variables. For each series i, s_i^2 is the sample variance of the residuals from a univariate AR(4) model. We have four hyperparameters to specify: λ_1 is an overall tightness parameter, λ_2 represents a cross-variable specific variance parameter, λ_3 is a scaling coefficient controlling the speed at which coefficients for lags greater than 1 converge to 0 with greater certainty, and λ_4 is a large (potentially infinite) variance parameter for the deterministic part of the model. We set $\lambda_1 = 0.1$, $\lambda_2 = 0.5$, $\lambda_3 = 1$, and $\lambda_4 = 100$, which are standard values in the literature (Dieppe et al., 2018). This is our baseline specification.

In the robustness analysis, we also utilize a hierarchical global-local shrinkage prior setup based on the horseshoe (HS) prior following Carvalho et al. (2010), where the prior variance is specified as

$$\underline{v}_r = \lambda^2\psi_r^2, \quad \lambda \sim C^+(0,1), \quad \psi_r \sim C^+(0,1), \tag{B.5}$$

where C^+ denotes the half-Cauchy distribution. The parameter λ denotes the global shrinkage parameter, which exerts shrinkage on all coefficients. The parameter ψ_r allows for individual, coefficient-specific shrinkage. Both parameters are estimated, and we do not have to specify any hyperparameters. We use the procedure outlined in Makalic and Schmidt (2016) to sample from the corresponding conditional posterior densities. Note that this specification only holds for the VAR coefficients, but for the deterministic and exogenous terms, we assume a diffuse Gaussian prior $c_i \sim \mathcal{N}(0, 10^5)$ and $\Pi_{ij} \sim \mathcal{N}(0, 10^5)$ $(i = 1, \ldots, n,\ j = 1, \ldots, n_e)$.

The prior distribution of the covariance matrix is

$$\Sigma \sim iW(\underline{\nu}, \underline{\boldsymbol{S}}), \tag{B.6}$$

where $iW(\nu, \boldsymbol{S})$ denotes the inverse Wishart distribution with prior degrees of freedom ν and prior scaling matrix $\boldsymbol{S}$. Following the recommendation in Kadiyala and Karlsson (1997), we specify $\underline{\nu} = n + 2$ and $\underline{\boldsymbol{S}} = \text{diag}(s_1^2, \ldots, s_n^2)$. The diagonal elements s_i^2, $i = 1, \ldots, n$, are the AR(4) residual variances defined above.

We briefly discuss the posterior simulation within an MCMC algorithm. All conditional posterior distributions are available in closed form, rendering a Gibbs sampler convenient. Although the conditional posterior distributions of the VAR coefficients and the covariance matrix are standard, we use the auxiliary sampler of Makalic and Schmidt (2016) to sample from the conditional posterior densities for the HS prior.

We also provide more details on identification via external instruments in VARs, as extensively discussed in Stock and Watson (2018) and Montiel Olea et al. (2021). We assume the usual mapping between the reduced-form errors and the structural innovations, $\boldsymbol{u}_t = \boldsymbol{S}\boldsymbol{\varepsilon}_t$. The structural impact matrix $\boldsymbol{S}$ is part of the decomposition of the covariance matrix $\boldsymbol{\Sigma} = \boldsymbol{S}\boldsymbol{S}'$. We are interested in identifying the structural impact vector $\boldsymbol{s}^{\text{ES}} \in \mathbb{R}^n$, defined through $\boldsymbol{S} = (\boldsymbol{s}^{\text{ES}}, \boldsymbol{s}^{-\text{ES}})$ (w.l.o.g.), where $\boldsymbol{s}^{-\text{ES}} \in \mathbb{R}^{n\times(n-1)}$. To do so, we first partition the vector of structural shocks as $\boldsymbol{\varepsilon}_t = (\varepsilon_t^{\text{ES}}, \boldsymbol{\varepsilon}_t^{-\text{ES}\prime})'$ (w.l.o.g.), where $\boldsymbol{\varepsilon}_t^{\text{ES}} \in \mathbb{R}$ denotes the respective energy supply shock and $\boldsymbol{\varepsilon}_t^{-\text{ES}} \in \mathbb{R}^{n-1}$ denotes all other shocks. We denote the external commodity price instrument by ω_t, which must be exogenous and relevant to be a valid instrument. These conditions can be summarized in the following equations:

$$\mathbb{E}[\omega_t \varepsilon_t^{\text{ES}}] = \alpha \neq 0, \tag{B.7}$$

$$\mathbb{E}[\omega_t \boldsymbol{\varepsilon}_t^{-\text{ES}}] = \boldsymbol{0}. \tag{B.8}$$

Eq. (B.7) states that ω_t is correlated with the energy supply shock (the instrument is *relevant*), while Eq. (B.8) states that ω_t is uncorrelated with the other shocks (the instrument is *exogenous*). The instrument is plausibly exogenous, as the high-frequency nature of the energy supply shocks captures surprises in very narrow time windows around announcements. It would therefore appear unlikely that other structural shocks can significantly affect the economy at the same time, so these other shocks should be uncorrelated with ω_t. Given the assumptions above, this allows us to identify $\boldsymbol{s}^{\text{ES}}$ up to scale by the covariance between ω_t and the reduced-form innovations $\boldsymbol{u}_t$:

$$\boldsymbol{\Gamma} = \mathbb{E}[\omega_t \boldsymbol{u}_t] = \mathbb{E}[\omega_t \boldsymbol{S}\boldsymbol{\varepsilon}_t] = \boldsymbol{\alpha}\boldsymbol{s}^{\text{ES}}, \tag{B.9}$$

which yields $\tilde{\boldsymbol{s}}^{\text{ES}} = \boldsymbol{\Gamma}/(\boldsymbol{e}_k'\boldsymbol{\Gamma}) = \boldsymbol{s}^{\text{ES}}/(\boldsymbol{e}_k'\boldsymbol{s}^{\text{ES}})$ where $\boldsymbol{e}_k$ is a vector of zeros with unity in the k-th position and selects the relevant commodity-price equation. This normalizes the impact response of the commodity price variable to one, which we then rescale to obtain a 10% commodity price increase. This equation implies performing the following first-stage regression of the reduced-form innovations on the instrument:

$$\boldsymbol{u}_t = \boldsymbol{\beta}\omega_t + \nu_t, \tag{B.10}$$

and for a demeaned instrument, we have that $\boldsymbol{\beta} = \boldsymbol{\Gamma}/\mathbb{E}[\omega_t^2]$.

B.2 Log-quantile density transformation

The functional modeling of a distribution is difficult due to two constraints: non-negativity and unit-integration. The literature has adopted the approximation via the logarithm of the density function, which naturally enforces the non-negativity constraint. However, this is not sufficient to enforce unit-integration. The problem originates from the fact that standard functional data methods (e.g., FPCA) are designed to model the behavior of random functions in the Hilbert space L^2, while the space of densities is only a subset of L^2.

Petersen and Müller (2016) propose representing densities through their *log-quantile density* (LQD) functions in the Hilbert space L^2. Let us define the function $g(\cdot)$ as the mapping that defines the LQD associated with the probability density function $p_t(x)$. The LQD is the logarithm of the first derivative of the *quantile function* $Q_t(z) = F_t^{-1}(z)$:

$$f_t(z) = g\{p_t(\cdot)\}(z) = \log\Big(\frac{d}{dz}Q_t(z)\Big) = \log Q_t'(z) = -\log p_t(Q_t(z)), \qquad z \in [0,1]. \tag{B.11}$$

Although the LQD representation avoids directly imposing the density constraints, an exact LQD on the fixed support $[L,U]$ satisfies $\int_0^1 \exp(f_t(z))dz = U - L$. Furthermore, the distribution of interest $p_t(\cdot)$ can be derived from $f_t(z)$ using the inverse transformation.

Derivation. Let us define the *probability density function* (pdf) $p_t(x)$ on $[L,U]$, where non-negativity $p_t(x) \geq 0$ and unit-integration $\int_L^U p_t(x)dx = 1$ hold. Define the *cumulative distribution function* (cdf) as $F_t(x) = \int_L^x p_t(s)ds$. The mapping $F_t : [L,U] \to [0,1]$ is non-decreasing and differentiable with $F_t'(x) = p_t(x)$. Let us further define the *quantile function* (or inverse cdf) as $Q_t(z) = F_t^{-1}(z)$ where $z \in [0,1]$. The mapping $Q_t : [0,1] \to [L,U]$ is non-decreasing as well and has boundary constraints $Q(0) = L$ and $Q(1) = U$. By definition, $F_t(Q_t(z)) = z$. Now, differentiate this identity with respect to z:

$$\begin{aligned}
\frac{d}{dz}F_t(Q_t(z)) &= \frac{d}{dz}z \\
F_t'(Q_t(z))Q_t'(z) &= 1 \\
p_t(Q_t(z))Q_t'(z) &= 1 \\
Q_t'(z) &= \frac{1}{p_t(Q_t(z))}
\end{aligned}$$

This allows us to express the log-quantile density (LQD) transformation as the logarithm of the first derivative of the quantile function:

$$f_t(z) := \log Q_t'(z) = -\log p_t(Q_t(z)), \qquad z \in [0,1]. \tag{B.12}$$

We model the transformed functions in $L^2([0,1])$ and enforce the support constraint through the normalized reconstruction described below.

To find the inverse mapping, we have that $Q_t'(z) = \exp(f_t(z))$. We simply have to find $Q_t(z) = Q_t(0) + \int_0^z Q_t'(s)ds$. For an exact LQD function, the boundary conditions imply $\int_0^1 \exp(f_t(z))dz = U - L$. Approximation and model-implied changes in the transformed function need not preserve this constraint. We therefore use a normalized reconstruction that ensures the reconstructed quantile function satisfies $Q_t(0) = L$ and $Q_t(1) = U$, restoring the prescribed support $[L, U]$. The corresponding reconstructed quantile function is

$$Q_t(z) = L + (U - L)\frac{\int_0^z \exp(f_t(u))du}{\int_0^1 \exp(f_t(u))du},$$

where the denominator is a normalizing constant. Differentiate:

$$Q_t'(z) = (U - L)\frac{\exp(f_t(z))}{\int_0^1 \exp(f_t(u))du}.$$

This allows us to recover the density along the quantile curve:

$$p_t(Q_t(z)) = \frac{1}{Q_t'(z)} = \frac{\int_0^1 \exp(f_t(u))du}{(U - L)\exp(f_t(z))} \tag{B.13}$$

Example: We simulate data from a LogNormal distribution whose logarithm has a mean of zero and a standard deviation of 0.5, truncated to have support in $[0,4]$.

Figure B1: LQD transformation: example.

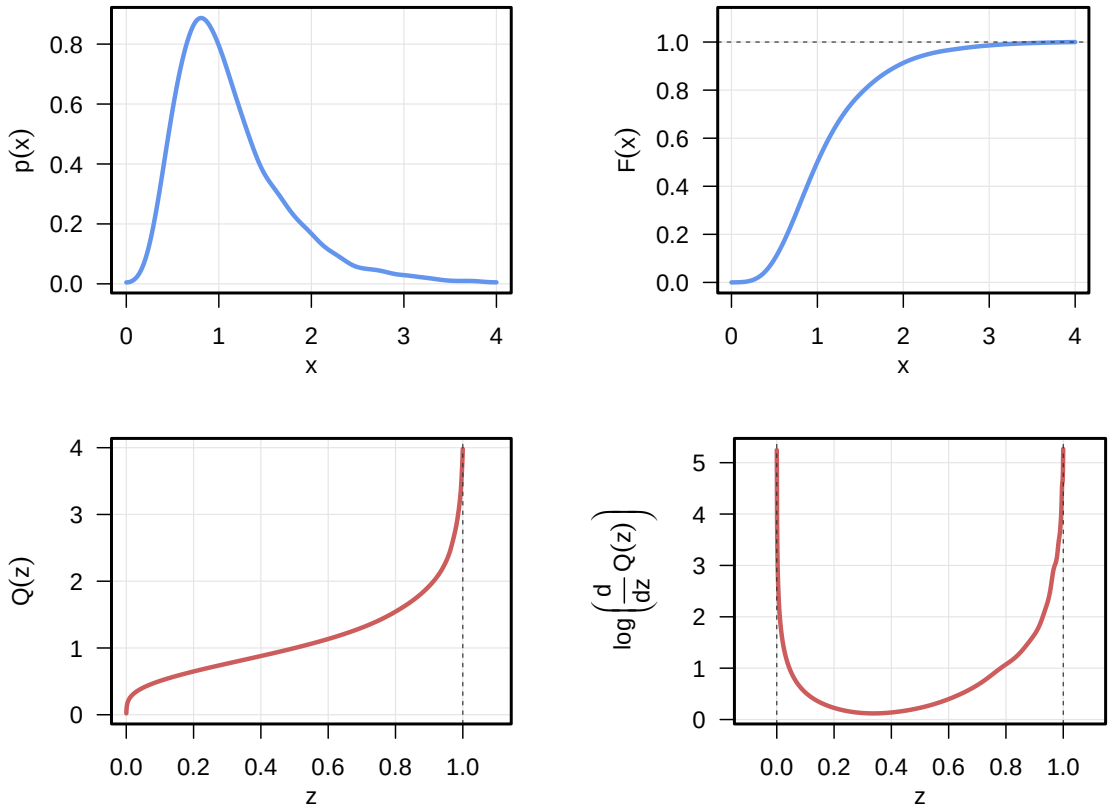

Notes: The figure displays four different representations of a LogNormal$(0,0.5)$ distribution restricted to have support in $\mathscr{X} = [0,4]$. The upper-left panel depicts the probability density function $p(x)$; the upper-right panel depicts the cumulative distribution function $F(x)$; the lower-left panel depicts the quantile function; and the lower-right panel depicts the log-quantile density function.

C. Additional empirical results

Figure C1: Time series of functional principal components.

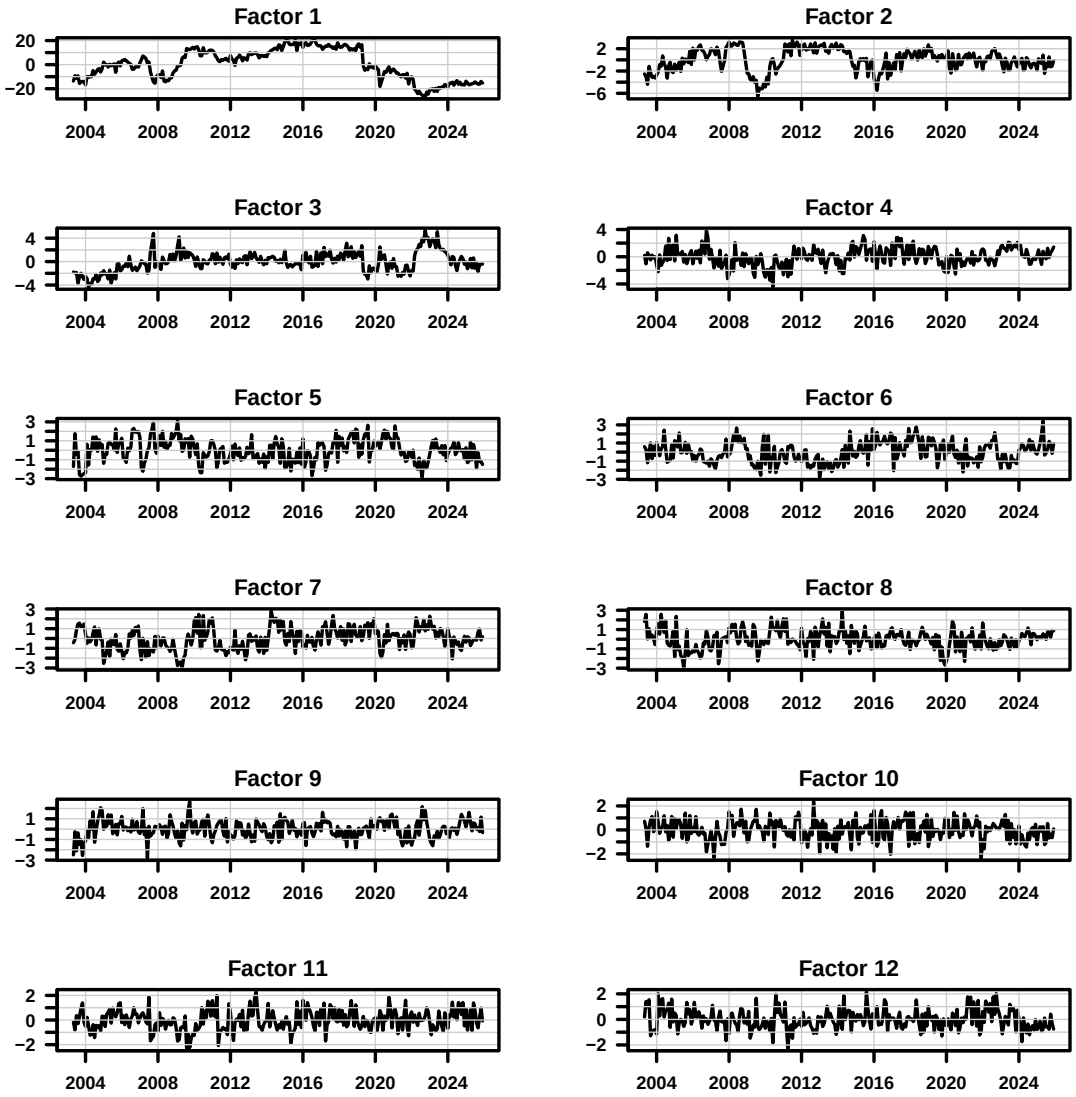

Notes: The figure shows the time series of the first twelve FPCs obtained for the inflation expectation data. The method and particular specification choices are explained in Section 3.

Figure C2: Sequences of oil supply news shocks and natural gas supply shocks.

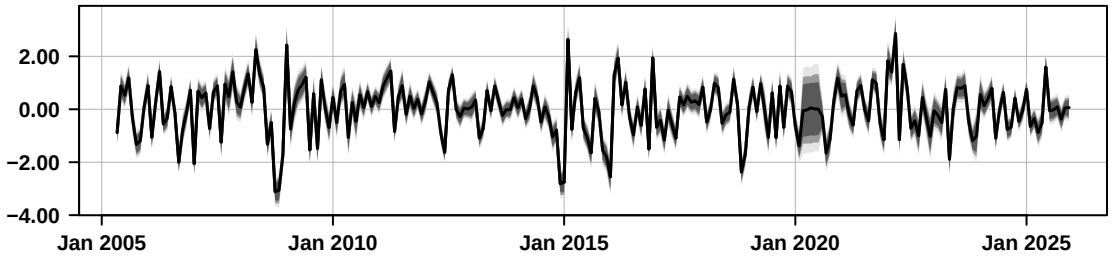

(a) Oil supply news shock.

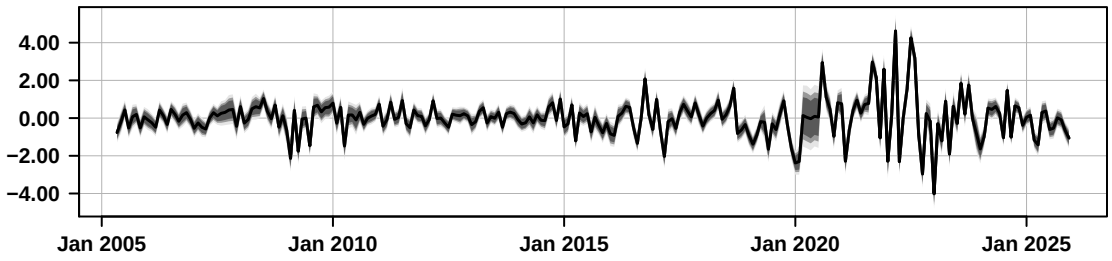

(b) Natural gas supply shock.

Notes: The figure shows the estimated sequences of oil supply news shocks (panel a) and natural gas supply shocks (panel b). The estimation covers the period 2003M5 to 2025M12. The black solid lines refer to the median estimates. The grey areas refer to the 68%/80%/90% credible sets. The spell of zero shocks in 2020 is due to the inclusion of dummies that allow us to minimize the effect of the Covid-19 pandemic on our estimates.

Figure C3: Impulse responses of the FPCs of the inflation expectation distribution to an oil supply news shock.

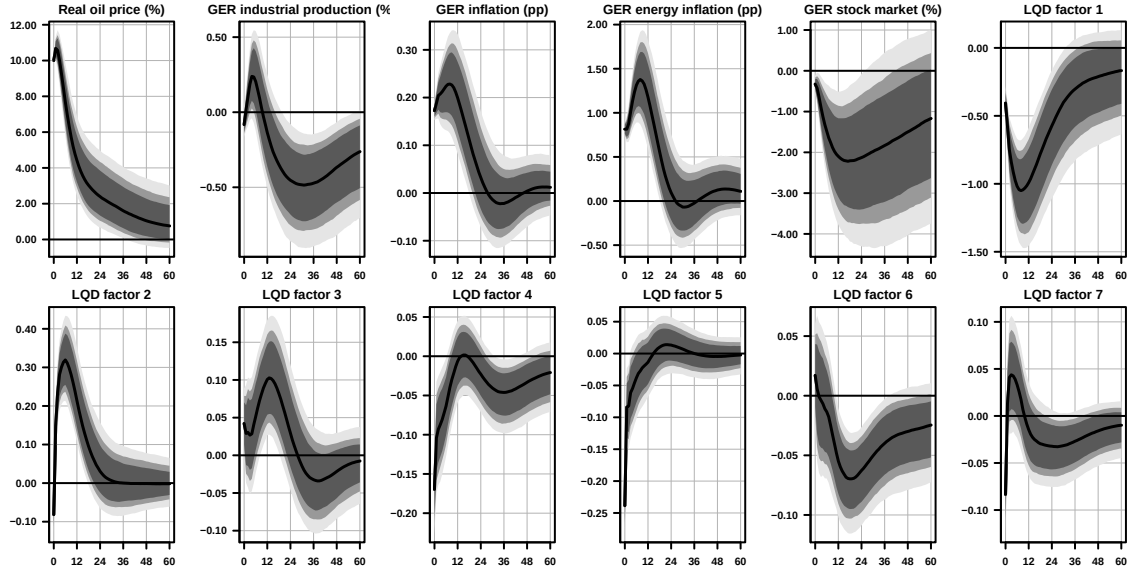

Notes: The figure shows the impulse responses of the FPCs of the inflation expectation distribution to an oil supply news shock. The shock is normalized to induce a 10% increase in the real price of oil. The estimation covers the period 2003M5 to 2025M12. The black solid lines refer to the median estimates. The grey areas refer to the 68%/80%/90% credible sets. The horizontal axis measures the impulse horizon in months. The vertical axis denotes deviations from trend/steady state in percent (real oil price, industrial production, stock market), percentage points (inflation, energy inflation), and factor units (LQD factors).

Figure C4: Impulse responses of interval probabilities for the inflation expectation distribution to an oil supply news shock.

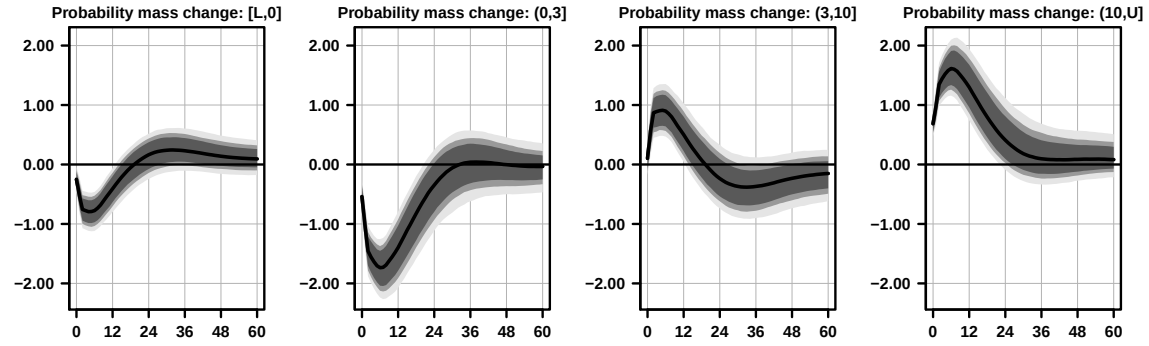

Notes: The figure shows the impulse responses of selected interval probabilities for the inflation expectation distribution to an oil supply news shock. The shock is normalized to induce a 10% increase in the real price of oil. The estimation covers the period 2003M5 to 2025M12. The black solid lines refer to the median estimates. The grey areas refer to the 68%/80%/90% credible sets. The horizontal axis measures the impulse horizon in months. The vertical axis denotes deviations from steady state in percentage points.

Figure C5: Impulse responses of macroeconomic aggregates and moments of the inflation expectations distribution to an oil supply shock.

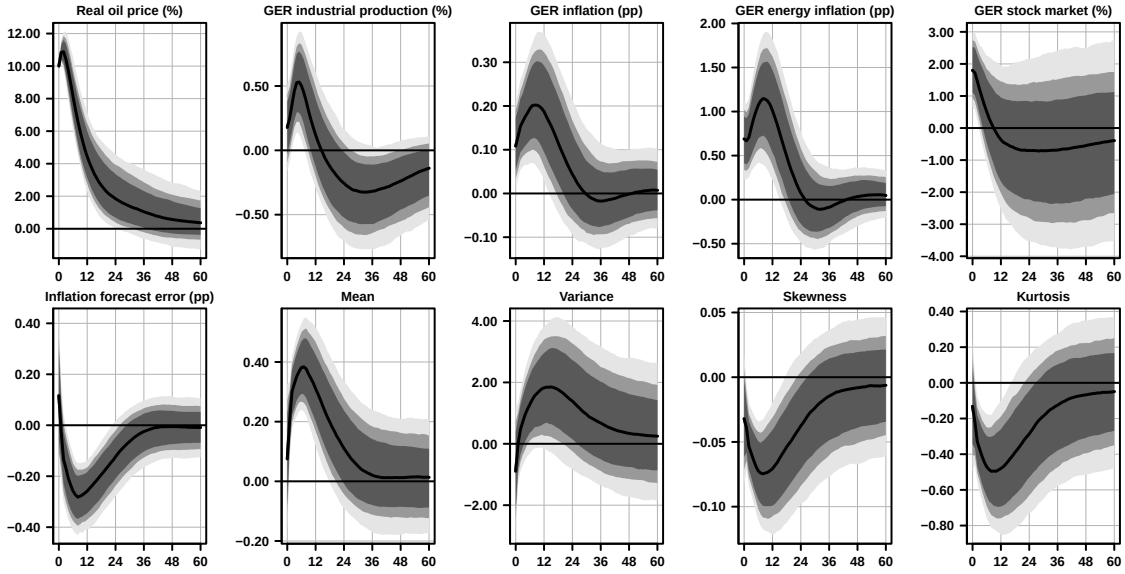

Notes: The figure shows the impulse responses of macroeconomic aggregates and the first four moments of the inflation expectation distribution to an oil supply shock. The shock is normalized to induce a 10% increase in the real price of oil. The estimation covers the period 2003M5 to 2025M12. The black solid lines refer to the median estimates. The grey areas refer to the 68%/80%/90% credible sets. The horizontal axis measures the impulse horizon in months. The vertical axis denotes deviations from trend/steady state in absolute values (skewness, kurtosis), percent (real oil price, industrial production, stock market), percentage points (inflation, energy inflation, mean) or squared percentage points (variance). The forecast error is defined as $FE_t = \pi_{t+12} - \bar{x}_t$, where $\bar{x}_t$ denotes the mean inflation expectation at time t.

Figure C6: Responses of the inflation expectation distribution to an oil supply shock.

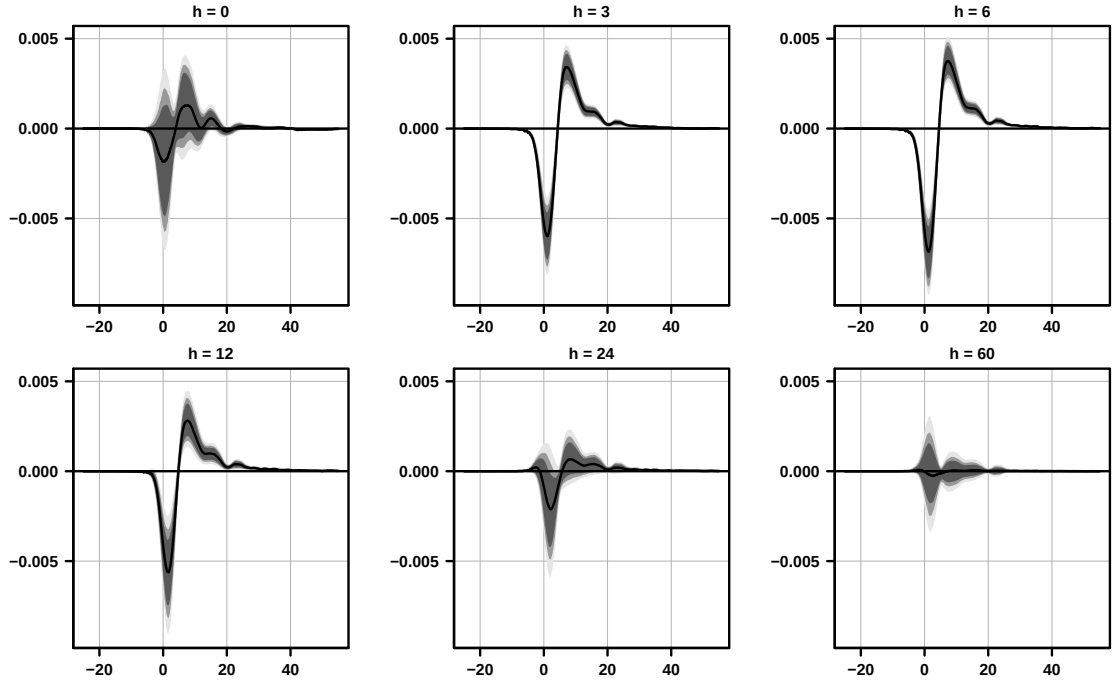

Notes: The figure shows the difference between the inflation expectation distribution after an oil supply shock and the one prevailing in the steady state. The shock is normalized to induce a 10% increase in the real price of oil. The estimation covers the period 2003M5 to 2025M12. The black solid lines refer to the median estimates. The grey areas refer to the 68%/80%/90% credible sets. The horizontal axis measures inflation expectations in percent. The vertical axis denotes the effect on the density $p_t(x)$.

Figure C7: Impulse responses of the FPCs of the inflation expectation distribution to a natural gas supply shock.

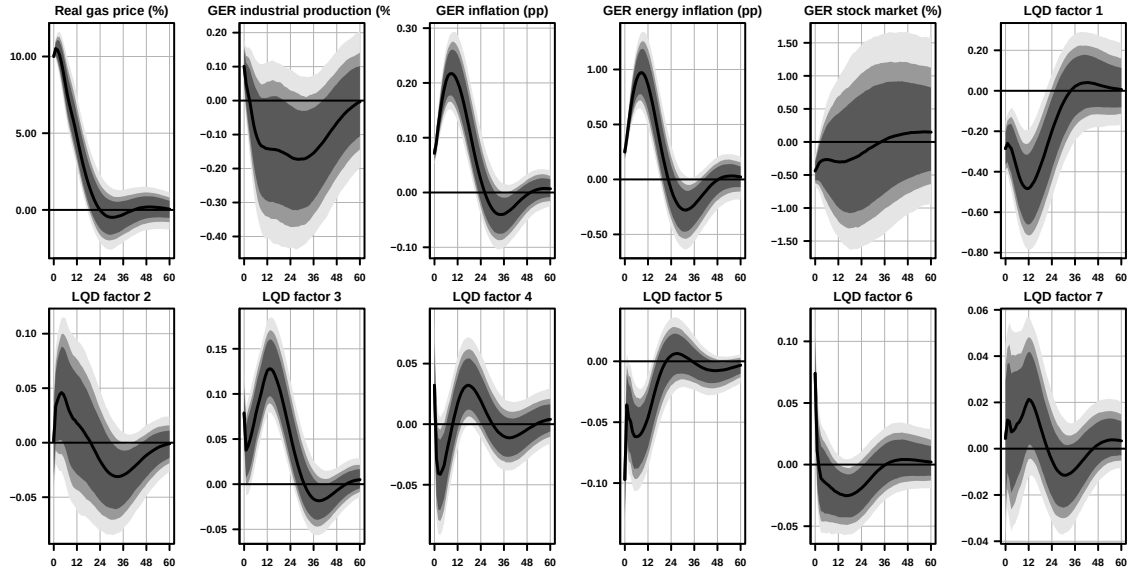

Notes: The figure shows the impulse responses of the FPCs of the inflation expectation distribution to a natural gas supply shock. The shock is normalized to induce a 10% increase in the real price of natural gas. The estimation covers the period 2003M5 to 2025M12. The black solid lines refer to the median estimates. The grey areas refer to the 68%/80%/90% credible sets. The horizontal axis measures the impulse horizon in months. The vertical axis denotes deviations from trend/steady state in percent (real gas price, industrial production, stock market), percentage points (inflation, energy inflation), and factor units (LQD factors).

Figure C8: Impulse responses of interval probabilities for the inflation expectation distribution to a natural gas supply shock.

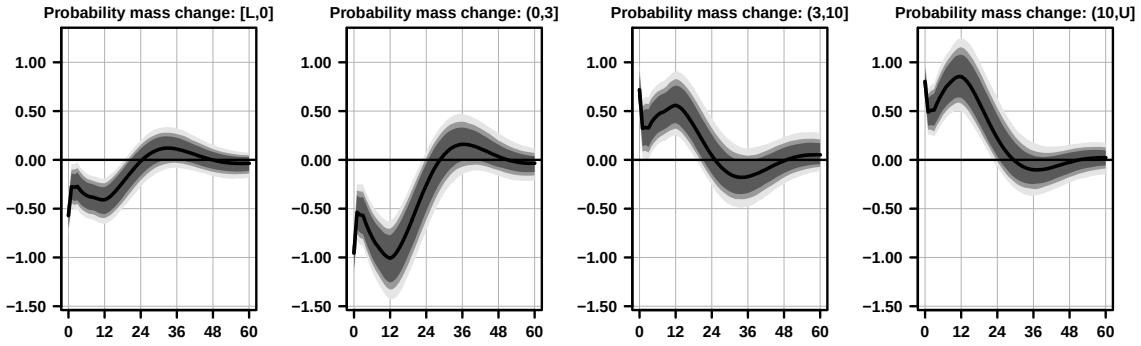

Notes: The figure shows the impulse responses of selected interval probabilities for the inflation expectation distribution to a natural gas supply shock. The shock is normalized to induce a 10% increase in the real price of natural gas. The estimation covers the period 2003M5 to 2025M12. The black solid lines refer to the median estimates. The grey areas refer to the 68%/80%/90% credible sets. The horizontal axis measures the impulse horizon in months. The vertical axis denotes deviations from steady state in percentage points.

Figure C9: Impulse responses of macroeconomic aggregates and moments of the inflation expectations distribution in the joint, large FVAR.

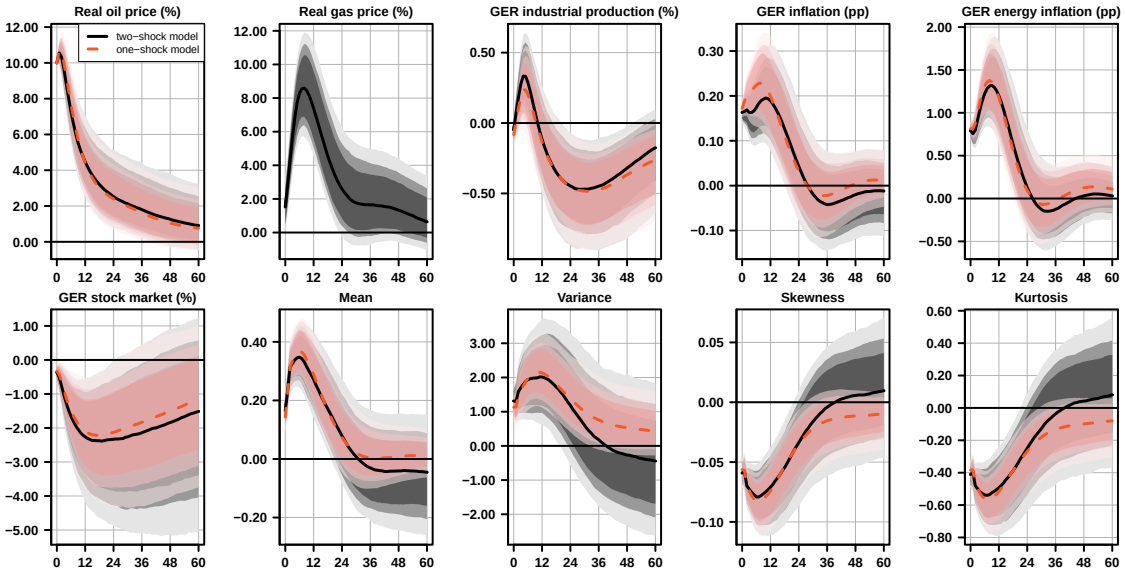

(a) Oil supply news shock.

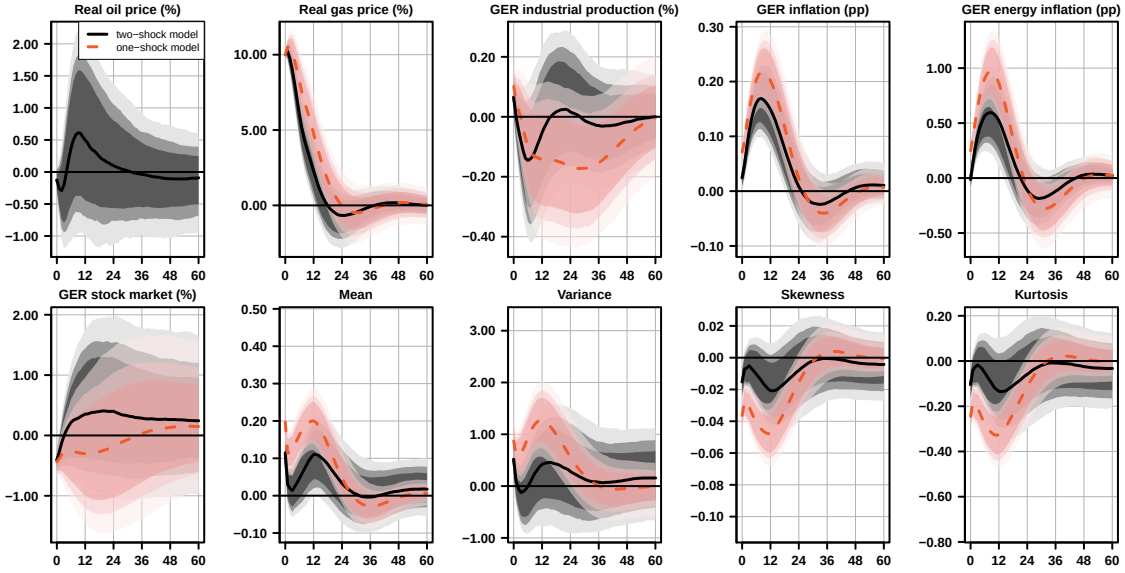

(b) Natural gas supply shock.

Notes: The figure shows the impulse responses of macroeconomic aggregates and the first four moments of the inflation expectation distribution to an oil supply news shock and a natural gas supply shock, respectively. Both shocks are normalized to induce a 10% increase in the real price of oil or natural gas. The estimation covers the period 2003M5 to 2025M12. The black solid lines refer to the median estimates. The grey areas refer to the 68%/80%/90% credible sets. The red dashed lines refer to the model with only one shock, while the red areas refer to the 68%/80%/90% credible sets. The horizontal axis measures the impulse horizon in months. The vertical axis denotes deviations from trend/steady state in absolute values (skewness, kurtosis), percent (real oil price, real gas price, industrial production, stock market), percentage points (inflation, energy inflation, mean) or squared percentage points (variance).

Appendix References

Carvalho C M., Polson N G., and Scott J G. (2010) "The Horseshoe Estimator for Sparse Signals". *Biometrika*, Vol. 97(2), pp. 465–480.

Dieppe A., Legrand R., and Van Roye B. (2018) *The Bayesian Estimation, Analysis and Regression (BEAR) Toolbox: Technical Guide*. Technical Document, BEAR Toolbox, European Central Bank. European Central Bank.

Kadiyala K R. and Karlsson S. (1997) "Numerical Methods for Estimation and Inference in Bayesian VAR Models". *Journal of Applied Econometrics*, Vol. 12(2), pp. 99–132.

Litterman R B. (1986) "Forecasting with Bayesian Vector Autoregressions—Five Years of Experience". *Journal of Business & Economic Statistics*, Vol. 4(1), pp. 25–38.

Makalic E. and Schmidt D F. (2016) "A Simple Sampler for the Horseshoe Estimator". *IEEE Signal Processing Letters*, Vol. 23(1), pp. 179–182.

Montiel Olea J L., Stock J H., and Watson M W. (2021) "Inference in Structural Vector Autoregressions Identified with an External Instrument". *Journal of Econometrics*, Vol. 225(1), pp. 74–87.

Petersen A. and Müller H.-G. (2016) "Functional Data Analysis for Density Functions by Transformation to a Hilbert Space". *The Annals of Statistics*, Vol. 44(1), pp. 183–218.

Stock J H. and Watson M W. (2018) "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments". *The Economic Journal*, Vol. 128(610), pp. 917–948.