DISCUSSION PAPER SERIES

DP18937

THE MACROECONOMIC EFFECTS OF INFLATION EXPECTATIONS: THE DISTRIBUTION MATTERS

Guido Ascari, Paolo Bonomolo and Alessandro Celani

INTERNATIONAL MACROECONOMICS AND FINANCE AND MONETARY ECONOMICS AND FLUCTUATIONS

CEPR

ISSN 0265-8003

THE MACROECONOMIC EFFECTS OF INFLATION EXPECTATIONS: THE DISTRIBUTION MATTERS

Guido Ascari, Paolo Bonomolo and Alessandro Celani

Discussion Paper DP18937
Published 21 March 2024
Submitted 20 March 2024

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- International Macroeconomics and Finance
- Monetary Economics and Fluctuations

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Guido Ascari, Paolo Bonomolo and Alessandro Celani

THE MACROECONOMIC EFFECTS OF INFLATION EXPECTATIONS: THE DISTRIBUTION MATTERS

Abstract

We augment a parsimonious monetary policy Bayesian VAR with heterogeneous expectations from the Michigan Survey to investigate the macroeconomic effects of shocks to the distribution of short term inflation expectations. A first surprising result is that (the Michigan survey) inflation expectations do not seem to be much influenced by macroeconomic developments, while the opposite is not true. Moreover, a comprehensive density impulse response function analysis shows that it matters to take into account the whole expectation distribution. First, it matters because considering only the first and second moment of the distribution leads to an underestimation of the macroeconomic effects of expectations shocks. Second, mean and variance shocks are stagflationary, while dispersion shocks might be recessionary. Third, the effects are sharper when the shock mass is condensed on the tails. Specifically, left-tail perturbations account for the largest effect of expectations shocks on macroeconomic fluctuations. It follows that central bank communication should focus on the tails: reducing the noise/dispersion might be more effective than anchoring the mean.

JEL Classification: E3, E52, C32

Keywords: Inflation, Expectations, Distribution

Guido Ascari - guido.ascari@unipv.it
De Nederlandsche Bank and University of Pavia and CEPR

Paolo Bonomolo - p.bonomolo@dnb.nl
De Nederlandsche Bank

Alessandro Celani - alecelani94@gmail.com
De Nederlandsche Bank

The Macroeconomic Effects of Inflation Expectations: The Distribution Matters*

Guido Ascari†
De Nederlandsche Bank
University of Pavia, CEPR

Paolo Bonomolo‡
De Nederlandsche Bank

Alessandro Celani§
De Nederlandsche Bank

March 20, 2024

Abstract

We augment a parsimonious monetary policy Bayesian VAR with heterogeneous expectations from the Michigan Survey to investigate the macroeconomic effects of shocks to the distribution of short term inflation expectations. A first surprising result is that (the Michigan survey) inflation expectations do not seem to be much influenced by macroeconomic developments, while the opposite is not true. Moreover, a comprehensive density impulse response function analysis shows that it matters to take into account the whole expectation distribution. First, it matters because considering only the first and second moment of the distribution leads to an underestimation of the macroeconomic effects of expectations shocks. Second, mean and variance shocks are stagflationary, while dispersion shocks might be recessionary. Third, the effects are sharper when the shock mass is condensed on the tails. Specifically, left-tail perturbations account for the largest effect of expectations shocks on macroeconomic fluctuations. It follows that central bank communication should focus on the tails: reducing the noise/dispersion might be more effective than anchoring the mean.

Keywords: Inflation, Expectations, Distribution, Monetary Policy.
JEL classification: E3, E52, C32.

*We wish to thank Fabio Canova, Fadi Hassan, Francesca Monti and Ricardo Reis for their comments and seminar participants at the Bank of Italy, the Dutch Central Bank, the Swiss Central Bank, seminars and conference participants at the 17th Conference on Computational and Financial Econometrics 2023. The views expressed are those of the authors and do not necessarily reflect official positions of De Nederlandsche Bank.

†Guido Ascari, De Nederlandsche Bank, Sparkelweg, Amsterdam and Department of Economics and Management, University of Pavia, Via San Felice 5, 27100 Pavia, Italy, Email: g.ascari@dnb.nl, guido.ascari@unipv.it.

‡Paolo Bonomolo, De Nederlandsche Bank, Sparkelweg, Amsterdam. Email: p.bonomolo@dnb.nl.

§Alessandro Celani, De Nederlandsche Bank, Sparkelweg, Amsterdam. Email: a.celani@dnb.nl.

1 Introduction

What is the macroeconomic impact of a shock that changes the short-term inflation expectation distribution?

The literature on inflation expectations is at the core of macroeconomics, because inflation expectations play a crucial role in theoretical models and policy discussions. Recently, this literature thrived with important theoretical contributions and empirical analysis, thanks to the increasing availability of surveys. The large recent literature on inflation expectations takes mainly a microeconomic perspective, focusing on understanding the formation and the determinants of inflation expectations.

While inflation expectations are obviously endogenous, one should not discard the possibility that they can also move exogenously or independently from macroeconomic developments. In their survey, D'Acunto et al. (2022) conclude that agents' expectations are biased and volatile in the time series, which suggests that inflation expectations might be subject to exogenous shocks. Moreover, this idea seems embedded in the way policy makers think about expectations in their speeches. These often feature discussions about the possibility of inflation expectations getting out of control, as if policy has to respond to counteract exogenous shifts in inflation expectations. This would not make sense in an environment of rational, or more general well-behaved, inflation expectation formation process.

Hence, following Ascari et al. (2023), we take a macroeconomic perspective and revert the common question about how expectations depend on the economic outlook to : how exogenous shifts in expectations affects macroeconomic outcomes? The literature on surveys indeed showed that exogenous variations in inflation expectations do affect agents' economic consumption and investment decisions, and that expectations about future inflation are associated with worse expected macroeconomic outcome (e.g., Coibion et al., 2019, 2023; Weber et al., 2022). Ascari et al. (2023) show the theoretical and empirical relevance of exogenous variations of inflation expectations. Specifically, a shock that increases the average short-term inflation expectation has negative macroeconomic effects, increasing inflation and decreasing output.

However, we know that the rich information set contained in survey expectations is not limited to the consensus forecast. Whether of households, firms, or professional forecasters, we observe a pervasive cross-sectional heterogeneity. Possible explanations include substantial inattentiveness to the temporal evolution of macroeconomic aggregates, sensitivity to salient prices or personal experiences, and a variety of different biases (e.g., D'Acunto et al., 2022). This dispersion in beliefs, or disagreement (Mankiw et al., 2003), exhibits substantial time variation. Reis (2022) convincingly argues that there is important information content contained in the higher order moments of the inflation expectation distribution.

Hence, we investigate the macroeconomic effects of various shocks to the short-term inflation expectation distributions, using a Bayesian VAR (BVAR) to study the joint evolution of a number of macroeconomic aggregates and the short-term inflation expectation distribution. In order to have a sufficiently large number of survey respondents to meaningful talk about higher moments and a long sample size, we use the Michigan Survey of Consumers (MSC). The main obstacle we face is how to incorporate the rich information set arising from individual expectations in a macroeconometric model. By assuming that the large cross-sections of the expectations originate from a nonparametric distribution, we approximate the theoretical quantile function with a rich number of empirical quantiles. The resulting dimensionality problem is tackled with a Bayesian hierarchical framework, whereas the pathology of quantile crossing is handled with a weighted Kernel smoothing procedure. This allows us to map a broad range of heterogeneity in the expectation formation process, thus, improving from studies employing just the consensus, and contributing to the emerging literature concerned with the joint estimation of a time series of densities and macroeconomic aggregates (Chang et al., 2023; Meeks and Monti, 2023).

We use this setup to study the cross-feedback effects between expectations and macroeconomy. Our identification assumptions are based on the natural discrepancy between the date on which the survey takes place and the release of the macroeconomic variables. A first notable, and somewhat surprising result, is that the inclusion of expectations is relevant to explain the macroeconomic variables, but not vice versa.[1] This result provides a strong rationale to study the main question of our paper stated at the beginning of this Introduction.

We therefore investigate how several exogenous movements of the inflation expectation distribution affect output and inflation, through a comprehensive density impulse response function analysis. First, we study the variation of the mean, the variance, as well as higher-order moments. Shocks that increase the first two moments of the distribution are stagflationary. Moreover, our analysis shows that incorporating the whole cross-sectional heterogeneity is important, as models based on just mean and variance would underestimate the macroeconomic effects of inflation expectation shocks to these moments. Regarding the possible transmission mechanism, we include the consumer sentiment variable from the MSC. This variable drops after a positive shock to the first and second moment of the distribution, signalling a bad expected future outcome by consumers, whose consequent behaviour could trigger the effects on inflation and output.

Second, we look at various shocks that change the skewness and the kurtosis of the distribution. Unlike location, dispersion is a more general concept and there are many ways to perturb the shape of the distribution. With quantiles in hand, we are able to generate movements related to specific parts of the distribution. We first look at symmetric shocks, showing that tails shock – that is, perturbation where more mass is distributed to the tails – are recessionary. Then, we investigate asymmetry to identify whether one part of the distribution is more relevant in inducing these effects. Our findings show that the left tail of the distribution is the only driver of the effects on the macroeconomic variables that we found in the symmetric case. On the contrary, movements in the right tail of the distribution have negligible effects. Coherently, consumer sentiment drastically drops after a left tail shock, while it reacts positively (even if negligibly) in case of a right tail shock.

Note that exogenous perturbation of distributional datasets requires modification in the way experiments are carried out. Another contribution we make is to show how to conduct dynamic analysis when an exogenous movement of the distribution is of interest. In particular, given an initial distribution of interest obtained perturbing the steady-state distribution, we compute

[1] This is true both in terms of forecast error variance decomposition and in terms of the reduced -form shrinkage parameters of the expectation versus the macroeconomic variable block of the BVAR. Note that the latter result does not depend on the identification assumptions.

the implied combination of exogenous distributional shocks. Since the number of exogenous shocks to the distribution is higher that the quantiles involved, this task requires a structural scenario analysis in the spirit of Antolín-Díaz et al. (2021).

These results call for an investigation about who are the households populating the left tail of the inflation expectation distribution. The MSC is a comprehensive survey, collecting inflation expectations, but also a variety of other information. It turns out that the probability of being in the left tail is positively related to have a college degree and to the level of income, suggesting that our findings are not due to poorly educated households with no spending ability.

Finally, we discuss some possible implications of our results for central bank communication. Specifically, which part of the distribution should central bank communication target, assuming that the central bank could change the shape of inflation distribution through effective communication? Should communication just focus on shifting the mean of the distribution towards the 2% target? We devise a combination of mean and dispersion shocks such that the distribution on impact is anchored at 2%, with different degrees of cross-sectional heterogeneity. As long as the dispersion around the mean is relatively low, shifting the mean towards 2% has positive macroeconomic effects. However, the results change if shifting the mean is accompanied by an increase in the dispersion. Then, economic conditions deteriorate substantially. Our results therefore suggest that, with respect to short-term inflation expectations, communication should focus on decrease the dispersion, and particularly moving mass from the left tail, where pessimistic expectations lie, rather than move the mean towards the 2% target. Interestingly, this echoes the discussion in Blinder et al. (2008) about the importance of central bank communication to reduce noise and thus raise the signal-to-noise ratio.

Literature. The literature on inflation expectations is vast, there are excellent surveys (i.e., Coibion et al., 2018; Weber et al., 2022; D'Acunto et al., 2022), such that there is no much point in summarizing it here. Moreover, as said, most of the literature focuses on the effects of various factors (e.g., macroeconomic devlopments, beliefs, personal experience, etc.) in shaping inflation expectations, while here we rather invert the causation asking whether and how exogenous changes in inflation expectations have an impact on macroeconomic variables.

Surprisingly, there are few papers in the literature investigating this question. Clark and Davig (2011) is an early work that looks at short-term inflation expectations in a VAR analysis, but it does not focus on the macroeconomic effects. Ascari et al. (2023) study the effects of a shock to the mean of short-term inflation expectations both theoretically and empirically using data from the Survey of Professional Forecasters in the US. They find that such a shock is stagflationary, as we do. Barrett and Adams (2022) analyzes the same topic with a focus on the empirical identification of such a shock, reaching opposite empirical conclusions. Neri (2023) asks what are the macroeconomic effects of changes in long-term inflation expectations in an empirical VAR on Euro area data. Given the focus on the long-term inflation expectations, Neri (2023) connects to a large literature studying the effects of 'de-anchoring'.

Contrary to the paper above, our paper focuses, on the whole distribution and not just on the consensus forecast. Reis (2022) was really an inspiring work for our analysis. He shows that inflation expectations seemed anchor in 1971 by looking at the cross-sectional mean of professional forecasters' or households' expectations, but looking at the whole distribution would have revealed a very different picture providing timely sign of a shift in inflation expectations. Reis (2022) has two important takeaways for our work: (i) movements in the the whole distribution of inflation expectations provide important signals about future inflation; (ii) there is a lot of information in the tails of the distribution. Motivated by these two points, our analysis reinforces the Reis (2022) conclusion, and not only for inflation: inflation expectations distribution matters for macroeconomic dynamics.

Meeks and Monti (2023) is particularly related to our work. They use functional principal component regression to fit an augmented Phillips Curve, where the distribution of short-term inflation expectations appears on the right-hand side. They find statistical relevance of considering the heterogeneity of expectations for inflation dynamics. Chang et al. (2023) propose a State-Space model, where the dynamics of macroeconomic aggregates and log densities form the state-transition equation. The infinite-dimensional log densities are discretized through cubic splines over a certain interval.

From a methodological standpoint, our work is mostly related to the emerging literature on the joint estimate of macroeconomic aggregates and distributional data derived from individual attributes. Statisticians are increasingly interested in analyzing samples of random objects that do not belong to vector spaces, such as univariate probability measures (Kokoszka et al., 2019; Chen et al., 2021; Zhang et al., 2022; Petersen et al., 2022; Zhu and Müller, 2023). On a broader level, one can conceive densities as a specific class of functions featuring non-negative and integration constraints. The more general statistical examination of functional data resulted in a large literature, comprehensively treated by Ramsay and Silverman (2005) and Horváth and Kokoszka (2012), among others. In economics, functional data emerge in the context of yield curves (Diebold and Li, 2006; Inoue and Rossi, 2019). Methods derived from functional data analysis could also be used to deal with probability measures, as in Meeks and Monti (2023) or Chang et al. (2023) quoted above. In these cases, proper densities are restored from the estimated functions by ex-post renormalizations, which usually generate negligible approximation errors.

Finally, to some extent, our work is related to an extensive recent literature on 'expectational shocks', where exogenous changes in expectations are assumed to drive economic fluctuations, through belief, sentiment, confidence or news shocks.[2]

In what follows, Section 2 presents the data, Section 3 the methodology, Section 4 the results, and Section 5 concludes.

2 Data

We estimate the model on monthly US data, using a number of macroeconomic variables and microdata from the MSC. The macroeconomic variables are: Real Oil Price (OP), Consumer Price Index (CPI), Industrial Production (IP), and Nominal Interest Rate (IR). The period

[2]This literature investigates, both empirically and theoretically, how news about future TFP (e.g., Beaudry and Portier, 2006, 2014; Barsky and Sims, 2011, 2012) or exogenous waves of optimism and pessimism due to sentiment (e.g., Benhabib et al., 2015; Angeletos et al., 2018) affect business cycle fluctuations.

considered is from January 1983 to December 2019, for a total of $T = 422$ observations.[3]

From the MSC, we use microdata on inflation expectations, and consumer sentiment. Each consumer is asked to respond about the future expected inflation. In this study we focus on expectations over the next 12 months, i.e. one year ahead inflation expectations.[4] The consumer sentiment is constructed in the MSC by aggregating five questions related to the financial and business conditions of the households, as well as short and long-term perspectives on the economy. Individual sentiments are then pooled together to create the Michigan Consumer Sentiment (SENT), which serves as a barometer for the whole economy.[5]

3 Methodology

We face the obstacle of incorporating heterogeneous beliefs into a macroeconometric model, resulting from the very high-dimensional nature of survey response data. We consider the large cross-sections of the expectations as originating from continuous distributions, whose functional form is unknown, and left unrestricted.

The first problem is related to the nature of the probability density function (pdf). Since pdfs do not constitute a vector space, being nonnegative and with a constrained integral, the application of typical linear functional data methods is inadequate. However, the density is just one of many characterizations of a distribution, like the cumulative density function (cdf) or the quantile function (qf). Among these, the qf is the least constrained, being restricted just to have a non-negative derivative. Moreover, an appealing feature of quantiles is the enhanced interpretation. It can be shown that the τ-th quantile of inflation expectations represents the median consensus under an asymmetric loss, where higher (if $\tau > 0.5$) or lower (if $\tau < 0.5$) values of expected inflation are weighed more. The no-asymmetry case coincides with the median, which corresponds to the popular consensus, routinely used in non-distributional applications. The researcher interested in the evolution of the "consensus" of agents with a fear of excessive inflation, would naturally weigh high inflation expectations more than the deflationary ones. This calculation will bring the researcher to some $\tau > 0.5$ quantile of the inflation expectation distribution. If one works with pdf and cdf this property of the data is lost, and the interpretation requires restoring the qf (or some moments of interest) indirectly. Moreover, to be modeled in a regression setting, they need to be transformed to span the real line, thus further losing the simple interpretation related to movements in the quantiles.

Quantile approximation of a distributional dataset is not exempt from drawbacks. First, the infinite-dimensional quantile function cannot directly enter a macroeconometric model. An approximation is required, which amounts to selecting a finite grid $\tau_1, .., \tau_{N_q}$, leading to a vector-valued time series of quantiles. The higher N_q, the higher the precision. A rich quantiles specification is needed to potentially capture any cross-sectional heterogeneity present in the data,

[3] We start in 1983 to avoid the potential effects on expectations of the change in the monetary policy regime during the Volcker disinflation period. The data can be downloaded from the St. Louis FRED database with the following ID: WTISPLC, INDPRO, CPIAUCSL, FEDFUNDS, UMCSENT.

[4] The questions refer to "prices in general" or "inflation", without specifying a particular measure. The survey is a short rotating panel with average number of respondent of 566 (min: 480, max: 1459). A summary of the main features of the survey data is available in Appendix C.

[5] A detailed description of the index construction is available at https://data.sca.isr.umich.edu/fetchdoc.php?docid=24770 .

and compete with sophisticated functional approximation methods.[6] Second, the structural quantiles delivered by the chosen model might cross, violating the monotonicity constraint. This feature is highly undesirable and generates a daunting trade-off. Third, even if estimated quantiles don't cross, they can't be used directly to recover proper densities. The initial approximation to compress the infinite-dimensional objects yields a discrete object, impossible to map with the original quantile function. In this paper, we tackle this problem by smoothing the estimated quantiles with a flexible kernel distribution, in the spirit of Galvez and Mencía (2014) and Galán (2020).[7] By construction, the smoothed quantiles are monotonically non-decreasing, and they belong to a known functional form that we can use to restore a pdf. The kernel distribution is appealing because it encompasses also multimodality, in addition to asymmetry and excess of kurtosis. The enhanced flexibility comes with the cost of selecting a bandwidth, which regulates the smoothness of the pdf. We select it in a data-driven fashion, as explained in details below, in the paragraph about kernel interpolation.

Temporal Discretization. Inflation expectations from the MCS are categorical variables, taking values in the set of integer numbers. We believe that there is an unobserved degree of uncertainty around the answer consumers give. For instance, when they answer 10, they actually might be thinking about something around 10. In order to introduce a degree of uncertainty, we need to smooth the dataset with a continuous distribution. Without loss of generality, we assume the uncertainty around a given number takes a Gaussian form, with a fixed variance. This implies that the optimal smoothing is the one obtained with a Gaussian kernel distribution.[8] A graphical representation of the time series of smoothed distributions is given in Figure 12 in Appendix D. From the estimated $f_1, ..., f_T$, we obtain the associated Kernel-based quantile functions $q_1, ..., q_T$. Then, for each t, we form the vector-valued time series of quantiles $\mathbf{q}_t$ by extracting from each q_t $N_q = 13$ values associated with the probabilities $\tau_1 = 2.5\%$, $\tau_{N_q} = 97.5\%$, and $\tau_i = \tau_{i-1} + \frac{\tau_{N_q} - \tau_1}{N_q - 1}$ for $i = 2, ..., N_q$.[9] These starting and end points ensure we trace out all the distribution, leaving behind just the extreme 5% of the observations. Morover, this sequence ensures that all the quantiles are equidistant, and the median always corresponds to $\tau_{(N_q+1)/2}$. The time series of smoothed quantiles are depicted in Panel (a) of Figure 1.[10]

The Time Series Model. After discretization, all the objects of interest are vector-valued,

[6] Chang et al. (2023) compare a low dimensional quantiles specification with a richer functional approximation, and show that the resulting dynamic analysis is comparable in median. However, the quantiles specification features higher uncertainty, making the confidence bands wider, thus challenging the interpretation. Using a simple local projection argument, they show that a poor quantile approximation is misspecified, with the misspecification vanishing as the number of quantiles increases.

[7] This problem is not new to the econometric literature. In fact, it arises also in the context of quantile regression, where the estimated conditional quantiles are noisy and potentially cross. Adrian et al. (2019, 2022) tackled this issue in a quantile regression context with a Skewed-t smoothing. The Skewed-t generalizes the Normal distribution allowing for asymmetry and excess of kurtosis, but not multimodality.

[8] Without any information about the degree of uncertainty, we select the bandwidth using Silverman's rule of thumb, which is equal to $1.06 \times \hat{\sigma}_t N_t^{-1/5}$, where σ_t and N_t are, respectively, the sample standard deviation of individual inflation expectations and the number of individuals at time $t = 1, ..., T$.

[9] We select 13 quantiles because this number guarantees that most of the stylized facts exhibited by the data are captured. It is in line with Chang et al. (2023), who select 10 basis functions to approximate the temporal distribution of US earnings.

[10] Given what said above, in Figure 1 the quantiles are defined by the τ_i's [2.5, 10.4, 18.3, 26.2, 34.1, 42, 50, 57.9, 65.8, 73.7, 81.6, 89.5, 97.5] and the corresponding means (see Panel (b)) are [-2.51, 0.63, 1.49, 2.1, 2.60, 3.08, 3.57, 4.11, 4.69, 5.4 6.64, 9.24, 16.12] .

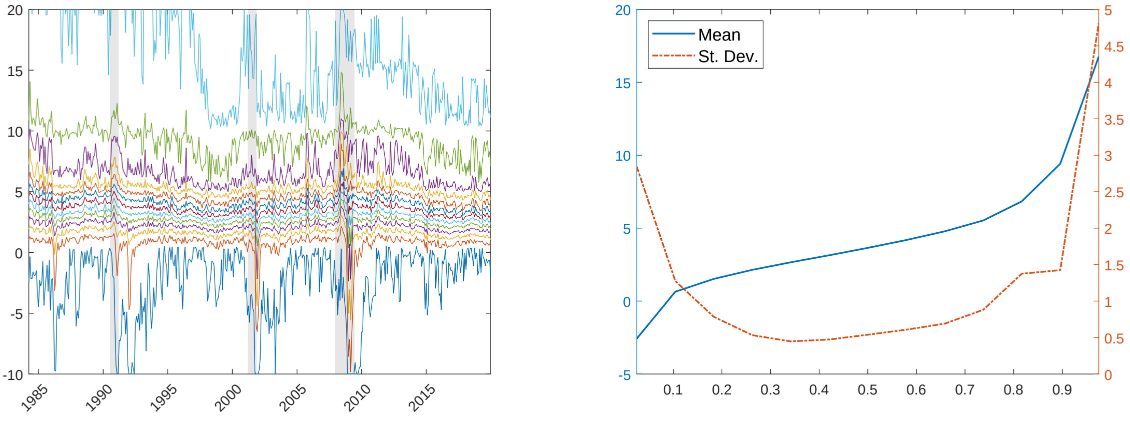

(a) (b)

Figure 1: Panel (a): Time series of the 13 empirical quantiles extracted from the MSC inflation expectation temporal distribution. Monthly data from January 1983 to December 2019. Panel (b): Empirical mean (solid blue line) and standard deviation (dashed-dotted orange) of the quantiles pictured in Panel (a).

and thus suitable to be jointly modeled. The vector of quantiles $\mathbf{q}_t$ and the index of consumer sentiment form the expectation block $\mathbf{e}_t$, of dimension $N_e \times 1$, where $N_e = N_q + 1$. All the remaining variables - i.e., OP, CPI, IP, and IR - form the macroeconomic block $\mathbf{y}_t$ of dimension $N_y \times 1$. We form the $N = N_e + N_y$ dimensional vector of endogenous variables $\mathbf{z}_t$ by vertically stacking the two vectors for each t: $\mathbf{z}_t = [\mathbf{e}_t', ..., \mathbf{y}_t']'$.

Let $\boldsymbol{\mu}_t$ be the long run component of $\mathbf{z}_t$, following the assumption that the deviations $(\mathbf{z}_t - \boldsymbol{\mu}_t)$ have stationary dynamics and null unconditional expectations. We model these deviations as a stable, structural VAR:

$$\boldsymbol{\Phi}(L)(\mathbf{z}_t - \boldsymbol{\mu}_t) = \boldsymbol{\epsilon}_t, \quad \boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}) \tag{1}$$

where $\boldsymbol{\Phi}(L) = \mathbf{I}_N - \boldsymbol{\Phi}_1 L - ... - \boldsymbol{\Phi}_P L^P$ is the usual polynomial in the lag operator L, and $\boldsymbol{\Phi} = [\boldsymbol{\Phi}_1, ..., \boldsymbol{\Phi}_P]$ are reduced form lagged parameters. Finally, $\boldsymbol{\epsilon}_t$ is the N dimensional innovation process, that is $\mathrm{E}(\boldsymbol{\epsilon}_t) = \mathbf{0}$, $\mathrm{E}(\boldsymbol{\epsilon}_t \boldsymbol{\epsilon}_t') = \boldsymbol{\Sigma}$, and $\mathrm{E}(\boldsymbol{\epsilon}_t \boldsymbol{\epsilon}_s') = \mathbf{0}$ for $s \neq t$, where $\boldsymbol{\Sigma}$ is positive definite. Our model is equivalent to the VAR in deviations from its steady states of Villani (2009), with the important difference that in this specification some of the steady states are allowed to change deterministically over time. This feature allows us to work with both stationary and non-stationary variables, and to distinguish between long and short dynamics through the trend-cycle decomposition.

The long run dynamics. We model the long-run component of the observable with regression splines. Regression splines are a subclass of piecewise polynomials possibly featuring continuity and smoothness, depending on the polynomial degree. Suppose to partition the time domain in $K + 1$ sections, delimited by a set of knots $\xi_1, ..., \xi_K$, where $\xi_1 \geq 1$, $\xi_K \leq T$, and $\xi_k > \xi_{k-1}$, for $k = 2, .., K$. An R-th order regression spline with knots ξ_k, is a piecewise polynomial of order R with continuous derivatives up to order $R - 2$. Let $\mu_{i,t}$ be the generic element of $\boldsymbol{\mu}_t$, we model it as:

$$\mu_{i,t} = \sum_{r=0}^{R_i} \theta_{i,r} t^r + \sum_{k=1}^{K_i} \theta_{i,R_i+k} (t - \xi_k)_+^R \tag{2}$$

where $\boldsymbol{\theta}_i = [\theta_{i,0}, ..., \theta_{i,R_i+K_i}]'$ is the $1 + R_i + K_i \times 1$ vector of free parameters to estimate for variable i. The expression in parentheses on the right-hand side is a short-hand notation for the following:

$$(t - \xi_k)_+^r = (t - \xi_k)^r \mathbb{1}(t - \xi_k) = \begin{cases} 0 & t \leq \xi_k \\ (t - \xi_k)^r & t > \xi_k \end{cases}$$

The benefit of this specification is that it preserves high parsimony (complex temporal dynamics can be captured with relatively low values of $R_i + K_i \geq T$) and linearity in the parameters, so that the standard estimation techniques can be deployed. By construction, splines are continuous for $R_i \geq 1$, and smooth for $R_i \geq 2$ across the whole domain. When $K_i = 0$, equation (2) boils down to a standard global polynomial of order R_i. When $K_i = R_i = 0$, $\mu_{i,t}$ reduces to a constant. In general, by tailoring R_i, K_i, and the position of the knots, the researcher controls for the desired level of fit. It is interesting to note that when $R_i = 3$, $\mu_{i,t}$ corresponds to a cubic spline, which is particularly appealing because it is the lowest-level polynomial for which the knot discontinuity is not visible to the human eye. As a consequence, the task of choosing the locations of the knots is minimized. In Appendix A we show how to represent compactly the long run component as $\boldsymbol{\mu}_t = \mathbf{H}_t \boldsymbol{\theta}$, where the $q = N + \sum_i R_i + K_i$ dimensional vector $\boldsymbol{\theta} = [\boldsymbol{\theta}_1', ..., \boldsymbol{\theta}_N']'$ vertically stacks all the free coefficients to estimate.

The expectation variables are I(0), and thus we assume $R_i = K_i = 0$ for $i = 1, ..., N_e$. On the other hand, all the macroeconomic variables exhibit non-stationary behaviors, and thus we allow for time-varying trends. In particular, we employ a cubic spline, with $K_i = 3$ knots placed at equally spaced locations in the time domain.[11] With this convenient representation, the model can be estimated by straightforward modification of the Gibbs sampler developed by Villani (2009).

Priors and hyperpriors. Bayesian inference requires a prior distribution on the reduced form parameters $\boldsymbol{\theta}$, $\boldsymbol{\phi} = \text{vec}(\boldsymbol{\Phi})$, and $\boldsymbol{\Sigma}$.

We reflect our belief that each variable is trend stationary, i.e. $\text{E}(\mathbf{z}_{i,t}) = \underline{\boldsymbol{\mu}}_{i,t}$, where $\underline{\boldsymbol{\mu}}_{i,t}$ represents the prior trend, as a function of the prior values of $\boldsymbol{\theta}_i \sim \mathcal{N}(\boldsymbol{\theta}_{i,0}, \boldsymbol{\Xi}_{i,0})$. To rule out possible posterior unstable dynamics, we center each $\boldsymbol{\theta}_{i,0}$ on the values resulting from the univariate spline fit for each $\mathbf{z}_{i,t}$. The prior variance is given by $\boldsymbol{\Xi}_{i,0} = \lambda_\theta \sigma_i \mathbf{I}_{1+R_i+K_i}$, where σ_i are the residuals from the univariate fit, and $\lambda_\theta > 1$ is a slab hyperparameter controlling the uncertainty around the prior mean. This procedure can be viewed as a generalization of pre-processing the variables by spline de-trending, obtained as the limit $\lambda_\theta \to 0$.

The VAR coefficients are a priori independent across equations, and Normally distributed:

[11] The selected specification ensures rather stable dynamics: in fact, we find that less than 40% of the draws are nonstationary, implying a maximum absolute eigenvalue of the resulting companion form above 1. Moreover, we point out that most of the instability is due to the Interest Rate. In an auxiliary specification without this variable, most of the draws ensure stable dynamics.

	OP cost push	Inflation expectations	Fundamental Macro
OP	+	0	0
Expectation block	/	+	0
Other Macro	/	/	+

Notes. Entries in this table show the restrictions imposed: + positive sign, 0 no effect, / no restriction. Rows represent the variables and column the shocks. Inflation expectation shocks and other fundamental macroeconomic shocs are block identified. The OP is affected on impact (positively) only by its own shock. The macro block shock does not affect on impact any of the other variables.

Table 1: Identification assumptions.

$\boldsymbol{\phi} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega})$. Let $\boldsymbol{\Phi}_{p,ij}$ be the coefficient reflecting the p-th lagged effect of variable j on equation i. We set the prior variance as:

$$\mathrm{V}(\boldsymbol{\Phi}_{p,ij}) = \begin{cases} \frac{\lambda_1 \sigma_i}{p^{\lambda_p} \sigma_j} & i = j \\ \frac{\lambda_1 \lambda_2 \sigma_i}{p^{\lambda_p} \sigma_j} & i \neq j \end{cases}$$

where σ_i and σ_j are defined above, λ_1, λ_2 and λ_p are hyperparameters. The first two are determined a priori, ie $\lambda_2 = 0.5$ and $\lambda_p = 1$, and the latter is estimated, given the large dimension of the system. In particular, we borrow the idea from Chan (2021) and assume that λ_1 is a priori Gamma distributed with shape $\gamma_1 = 1.02$ and scale $\gamma_2 = 10$. These values imply the prior mode of λ_1 is 0.02, with a rather slab variance. The prior setting concludes by assuming the usual Inverse-Wishart for $\boldsymbol{\Sigma}$:

$$\boldsymbol{\Sigma} \sim \mathcal{IW}(\boldsymbol{\Psi}, d_0)$$

where $\boldsymbol{\Psi}$ and d_0 are, respectively, the prior scale matrix and prior degrees of freedom. The posterior distribution of the model in equation (1) is intractable, and so direct sampling is unfeasible. In appendix A we outline the steps of the MCMC sampler used to draw from the posterior distribution of the parameters of interest.

Identification. We impose a set of minimal identification assumptions based on the natural discrepancies between the date the survey is conducted, and the date in which macroeconomic variables are released. The Michigan survey usually takes place between the third and the fourth week of each month, whereas all the macroeconomic variables of the current month are released weeks later. Hence, when consumers form their expectations, and answer the survey questions, they don't know which will be the value of the aggregate variables. As a consequence, they can't be influenced by them during the current month. This feature of the data allows us to impose that all the macroeconomic variables have a zero restricted impact on the expectation equations. However, this assumption does not hold for the oil price, which we treat as the most exogenous variable in the system. It is reasonable to assume that oil prices might instantaneously impact inflation expectations and sentiment because consumers are likely to observe the petrol price -

highly correlated with oil prices - at the petrol station while filling their car tanks and because the literature suggests that the price at the pump is a very salient price for consumers. The identification assumptions are illustrated in Table 1.

Kernel interpolation. The mapping between the infinite-dimensional densities to the vector of quantiles enables us to leverage standard time series econometric methods such as the VAR. This drastically simplifies the issue of modeling functional and vector-valued objects together. Nevertheless, the map is not bijective, and from the discrete quantiles obtained from the VAR, we cannot restore proper pdfs. Moreover, although the choice of using a large number of quantiles reduces the misspecification, the higher flexibility comes at the cost of potential quantile crossing. This clearly violates the monotonicity assumption underlying the quantile function. The problem cannot be solved unless inequality constraints in the parameter's space are introduced, which would dramatically hamper the estimation of the model.

Here, we rely on a post-estimation procedure based on a weighted kernel interpolation, as in Galvez and Mencía (2014); Galán (2020). A similar approach was used by Adrian et al. (2019), who invoked a Skewed-t distribution. A parametric fitting, although practical, implies strong assumptions on the functional form. In particular, the expectation distributions face also multimodality, a feature not supported by the Skewed-t. The kernel distribution is an appealing nonparametric alternative, able to fit any type of cross-sectional dispersion featured by the data. The degree of flexibility depends on the bandwidth $\mathcal{B}$, that we choose in a data-driven fashion.

Suppose we are interested in smoothing the steady state of the quantiles obtained from the BVAR. At the m-th iteration of the sampler, a valid draw from the noisy posterior distribution of the quantiles steady states is available: $[\mu_1^{(m)}, ..., \mu_{N_q}^{(m)}]$, where the superscript $m = 1, ..., M$ indicate the m-th draws. We adopt the following strategy: we smooth this noisy vector with a weighted Kernel distribution, with predetermined bandwidth $\mathcal{B}$, and estimated weights. The associated smoothed quantiles steady state cdf at iteration m is given by:

$$\mathrm{F}_{ss}^{(m)}(x|\mathcal{B}) = \sum_{i=1}^{N_q} \hat{w}_i \Phi\left(\frac{x - \mu_i^{(m)}}{\mathcal{B}}\right) \tag{3}$$

where $\Phi(\cdot)$ is the standard Gaussian cdf. The vector of weights $\hat{\mathbf{w}} = [\hat{w}_1, ..., \hat{w}_{N_q}]'$ are obtained by minimizing the squared distance between the quantile level and the associated cdf:

$$\hat{\mathbf{w}}^{(m)} = \arg\min_{\mathbf{w}} \quad \sum_{i=1}^{N_q} \left[\tau_i - \mathrm{F}_{ss}^{(m)}\left(\mu_i^{(m)}|\mathcal{B}\right)\right]^2, \quad \text{s.t.} \quad \mathbf{w} \geq \mathbf{0}, \quad \boldsymbol{\iota}'\mathbf{w} = 1 \tag{4}$$

Equation (4) identifies a quadratic programming problem, which can be easily solved with standard software. In our setting, the small dimension of N_q ensures a negligible computational cost, such that the estimation time of the model with or without smoothing is not distinguishable. The linear equality and inequality constraints guarantee that the solution to this problem exists.[12] The bandwidth $\mathcal{B}$ is crucial, and determines the jaggedness of the distributional output. We select it in a data-driven fashion to match relevant features of the data. In particular,

[12]In fact, (4) identifies a framework in which N_q explanatory variables are used to estimate N_q parameters. The constraints alleviate the problem of inverting a near singular moment matrix.

we select it as the value that minimizes the empirical Wasserstein distance of two probability measures: the one given by the posterior median of the smoothed steady state cdf, and the one implied by smoothing the pooled dataset, following the procedure in the temporal discretization subsection.

Here we have outlined the procedure considering the steady state, given that it provides a simple and intuitive example. However, we use this procedure any time a quantile-based object is studied. Thus, also the dynamic analysis results we report are smoothed using this procedure.

Density conditional forecasts. In our experiments, we investigate the macroeconomic effects of shocks to the inflation expectation distribution, summarized by the N_q quantiles. Although they enter the VAR separately, quantiles represent a single object, because we want to see the effects of perturbing the whole distribution of inflation expectations. Of course, there are many ways one can perturb a distribution. For example, think about the following question: What happens to the macroeconomic aggregates if an exogenous shock in the consumer's expectations hits the economy, and the perturbation is such that the new distributions' location moves rightward of 1%? That is, the shock produces an increase in inflation expectation of 1%, without affecting any higher order moment of the distribution. To analyze this, as well as any other scenario that involves a movement in the distribution, we need two ingredients. The first one is simply the desired movement of the distribution which implies a movement in all the N_q quantiles, given the one-to-one match between pdf and quantile function. Our example of a positive 1% variation in the location of the distribution implies a positive, 1% movement of all the quantiles.

The second ingredient is the composition of the structural shocks. Remember that given our identification restrictions, the shocks to expectations are block identified. These correspond, at each time t, to N_e reduced form shocks: the N_q shocks to the quantile equations plus the shock in the equation of sentiment. The latter is also part of the expectation block, as it is part of the MSC, and thus it directly affects consumer's expectations reflecting the consumer's perceived stance of the economy. Thus, our structural shock to the expectations is a linear combination of these N_e reduced form shocks. Since we want to find a linear combination of N_e shocks that satisfy $N_q < N_e$ restrictions, there is an infinite number of solutions. These features pave the way for the structural scenario analysis in the spirit of Antolín-Díaz et al. (2021): we find the posterior distribution of the N_e shocks realizations that satisfy the N_q restrictions, imposing all the other non-expectation shocks to be zero. Intuitively, the infinite solutions of our problem are subject to a probability distribution that reflect the correlations in the data. This procedure, which is explained in details in Appendix B, is very general and allows us to experiment with a variety of economically meaningful shocks to the distribution of inflation expectations. We can analyse variations in the common statistics like the various moments, and, more interesting, we can study the effects of exogenous marginal movements of specific sections of the distribution.

4 Results

Steady states. Figure 2 displays the estimated steady-state distribution of inflation expectations. The wide range of the answers (roughly from -5 to 20 percent) and the multimodality

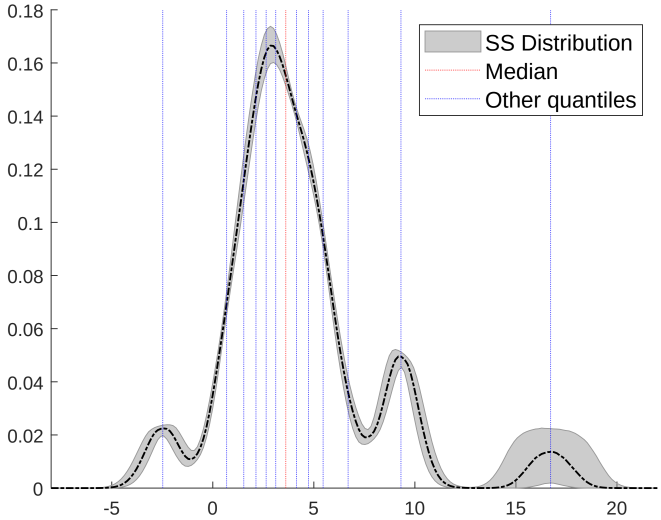

Figure 2: Kernel smoothed posterior SS distribution: median (dashed-dotted black line), and 68% credible interval (grey area). The red and blue lines represent the median posterior of the quantiles (see footnote 10).

of the distribution highlight the importance of the heterogeneity of beliefs. The median of the distribution is about 3.6%, definitely higher than the central bank inflation target. The long-run component of the macroeconomic variables, along with the associated gaps, are pictured in Figure 13 in Appendix D.

The remainder of this Section presents the main results. Section 4.1 studies the cross-influence between the expectation and to the macro blocks of the model providing indirect empirical support to our identification strategy. Section 4.2 looks at the relative importance of shocks to the two blocks for the dynamics of the variables in the BVAR. Section 4.3 contains the main results of the paper, that is, the analysis, through density impulse response functions, of the effects on the macroeconomic variables of different kind of shocks to the inflation expectation distribution: mean, variance, kurtosis and tail shocks. Finally, Section 4.4 investigates the effects of communication shocks.

4.1 Cross-influence between blocks of the model

Before moving to the main results, here we want to further investigate the cross-influences between the two blocks of variables by focusing on the lagged coefficients. In particular, we are interested in empirically assessing to what extent the macroeconomic variables influence expectations and vice versa. In a frequentist setting, one possible way to investigate this question is to conduct a Granger-causality test (F-test) on blocks of parameters. Here we take a more probabilistic approach, by exploiting the hierarchical structure of our model. In particular, we relax the assumption of a common λ_1, by assigning different shrinkages to parameters belonging to two different portions of the system: the ones related to the lagged effect of the macroeconomic variables on the expectation equations ($e \leftarrow m$), and the ones related to the lagged effect of the expectation variables on the macro ($m \leftarrow e$). To fix ideas, consider that each $\mathbf{\Phi}_i$ can be

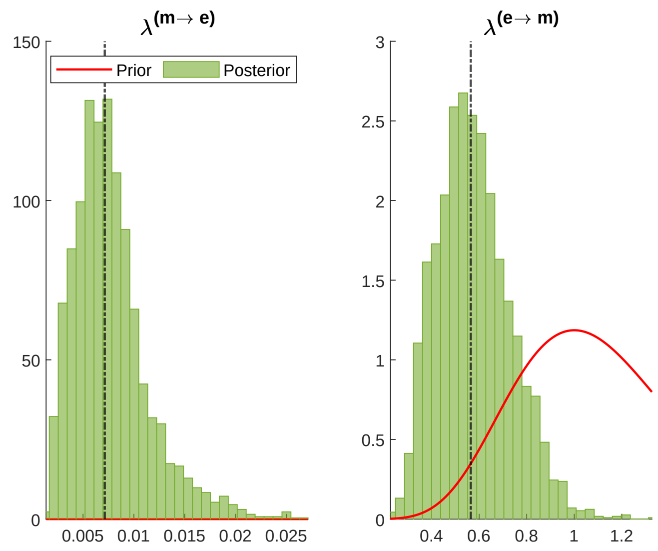

Figure 3: Expectations are important predictors of macroeconomic variables and not vice-versa. Prior and posterior distribution of the two shrinkage parameters: $\lambda^{m\to e}$ (macro on expectations) and $\lambda^{e\to m}$ (expectations on macro).

partitioned as follows

$$\boldsymbol{\Phi}_i = \begin{bmatrix} \boldsymbol{\Phi}_i^{(e\leftarrow e)} & \boldsymbol{\Phi}_i^{(e\leftarrow m)} \\ \boldsymbol{\Phi}_i^{(m\leftarrow m)} & \boldsymbol{\Phi}_i^{(m\leftarrow e)} \end{bmatrix}$$

where $\boldsymbol{\Phi}_i^{(e\leftarrow e)}$ and $\boldsymbol{\Phi}_i^{(m\leftarrow m)}$ are, respectively, the $[N_e \times N_e]$ and $[N_y \times N_y]$ submatrices embedding expectation to expectation and macro to macro lagged coefficient. The interesting blocks for us are the off-diagonal ones, $\boldsymbol{\Phi}_i^{(e\leftarrow m)}$ and $\boldsymbol{\Phi}_i^{(m\leftarrow e)}$, that incorporate cross-influences between the two blocks. Remarkably, the two blocks share the same number of parameters $N_e N_y$, and so can be compared. Let $\mathrm{C}_{p,ij} = \lambda_1 \lambda_2^{-p} \sigma_i \sigma_j^{-1}$ be the constant part of the Minnesota prior variance for each p, i, j, as defined above.[13] We now define a new prior variance or the two off-diagonal sub-matrices:

$$\mathrm{V}(\boldsymbol{\Phi}_{p,ij}^{(m\leftarrow e)}) = \lambda_{m\leftarrow e} \mathrm{C}_{p,ij}, \quad \mathrm{V}(\boldsymbol{\Phi}_{p,ij}^{(e\leftarrow m)}) = \lambda_{e\leftarrow m} \mathrm{C}_{p,ij}$$

both $\lambda_{m\leftarrow e}$, and $\lambda_{e\leftarrow m}$ are equipped with a Gamma prior with shape and rate equal to 11 and 10, respectively. This ensures the prior mode is equal to 1, with a rather slab variance, to let the data speak. The differences in the magnitude of their posterior distributions inform us about the relative predictive power of the lagged values of the two respective blocks of variables. If the posterior of $\lambda_{m\leftarrow e}$ and $\lambda_{e\leftarrow m}$ share similar magnitudes we can conclude that the lagged dependence of the two blocks of variables is comparable. If there is a statistically significant discrepancy, a ranking emerges, suggesting that one of the two classes is useful as a lagged predictor for the other, but not vice versa.

Figure 3 compares the prior and the posterior distributions of the two shrinkage parameters. What clearly emerges is that the posterior of $\lambda^{m\leftarrow e}$ is roughly 100 times higher than the one of

[13] In what follows, we assume λ_1 is known, and is the output of a first-stage estimation, from which we get its posterior mode. With λ_1 in hand, we re-estimate the model with the two shrinkages.

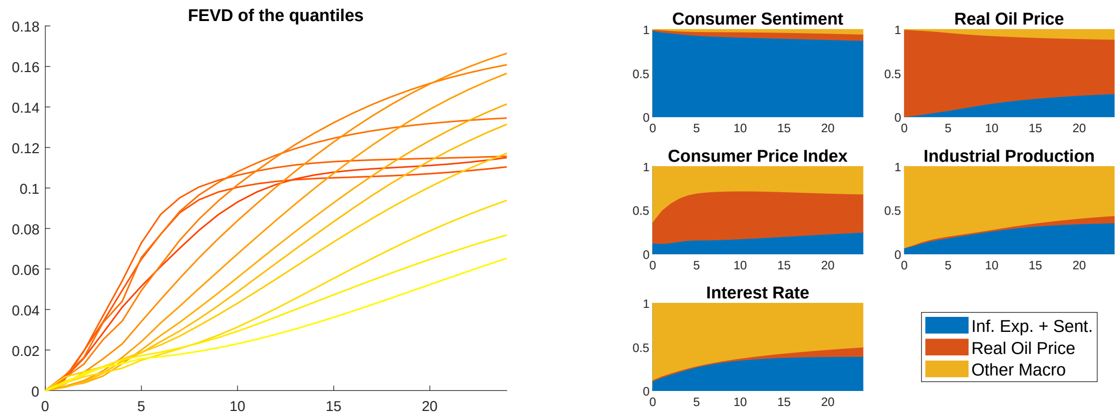

(a) (b)

Figure 4: Panel (a): Contribution of the three macroeconomic variables to the FEVD of the 13 Quantiles at different horizons (x-axis). Different colors indicate the location of the quantiles in the distribution. Starting from the left tail (red), to the right tail (yellow). Panel (b): FEVDs of the macroeconomic variables at different horizons (x-axis).

$\lambda^{e \leftarrow m}$. When the two sub-portions of the system are allowed to have different shrinkages, the model convincingly shrinks towards 0 all the macro to expectations parameters. Independently on the identification assumptions, macroeconomic variables are not relevant predictors for the expectation ones.

This exercise is similar, in spirit, to what a Granger causality test would suggest in a frequentist setting. The evidence strongly supports the recursive identification scheme we described above, originally based on the timing of the information set. Let us stress that these results are only used to corroborate the zero restrictions on the impact matrix: the analysis we describe below is based on a model in which we assume a common shrinkage parameter.

4.2 The important role of exogenous variations in inflation expectations

Expectations are barely affected by macroeconomic shocks. In Figure 4 we report the forecast error variance decomposition (FEVD) of the two blocks of variables of interest. In Panel (a) we depict the quota of FEVD – at different horizons on the x-axis – for each of the 13 quantiles explained by the macroeconomic variables, except for the oil price. The quantiles are depicted with different nuances of red and yellow to understand the location of a given quantile. In particular, red lines represent the quantiles in the right tail, whereas the yellow lines represent the left tail ones. The lines start at zero given the assumption that macroeconomic variables don't affect expectations on impact. At a first glance, what stands out is the relatively low magnitude of the decompositions, reflecting a quasi-exogeneity of the expectations with respect to the macroeconomic aggregates. This is a very surprising result given the fundamental tenant in macroeconomic models that inflation expectations are fully endogenous to the macroeconomic developments. Instead, our result suggest that consumers in the MSC give little attention to the standard macroeconomic variables in forming their inflation expectations. This is somewhat consistent with the recent large literature on survey expectations that showed that inflation expectations are far from rational, subjects to various biases and affected more by personal experiences or few salient prices, rather than by macroeconomic aggregates.[14]

An additional interesting feature is the ranking that emerges by analyzing the single quantiles FEVD. Although the responsiveness to macroeconomic fluctuations is generally low, it differs across quantiles. In particular, the quantiles of the right tail (yellow lines) are the least affected, with values below 10%. The frame is different for the left tail. Here the FEVD of the macroeconomic variables rise in the first horizons, and then stabilize on values slightly higher than the ones on the right tail. The more central quantiles (orange lines) feature the higher macro-based FEVD. This pattern suggests that the expectation formation process of consumers might be heterogeneous. In particular, consumers with a deflationary and low inflationary view of the economy might be influenced by business cycle movements, although partially. It makes also intuitive sense that the central quantiles, which represents the least biased inflation expectations in some sense, are the most responsive to macroeconomic changes. In contrast, consumers who fear high or extreme inflation are almost entirely exogenous to macroeconomic fluctuations.

Exogenous variations in expectations affect macroeconomic variables. Panel (b) of Figure 4 exhibits an opposite scenario. Here, we show the complete FEVD for the macroeconomic variables and the sentiment. We partition the FEVD in three main contributions to streamline the interpretation: the expectations block (blue area), the real oil price contribution (orange area), and the macroeconomic block (yellow area). Notably, there are two main results: exogenous shocks to inflation expectations are an important determinant of variations in macroeconomic variables. Their relevance increases with the horizon. The respective share of the variance of the two year forecast error is roughly 24% percent for CPI, 35% percent for Industrial Production and 36% percent for the interest rate.

This result suggests that the importance of inflation expectations goes beyond their role in the propagation mechanism of macroeconomic shocks: there is an exogenous source of variation of expectations that affects directly macroeconomic variables. Next, we move to analyse the dynamic implications of such shocks, and we will stress the importance of considering the whole distribution of inflation expectations.

4.3 Dynamic implications of shocks to inflation expectations: the distribution matters

This Section contains the main findings of this paper. We perturb the steady-state distribution of inflation expectations, shown in Panel (a) of Figure 2, using expectation shocks and study the dynamic implications on all the variables in the model. We conduct several experiments analysing the marginal effects of a change in the mean, in the dispersion and in some selected portions of the distribution.

[14] In Panel (a) we show the contribution to the FEVD of shocks to the macroeconomic block, excluding oil prices, to highlight how consumers barely take into account this information when forming inflation expectations. Figure 14 in Appendix D shows the FEVD also including oil price shocks: the latter are an important source of variation, but the exogenous shocks to expectations remain by far the most important determinant of consumers' beliefs.

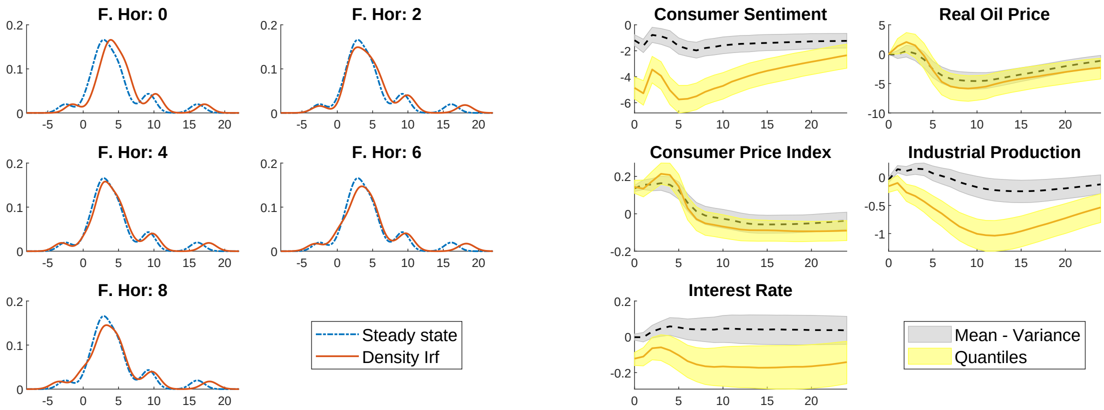

(a) (b)

Figure 5: Panel (a): Smoothed DIRFs of the positive mean shock for some selected forecast horizons $0, 2, .., 8$. SS distribution (dashed-dotted blue lines), DIRFs (solid orange lines). Panel (b): IRFs of macroeconomic variables for the positive mean shock: dashed black lines (mean-variance specification), solid yellow lines (quantile specification). Shaded areas delimit the 50% credible bands (this bandwidth will be also used for all the subsequent experiments).

4.3.1 Mean shock

The first experiment we run is the distributional counterpart of the IRFs one would study if the model contains a synthetic measure of expectations, like the mean. We ask the following question: what is the macroeconomic effect of a shock that increases inflation expectations of 1%? Recall that any $\alpha\%$ increase in the location of the distribution corresponds to an $\alpha\%$ increase in all the quantiles. Having said that, here, we generate a scenario where all the quantiles increase by the same amount, and the distribution of inflation expectations translates to the right keeping its shape unchanged. Panel (a) of Figure 5 shows the smoothed distributional impulse response functions (DIRFs) for some selected forecast horizons: after the initial shift to the right, the distribution slowly converges back to its steady state. The effects on the other macroeconomic variables are shown by the yellow line in Panel (b). An exogenous increase in the mean of inflation expectations is stagflationary, inducing a rise in inflation and a decrease in industrial production. This result is consistent with Ascari et al. (2023) who obtain similar evidence in BVAR on aggregate variables, using a different dataset and identification assumptions. Moreover, the strong negative response of sentiment suggests that the effects on the economy work through the association of higher inflation with a pessimistic perception about the economy. This is in line with robust evidence from surveys pointing to households consistently having a supply-side view of inflation, such that they become more pessimistic about the economic outlook when their inflation expectations increases (see, e.g., Kamdar, 2019; Candia et al., 2020; Binder, 2020; Binder and Makridis, 2022). Moreover, Coibion et al. (2023) exploit a randomized controlled trial approach on Dutch households data to provide causal evidence from inflation expectations to actual spending. Households who have exogenously higher inflation expectations become much more pessimistic about the state of the economy and their future

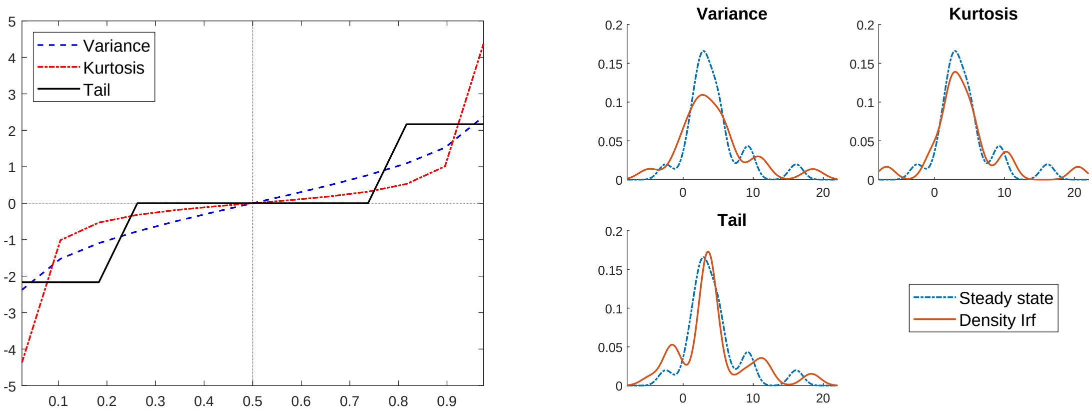

(a) (b)

Figure 6: Panel (a): Three examples of shocks to the quantiles inducing an increase in dispersion: A variance shock (dashed blue line), a kurtosis shock (dashed-dotted red line), and a tail shock (solid black line). Panel (b): Corresponding perturbed distributions (solid orange lines), and SS distributions (dashed-dotted blue lines), reflecting the three shocks depicted in Panel (a).

income growth, and they sharply reduce their spending on durable goods.[15]

In Panel (b) we also compare our results with the IRFs obtained through a VAR in which, instead of the 13 quantiles, we include the mean and the variance of the expectation distribution (black dashed lines). Both the parsimonious and the rich specification agree on the fact that an exogenous increase in inflation triggers a stagflationary effect. However, while the two models exhibit similar qualitative responses, there is a statistically significant and marked difference in the quantitative responses of sentiment and industrial production. Specifically, the parsimonious specification dramatically underestimates the magnitude of the recession induced by the shock. As a response to the same perturbation, the effect on industrial production in the richer model is at least five times bigger, with a long-lasting effect. The credible bands confirm the statistical significance of the result. Therefore, the distribution matters in this exercise, because not taking into account the whole distribution would lead to an underestimation of the effects of a shock to the mean of the inflation expectations. The results suggest that there might be a pass-through effect from expectations to real economy, via sentiment, which can be detected only if the entire distribution is accounted for.

4.3.2 Dispersion shocks

The mean shifting shock informs us about the effect of a hypothetical scenario in which all the households, exogenously, believe inflation will be higher within a year. Another relevant scenario is represented by an exogenous variation of the disagreement, without a change in the consensus. What are the macroeconomic effects of a shock that increases the dispersion of the distribution, that leaves the mean unchanged? These changes could also be interpreted as expectation uncertainty shocks.[16]

[15] A similar result is also find in Bachmann et al. (2015), while D'Acunto et al. (2016) find the opposite result.

[16] While we are aware that a more precise measure of uncertainty would be the standard deviation taken from the probability distribution of a single household, the MSC does not provide those. Moreover, when available in

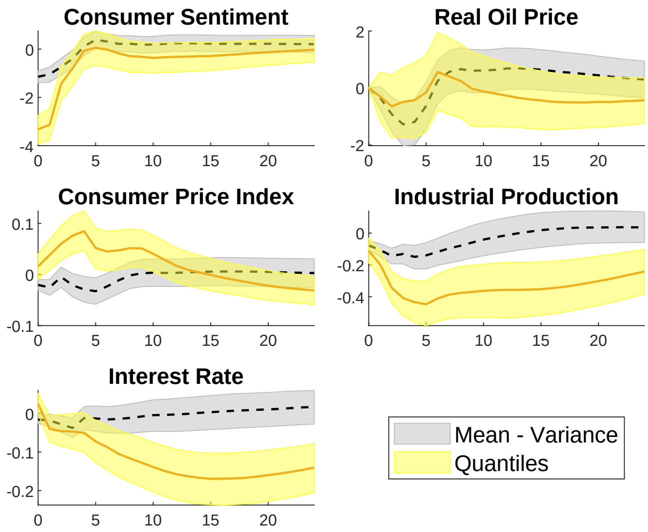

Figure 7: IRFs of macroeconomic variables to a the positive variance shock: dashed black lines (mean-variance specification), solid yellow lines (quantile specification).

Dispersion is a broad concept, involving any possible force that stretches or squeezes the distribution. In the simple Gaussian case, the dispersion is intimately connected to the variance. For more flexible parametric distributions, dispersion is a multivariate function, determined by the interplay of more parameters. In our general framework, lacking of any parametric assumption, an increase in dispersion is determined by any force that symmetrically drags the quantiles far from the median, leaving the median untouched. The movement has to be non-overlapping across the quantiles, otherwise violating the constraints featuring the definition of probability.

We chose to study the effects of three different dispersion shocks: a variance, a kurtosis, and a pure tail shock. Panel (a) of Figure 6 shows the deviations of each quantile from the steady-state distribution for the three experiments, and, for the sake of clarity, Panel (b) plots the corresponding perturbed distribution compared with the steady state. All the shocks are median-preserving and are scaled such that the overall magnitude of the quantile movements is the same across experiments. The variance shock moves probability mass more uniformly from the bulk of the distribution to the tails, while in the other two cases we assign progressively more weight to the tails while preserving the bulk. In particular, the tail shock (black line) is a simple step function that moves only the first three and the last three quantiles leaving all the others at the center unchanged. The comparison between the effects of these shocks will help us disentangling the respective role of the more peripherical and more central parts of the distribution.

We focus first on the variance shock. Figure 7 shows the effects on the macroeconomic variables (yellow lines). The variance shock is stagflationary too, increasing prices and decreasing output. Sentiment drops, a symptom that higher uncertainty is perceived with pessimism by

surveys, these data have several rounding and time-inconsistency issues making estimations of uncertainty not reliable. Hence, very often then the literature assumes that an increase in the dispersion of individual expectations across households proxies an increase in the expectation uncertainty of the households.

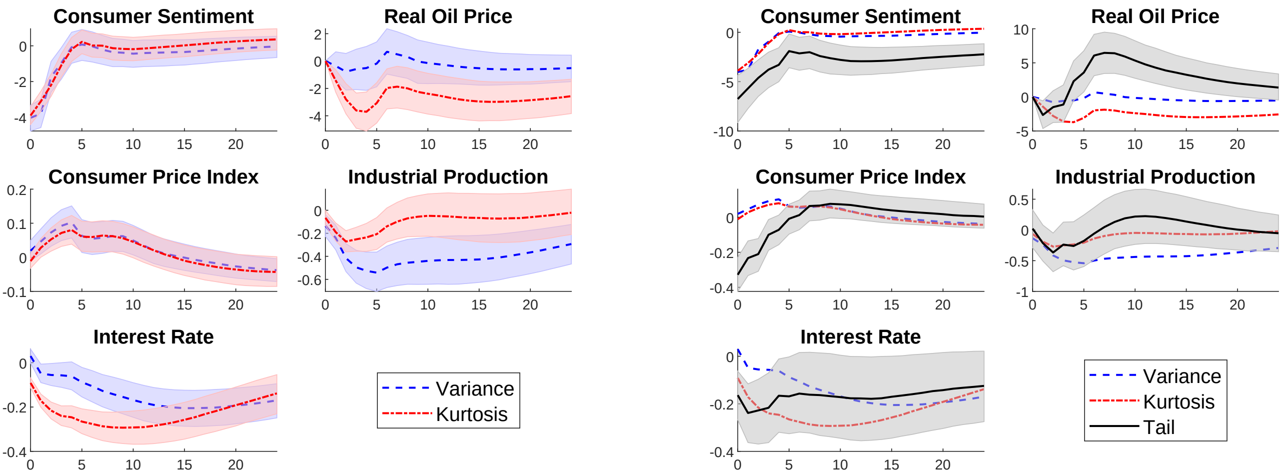

(a) (b)

Figure 8: Panel (a): IRfs of macroeconomic variables as a result of a variance (dashed blue lines) and a kurtosis (dashed-dotted red lines) shock. Panel (b): Results from Panel (a) with the inclusion of the tail shock (solid black lines). Note: Confidence bands of the variance and kurtosis shock, shown in Panel (a), are omitted in Panel (b) to help readability.

households. As for the mean shock, we can compare the effects of the variance shock with the analogous one obtained with the more parsimonious specification in which, instead of the 13 quantiles, we include only the mean and the variance (black lines). For all the variables except the Consumer Price Index, the effects have the same sign. However, as for the mean shock, the fluctuations in case of the richer model with the quantiles are much larger, leading to the same conclusion: the amplification of the responses in the specification with the quantiles highlights the important role of the whole distribution in the transmission mechanism of expectation shocks.

The second dispersion shock we consider is an exogenous increase in the kurtosis of the distribution: the effect on the tail is relatively larger than the one following a variance shock, while smaller on the bulk of the distribution. Panel (a) of Figure 8 compares the IRFs of macroeconomic variables after the kurtosis and the variance shocks. In case of kurtosis shocks, interest rate responds more and output drops less, but overall the responses are similar between the two shocks.

The differences in the IRFs become more pronounced in case of tail shock, as shown in Panel (b) of Figure 8. When the central quantiles remain at their steady state values and only the tails move outward, we have a larger drop in consumer sentiment and most notably the Consumer Price Index decreases substantially, displaying the opposite sign with respect to the other two dispersion shocks. A tail shock is recessionary. The macroeconomic effects resemble the one of a negative demand shock, with a drop in both output and inflation. This difference shows how the very same interpretation of a dispersion (uncertainty) shock, as demand or a supply shock, depends on the more precise definition of dispersion (uncertainty), and on the exact movements of the distribution of beliefs.

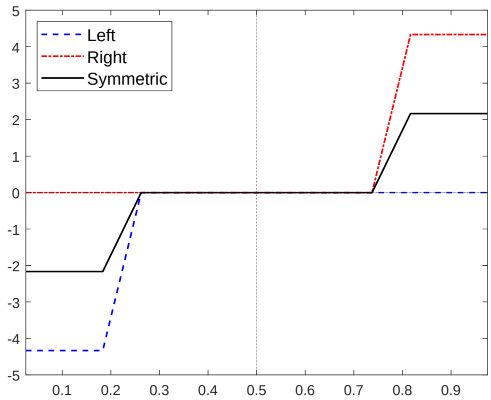

(a)

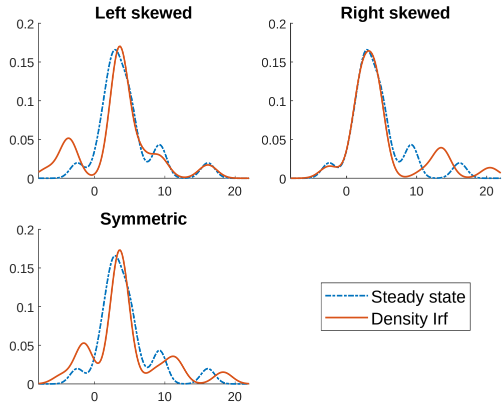

(b)

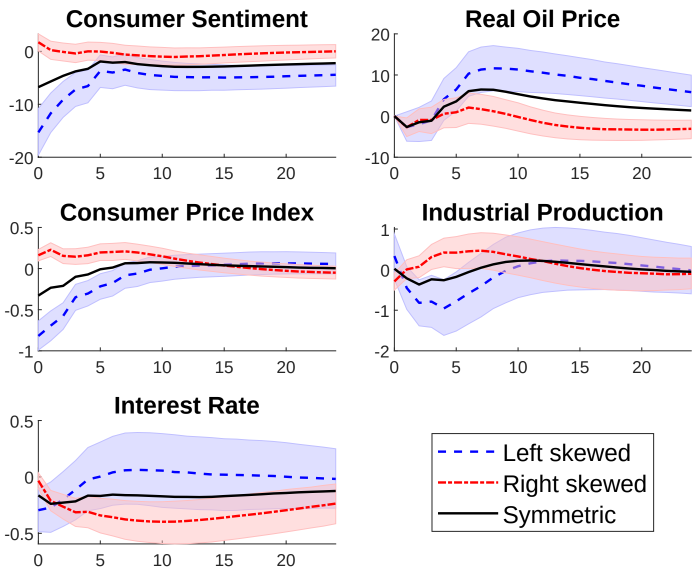

(c)

Figure 9: Panel (a): Three examples of shocks to the quantiles inducing an *asymmetric* increase in dispersion: Left tail shock (dashed blue line), right tail shock (dashed-dotted red line), and symmetric tail shock (solid black line). Panel (b): Corresponding perturbed distributions (solid orange lines), and SS distributions (dashed-dotted blue lines), reflecting the three shocks depicted in Panel (a). Panel (c): IRFs of macroeconomic variables to the three tail shocks: left tail shock (dashed blue line), right tail shock (dashed-dotted red line), and symmetric tail shock (soldi black line). In Panel (b), confidence bands of the symmetric tail shock, shown in Panel (b) of Figure 8, are omitted to help readability.

4.3.3 Asymmetric tail shocks

From the previous subsections we got two main results: (i) variance and kurtosis shocks are stagflationary; (ii) a tail shock is recessionary. Hence, these two dispersion shocks have both a negative effects on IP, but an opposite effect on the CPI. How does this difference arise? It is natural to conjecture that the left tail should be responsible for the negative effect on the CPI, unless tails generate puzzling effects – for instance, a decrease in inflation expectations of the left tail households implies higher inflation. Hence, movements in the left tail should account for larger macroeconomic effects, compared to the ones in the right tail.

To investigate this, we need to analyze the marginal effects of the left and right tails, by separately looking at two asymmetric shocks: the left one, and the right one. The left (right) tail shock is a negative (positive) shock of the first (last) three quantiles, with equal magnitude. The effects of a marginal movement of the tails are obtained by constraining to zero at impact all but the first/last three quantiles. As for the other experiments, the magnitude is chosen such that the overall magnitude of the three shocks (including the symmetric one) is the same. This construction ensures that the symmetric tail shock is included in the two asymmetric ones, thus easing the interpretation of the results. Panel (a) of Figure 9 shows the perturbation applied to each quantile in each of the three scenarios, whereas Panel (b) shows the resulting perturbed smoothed distribution on impact in comparison with the steady state distribution. The latter clarifies what is going on in this exercise. Being the inflation expectation distribution asymmetric, it is clear that the effects of the two shocks are not mirroring. In particular, the right tail is substantially thicker than the left one, and by perturbing the last three quantiles we are confident we do not touch the bulk. This is quite evident from the second plot of Panel (b). This is not the case for the left tail shock. The movement of the first three quantiles includes also part of the bulk. In this case, we choose to be agnostic by picking the same number of quantiles to generate the left and right tail shock.

Panel (c) shows that the two asymmetric tails shocks have very different effects on the macroeconomic variables. Specifically, CPI drops on impact after a left tail shock, while it raises – although negligibly – after a right tail one. IP gets definitely in recessionary territory (after an initial increase) following a left tail shock, while the opposite happens following a right tail one – although again only marginally and not significantly. Thus, the negative effect of the pure tail shock (black line) is mainly attributed to the left tail shock. Regarding the transmission mechanism, consumer sentiment reacts very differently in the two cases. It drastically drops after a left tail shock, symptom of a bad expected future outcome, while it surprisingly increases in response to a right tail shock – although again negligibly. In both cases, monetary policy eases, a response that squares with the responses of IP and CPI for the left tail shock, less so for he right tail one. Therefore, the responses of CPI and IP confirm our prior conjecture, and the answer to the initial question about which tail causes the symmetric tail shock to be recessionary is clear: it is all about the left tail shock. The responses to a right tail shock are muted, while the ones to a left tail shock are definitely contractionary, both for sentiment and for the macroeconomic variables.

Finally, Appendix D analyzes the asymmetric version of the variance and kurtosis shock. Figures 15 shows that there is no difference between the symmetric and asymmetric shocks in

$$\mathrm{P}\left[\mathrm{E}_{i,t}(\pi_{t,t+12}) \in A\right] = \mu_t + \mathbf{x}_{i,t}^{'}\boldsymbol{\beta} + \epsilon_{i,t}$$

Variable	$A = (-\infty, q_3]$	$A = (q_1, q_3]$	$A = (q_2, q_3]$
Hold diploma	-0.0051^{**}	-0.0051^{**}	0.0073^{***}
	(0.0023)	(0.0023)	(0.0015)
Hold degree	0.0109^{***}	0.011^{***}	0.0068^{***}
	(0.001)	(0.001)	(0.0007)
Income quartiles	0.0194^{***}	0.0211^{***}	0.0076^{***}
	(0.0018)	(0.0017)	(0.0011)
Sentiment	0.0322^{***}	0.0303^{***}	0.0090^{***}
	(0.0013)	(0.012)	(0.0008)
R^2	0.02708	0.02674	0.00854
Observations	180226	180226	180226
Effects	Time	Time	Time
Controls	Yes	Yes	Yes
VCV Robust	Yes	Yes	Yes

Table 2: Pooled OLS regression on inflation expectations. The dependent variables take value of 1 if individual i has an inflation expectation that falls within the set A, 0 otherwise. $q_1 = -2.51$, $q_2 = 0.63$, and $q_3 = 1.41$ represent the first three empirical quantiles employed in the study. To the sake of parsimony, here we show the results of some relevant variables. We propose the full specification in Table 3 in Appendix D.

case of variance shocks. Figure 16 shows that the responses to a kurtosis asymmetric shocks are similar to the case of a tail shock commented above. A left kurtosis shock is definitely contractionary, while a right kurtosis shock is inflationary with no effects on IP.

4.3.4 Who populates the left tail?

The empirical evidence that emerged from the asymmetric tail shocks stimulates further investigation. Why left tail exogenous movements generate such big macroeconomic fluctuations? Trying to answer this and further questions about a possible channel is crucial for a deeper understanding of the evidence. We study the individual characteristics of the households forming the MSC. The MSC is a comprehensive survey, collecting inflation expectations, but also a variety of other information, including income and wealth, demographic statistics, personal finance, and so on. In particular, we want to assess if the probability of having low inflation expectations (i.e., being in of the left tail of the inflation expectation distribution) depends on some individual features, that in turn may represent a possible propagation channel for the macroeconomic fluctuations. Table 2 shows the results from three Pooled OLS regressions with time-fixed effects.[17] The dependent variable is a binary transformation of inflation expectations that takes the value of 1 if the individual i inflation expectation is less or equal to the third

[17] Individual fixed effects are not considered given that the MSC is a rotating panel, and the same household remains in the survey for not more than three periods. Thus, it is impossible to track the story of an individual across time.

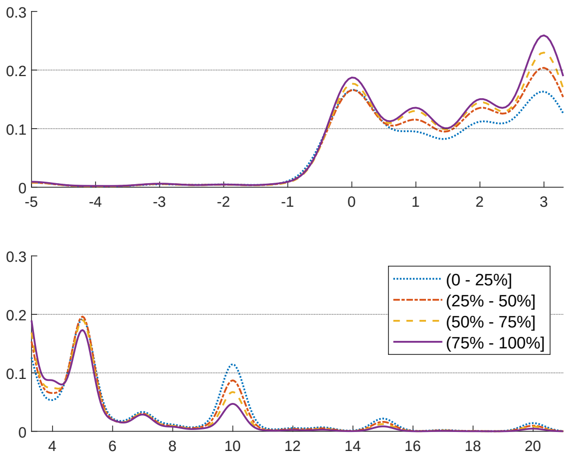

(a)

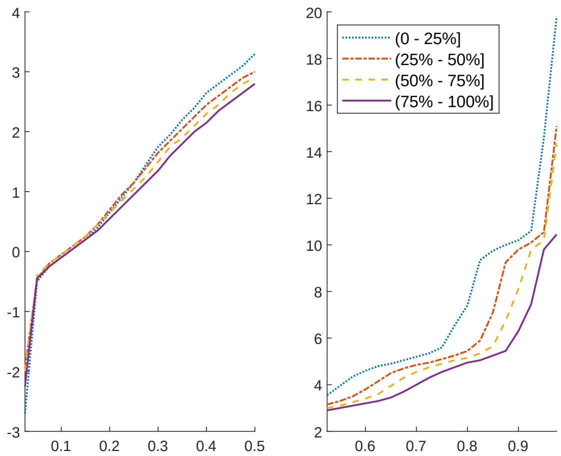

(b)

Figure 10: Panel (a): Kernel smoothed cross-sectional inflation expectation distribution of the pooled dataset, filtered according to four quartiles of income: dotted blue lines (bottom 25%), dashed-dotted orange lines (25% - 50%), dashed yellow lines (50% - 75%), solid purple lines (upper 25%). For ease of visualization, we split the plot into two subplots: the upper one presents the first half of the distribution (up to the median), whereas the bottom one the second half of the distribution. Panel (b): Quantile function associated with the four densities of Panel (a). Here the division in two subplots is made according to $\tau \in (0, 1)$ being lower or bigger than 0.5.

quantile. Column 1 is the benchmark case, where A represents the full left tail. The last two columns prove the results are robust when filtering out excessively low values of inflation expectations. Remarkably, the probability of being in the tails is positively related to the probability of having a college degree and high income. Moreover, despite they show a downward bias in inflation being in the left tail, they tend to be more optimistic, in terms of sentiment. The magnitude and robustness of the coefficients suggest that people in the left tail of the distribution are relatively more educated and have relatively higher capacity to spend, given the higher income, providing some indirect support for the transmission mechanism of expectation shocks to the macroeconomic variables that we found in the BVAR analysis.

4.4 Central bank's communication

In this subsection, we investigate the potential implications of our results for monetary policy communication. Despite households are generally inattentive to monetary policy, our results suggest that there might be scope for a communication strategy, that should aim not only to convey the consensus (mean) or target, but to affect the whole distribution or, better, some targeted location in the distribution. Assume communication is effective so that it can influence households' expectations, one could ask whether it is more important to change the mean of the distribution towards the 2% target or to decrease the disagreement/uncertainty (dispersion shocks) or to target some specific part of the distribution (asymmetric tail shocks).

In order to investigate this, we study the effects of exogenous distribution shocks that discipline the expectations toward the 2% inflation target, as if these shocks were the result of

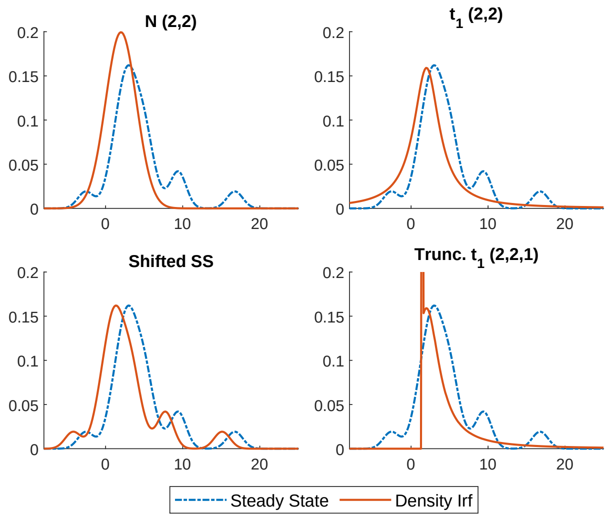

(a)

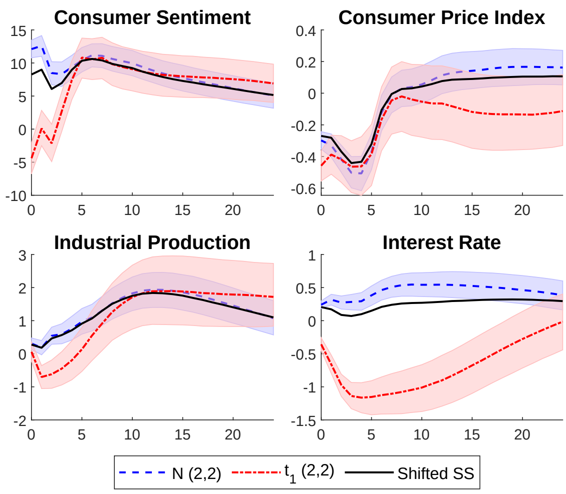

(b)

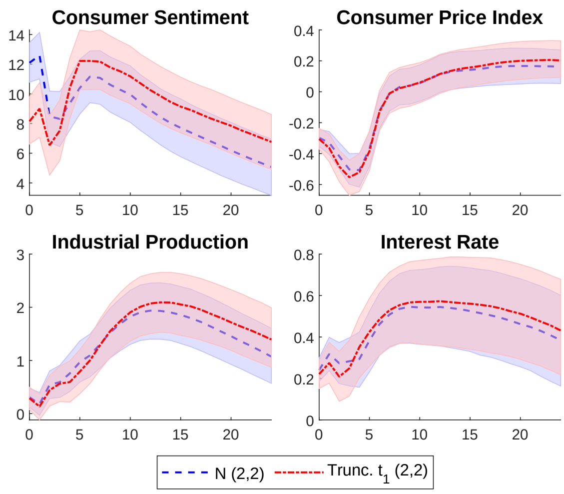

(c)

Figure 11: Panel (a): Perturbed distributions corresponding to the four communication shocks. In order: $\mathcal{N}(2,2)$, $t_1(2,2)$, Truncated $t_1(2,2)$, and Shifted SS (shape preserving mean shock). Panel (b): macroeconomic IRFs forthree shocks: Gaussian (dashed blue line), Student's t (dashed-dotted red line), shifted SS (solid black line). Credible bands of the Shifted SS are omitted to ease interpretation, given the overlap of the IRFs. Panel (c): macroeconomic IRFs for two shocks: Gaussian (dashed blue line) and truncated Student's t (dashed-dotted red line).

effective communication. More specifically, we devise four different scenarios, featuring different effects on the shape of the distribution on impact. The first scenario corresponds to the – possibly ideal – case in which the communication shock makes the initial distribution of expectations equal to a Normal with mean and variance equal to 2. In this case the cross-sectional dispersion is limited – the signal-to-noise ratio here is $\mu/\sigma = 1$ – and fairly disciplined. The second scenario assumes that the distribution after the shock becomes a Student's t with mean and variance equal to 2 and one degree of freedom. While keeping the same first and second moments as in the first scenario, this scenario features a pronounced tail effect, in the form of leptokurtosis. Ceteris paribus, any discrepancies between the IRFs to macroeconomic variables of the first two scenarios will indicate the statistical relevance of dispersion, as suggested by the previous results, also in the case when communication effectively moves expectations in line with the target on average. In the third scenario we translate the steady state distribution to the left so that the mean is equal to 2%. This case describes a central bank communication policy which tries to steer the mean of the distribution, disregarding the shape. The fourth scenario, instead, envisages a communication strategy that just concentrates on wiping the left tail of the distribution out, by truncating the Student's t distribution considered in the second scenario. As opposed to the experiments in the previous subsections, here the shock renormalization doesn't hold anymore, given that each of these four scenarios requires different perturbations.

The first panel of Figure 11 provides, as before, a graphical representation of the movements on impact of the distribution of expectations under the four scenarios, whereas Panels (b) and (c) show the IRfs of macroeconomic variables. Looking at Panel (b), the IRFs to a shock that makes the initial distribution a Normal $N(2,2)$ – dashed blue lines – are very similar to those obtained when the shock produces a shift in the whole distribution – solid black lines. The effect on CPI is negative and on IP is positive, as well as the one on sentiment, suggesting that anchoring inflation expectations to target – from above; recall that the median of the steady state distribution is 3.6% – is beneficial for the real economy, in line with what we saw in the case of a mean shock. The responses we get under the second scenario featuring a Student's t with the same first and second moment but fatter tails – dashed-dotted red lines –, differ from the other two cases. Specifically, consumer sentiment drops on impact and the initial response of industrial production is negative. These differences stress the importance of the tails of the distribution: it is not sufficient to cause a shift of the mean towards the target, but communication should avoid having mass on the tails of the distribution.

Given the results in the previous subsections, however, the next question is: are the two tails alike, or should communication target a specific tail? The next experiment thus truncates the Student's t distribution in Panel (b) such that the masses in the first and second quantiles are shifted to the third quantile of the distribution – see the difference between the top-right and the bottom-right panels in Panel (a). Panel (c) shows that in this case the IRFs are very similar to the case of the Normal $N(2,2)$ analyzed in the first scenario: the response of sentiment and IP are now positive. The comparison between the IRFs in Panel (b) and (c) shows that the negative reaction of sentiment and IP under the Student's t distribution is entirely driven by the fat left tail of the distribution. Hence, a communication strategy that targets the left tail of the distribution, but leaves the right tail fat, causes the same reaction of an ideal communication, depicted here as a Normal $N(2, 2)$. This result suggests that a central bank should focus its communication on convincing consumers that inflation will not be too low. Since these values can be considered as evidently unrealistic, this strategy seems relatively simpler, but most of the improvements in the real activity seems to come through it.

5 Conclusion

We develop a Monetary Policy BVAR augmented with heterogeneous beliefs from the MSC to shed light on the macroeconomic effects of shocks to the entire distribution of short-term inflation expectations. In line with Reis (2022), the main message of this paper is that there is important information content in the higher order moments of the inflation expectation distribution that should be considered in applied work. The periphery of the distribution, often overlooked in parametric approaches, emerges as a location driving substantial macroeconomic fluctuations.

Our analysis reveals the following main results. First, and somewhat surprisingly, households' inflation expectations affect the macroeconomic variables, but not much vice versa. Second, shocks that increase the first two moments of the distributions are stagflationary, and models based on just mean and variance – without accounting for the whole cross-sectional heterogeneity – would underestimate the macroeconomic effects of these shocks. Third, the transmission mechanism of these shocks goes through consumer sentiment, and thus through the relationship between inflation expectations and the perceived future outcome by consumers, in line with microeconomic survey evidence (e.g., Coibion et al., 2019, 2023; Weber et al., 2022). Fourth, shocks to the tails of the inflation expectation distribution are recessionary. Fifth, by introducing asymmetry, we give further evidence on the specific portions of the distribution that trigger the results. In particular, it emerges that left-tail perturbations account for the largest macroeconomic effects. Sixth, we show that households populating the left tail are relatively more educated and have relatively higher income, suggesting that our findings are not due to poorly educated households with no spending ability. Finally, we discuss some implications of our results for central bank communication, revealing a possible complex trade-off for central banks. Communication should focus on decrease the dispersion, and particularly moving mass from the left tail, where pessimistic inflation expectations lie, rather than move the mean towards the 2% target. In other words, a communication targeting pessimistic inflation expectations might be most effective.

Methodologically, our work contributes to the growing literature on addressing dimensionality challenges when incorporating heterogeneous beliefs into macroeconometric models. Our approach, grounded on empirical quantiles, enables us to directly perturb the expectations of specific classes of agents for meaningful economic insights.

References

Adrian, T., N. Boyarchenko, and D. Giannone (2019): "Vulnerable Growth," *American Economic Review*, 109, 1263–89.

Adrian, T., F. Grinberg, N. Liang, S. Malik, and J. Yu (2022): "The Term Structure of Growth-at-Risk," *American Economic Journal: Macroeconomics*, 14, 283–323.

Angeletos, G.-M., F. Collard, and H. Dellas (2018): "Quantifying Confidence," *Econometrica*, 86, 1689–1726.

Antolín-Díaz, J., I. Petrella, and J. F. Rubio-Ramírez (2021): "Structural scenario analysis with SVARs," *Journal of Monetary Economics*, 117, 798–815.

Ascari, G., S. Fasani, J. Grazzini, and L. Rossi (2023): "Endogenous uncertainty and the macroeconomic impact of shocks to inflation expectations," *Journal of Monetary Economics*.

Bachmann, R., T. O. Berg, and E. R. Sims (2015): "Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence," *American Economic Journal: Economic Policy*, 7, 1–35.

Barrett, P. and J. J. Adams (2022): "Shocks to Inflation Expectations," Working Paper 2022/072, International Monetary Fund.

Barsky, R. B. and E. R. Sims (2011): "News shocks and business cycles," *Journal of Monetary Economics*, 58, 273–289.

——— (2012): "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence," *American Economic Review*, 102, 1343–77.

Beaudry, P. and F. Portier (2006): "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, 96, 1293–1307.

——— (2014): "News-Driven Business Cycles: Insights and Challenges," *Journal of Economic Literature*, 52, 993–1074.

Benhabib, J., P. Wang, and Y. Wen (2015): "Sentiments and Aggregate Demand Fluctuations," *Econometrica*, 83, 549–585.

Binder, C. and C. Makridis (2022): "Stuck in the Seventies: Gas Prices and Consumer Sentiment," *The Review of Economics and Statistics*, 104, 293–305.

Binder, C. C. (2020): "Long-run inflation expectations in the shrinking upper tail," *Economics Letters*, 186, 108867.

Blinder, A. S., M. Ehrmann, M. Fratzscher, J. De Haan, and D.-J. Jansen (2008): "Central bank communication and monetary policy: A survey of theory and evidence," *Journal of Economic Literature*, 46, 910–945.

Candia, B., O. Coibion, and Y. Gorodnichenko (2020): "Communication and the Beliefs of Economic Agents," Working Paper 27800, National Bureau of Economic Research.

Chan, J. C. (2021): "Minnesota-type adaptive hierarchical priors for large Bayesian VARs," *International Journal of Forecasting*, 37, 1212–1226.

Chang, M., X. Chen, and F. Schorfheide (2023): "Heterogeneity and aggregate fluctuations," *Journal of Political Economy*, Forthcoming.

Chen, Y., Z. Lin, and H.-G. Müller (2021): "Wasserstein Regression," *Journal of the American Statistical Association*, 0, 1–14.

Clark, T. E. and T. Davig (2011): "Decomposing the declining volatility of long-term inflation expectations," *Journal of Economic Dynamics and Control*, 35, 981–999.

Coibion, O., D. Georgarakos, Y. Gorodnichenko, and M. van Rooij (2023): "How Does Consumption Respond to News about Inflation? Field Evidence from a Randomized Control Trial," *American Economic Journal: Macroeconomics*, 15, 109–52.

Coibion, O., Y. Gorodnichenko, and R. Kamdar (2018): "The Formation of Expectations, Inflation, and the Phillips Curve," *Journal of Economic Literature*, 56, 1447–91.

Coibion, O., Y. Gorodnichenko, and T. Ropele (2019): "Inflation Expectations and Firm Decisions: New Causal Evidence," *The Quarterly Journal of Economics*, 135, 165–219.

Diebold, F. X. and C. Li (2006): "Forecasting the term structure of government bond yields," *Journal of Econometrics*, 130, 337–364.

D'Acunto, F., D. Hoang, and M. Weber (2016): "The Effect of Unconventional Fiscal Policy on Consumption Expenditure," Working Paper 22563, National Bureau of Economic Research.

D'Acunto, F., U. Malmendier, and M. Weber (2022): "What Do the Data Tell Us About Inflation Expectations?" Working Paper 29825, National Bureau of Economic Research.

Galán, J. E. (2020): "The benefits are at the tail: uncovering the impact of macroprudential policy on growth-at-risk," *Journal of Financial Stability*, 100831.

Galvez, J. and J. Mencía (2014): "Distributional linkages between European sovereign bond and bank asset returns," *Available at SSRN 2523863*.

Horváth, L. and P. Kokoszka (2012): "Inference for Functional Data with Applications," .

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

Kamdar, R. (2019): "The Inattentive Consumer: Sentiment and Expectations," 2019 Meeting Papers 647, Society for Economic Dynamics.

Kokoszka, P., H. Miao, A. Petersen, and H. L. Shang (2019): "Forecasting of density functions with an application to cross-sectional and intraday returns," *International Journal of Forecasting*, 35, 1304–1317.

Mankiw, N. G., R. Reis, and J. Wolfers (2003): "Disagreement about Inflation Expectations," *NBER Macroeconomics Annual*, 18, 209–248.

Meeks, R. and F. Monti (2023): "Heterogeneous beliefs and the Phillips curve," *Journal of Monetary Economics*.

Neri, S. (2023): "Long-term inflation expectations and monetary policy in the euro area before the pandemic," *European Economic Review*, 154, 104426.

Petersen, A., C. Zhang, and P. Kokoszka (2022): "Modeling Probability Density Functions as Data Objects," *Econometrics and Statistics*, 21, 159–178.

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

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

Villani, M. (2009): "Steady-state priors for vector autoregressions," *Journal of Applied Econometrics*, 24, 630–650.

Weber, M., F. D'Acunto, Y. Gorodnichenko, and O. Coibion (2022): "The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications," *Journal of Economic Perspectives*, 36, 157–84.

Zhang, C., P. Kokoszka, and A. Petersen (2022): "Wasserstein autoregressive models for density time series," *Journal of Time Series Analysis*, 43, 30–52.

Zhu, C. and H.-G. Müller (2023): "Spherical autoregressive models, with application to distributional and compositional time series," *Journal of Econometrics*.

A Posterior sampling for the hierarchical time varying long run BVAR

This section gives the details of the posterior sampling algorithm for the model employed in this study. The posterior distribution of the parameters implied by the model in equation (1) is intractable. However, the conditional posterior distribution for the following partition of parameters $\{\boldsymbol{\Sigma}, \boldsymbol{\phi}, \boldsymbol{\theta}, \lambda_1\}$ is available. Thus, we simulate from the joint posterior with a Gibbs sampler that sequentially samples from: $\pi(\boldsymbol{\Sigma}|\mathbf{Z}, \boldsymbol{\phi}, \boldsymbol{\theta}, \lambda_1)$, $\pi(\boldsymbol{\phi}|\mathbf{Z}, \boldsymbol{\Sigma}, \boldsymbol{\theta}, \lambda_1)$, $\pi(\boldsymbol{\theta}|\mathbf{Z}, \boldsymbol{\Sigma}, \boldsymbol{\phi}, \lambda_1)$, and $\pi(\lambda_1|\mathbf{Z}, \boldsymbol{\phi}, \boldsymbol{\theta}, \boldsymbol{\Sigma})$.

- The conditional distribution of $\boldsymbol{\Sigma}|\mathbf{Z}, \boldsymbol{\phi}, \boldsymbol{\theta}, \lambda_1$ and $\boldsymbol{\phi}|\mathbf{Z}, \boldsymbol{\Sigma}, \boldsymbol{\theta}, \lambda_1$ are standard, and can be obtained as in Villani (2009).

- For the $\boldsymbol{\theta}$ step, we start rewriting compactly the model as follows. Let $\mathbf{h}_{i,t} = [1, t, ..., t^R, (t-\xi_1)_+^R, ..., (t-\xi_K)_+^R]'$ be the time t vector of basis function for variable i. We represent compactly the long run component as $\boldsymbol{\mu}_t = \mathbf{H}_t\boldsymbol{\theta}$, where

$$\mathbf{H}_t = \begin{bmatrix} \mathbf{h}'_{1,t} & 0 & \cdots & 0 \\ 0 & \mathbf{h}'_{2,t} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{h}'_{N,t} \end{bmatrix}$$

is an $[N \times R]$ matrix that collects all the basis functions, and $\boldsymbol{\theta} = [\boldsymbol{\theta}'_1, ..., \boldsymbol{\theta}'_N]'$ vertically stacks all the free coefficients to estimate. Let $\bar{\boldsymbol{\Phi}} = [\mathbf{I}_N \quad -\boldsymbol{\Phi}]$, $\tilde{\mathbf{z}} = \text{vec}(\bar{\boldsymbol{\Phi}}\bar{\mathbf{z}})$, and $\tilde{\mathbf{H}} = [\bar{\mathbf{H}}'_{P+1}\bar{\boldsymbol{\Phi}}' \quad ... \quad \bar{\mathbf{H}}'_T\bar{\boldsymbol{\Phi}}']'$, where

$$\bar{\mathbf{z}} = \begin{bmatrix} \mathbf{z}_{P+1} & \cdots & \mathbf{z}_T \\ \vdots & \ddots & \vdots \\ \mathbf{z}_1 & \cdots & \mathbf{z}_{T-P} \end{bmatrix}, \quad \bar{\mathbf{H}}_t = \begin{bmatrix} \mathbf{H}_t \\ \vdots \\ \mathbf{H}_{t-P} \end{bmatrix}$$

Then, the model for $\boldsymbol{\theta}$ is a standard linear regression model with general error covariance matrix:

$$\tilde{\mathbf{z}} = \tilde{\mathbf{H}}\boldsymbol{\theta} + \mathbf{e}, \quad \mathbf{e} \sim \mathcal{N}(\mathbf{0}, \tilde{\boldsymbol{\Sigma}})$$

where $\tilde{\boldsymbol{\Sigma}} = \mathbf{I}_{T-P} \otimes \boldsymbol{\Sigma}$. The full conditional posterior of $\boldsymbol{\theta}$ is:

$$\pi(\boldsymbol{\theta}|\mathbf{Z}, \boldsymbol{\phi}, \boldsymbol{\Sigma}, \lambda_1) = \mathcal{N}(\bar{\boldsymbol{\theta}}, \boldsymbol{\Xi})$$

where $\boldsymbol{\Xi} = (\boldsymbol{\Xi}_0^{-1} + \tilde{\mathbf{H}}'\tilde{\boldsymbol{\Sigma}}^{-1}\tilde{\mathbf{H}})^{-1}$, and $\bar{\boldsymbol{\theta}} = \boldsymbol{\Xi}(\boldsymbol{\Xi}_0^{-1}\boldsymbol{\theta}_0 + \tilde{\mathbf{H}}'\tilde{\boldsymbol{\Sigma}}^{-1}\tilde{\mathbf{y}})$.

- The full conditional posterior of λ_1 is given by

$$\lambda_1|\mathbf{Z}, \boldsymbol{\phi}, \boldsymbol{\Sigma}, \boldsymbol{\theta} = \mathcal{GIG}\Big(\gamma_1 - \frac{N^2P}{2}, 2\gamma_2, \sum_{i=1}^{N^2P}\frac{\theta_i^2}{2C_i}\Big)$$

where $\mathcal{GIG}(\cdot,\cdot,\cdot)$ represents the Generalized Inverse Gamma distribution. In this case, C_i is the constant part of the prior variance of each θ_i, ie $C_i = \frac{\sigma_1}{p^{\lambda_p}\sigma_j}$ for own lagged coefficients and $C_i = \frac{\lambda_2\sigma_1}{p^{\lambda_p}\sigma_j}$ for other lagged coefficients.

- In an auxiliary specification, we distinguish between macro to expectation shrinkage and vice-versa. In this case, the conditional posterior of $\lambda_{m\to e}$ is given by:

$$\lambda_{m\leftarrow e}|\mathbf{Z},\boldsymbol{\phi},\boldsymbol{\Sigma},\boldsymbol{\theta},\lambda_{e\leftarrow m} = \mathcal{GIG}\Big(11-\frac{N_yN_eP}{2},20,\sum_{i\in\mathcal{S}_{m\leftarrow e}}^{N^2P}\frac{\theta_i^2}{2C_i}\Big)$$

where $\mathcal{S}_{m\leftarrow e}$ is the index set that collects all the indexes i such that θ_i is a coefficient associated with the macro to expectation block. The conditional posterior of $\lambda_{e\leftarrow m}$ is analogous. The only difference is that $\mathcal{S}_{m\leftarrow e}$ is replaced by $\mathcal{S}_{e\leftarrow m}$, the index set that collects all the indexes associated with the expectations to macro block.

B Scenario analysis

This section explains how the scenario analysis is conducted. The goal is to draw the structural shock that generates the movements of the N_q quantiles we have in mind, using just the shocks of the quantiles and the sentiment. We start by rewriting compactly the SS-BVAR in equation (1) in structural form for $t = T+1$:

$$\boldsymbol{\Phi}_0\bar{\mathbf{z}}_{T+1} = \mathbf{c}+\mathbf{u}_{T+1}, \quad \mathbf{u}_{T+1}\sim\mathcal{N}(\mathbf{0},\mathbf{I}_N) \tag{B1}$$

where $\bar{\mathbf{z}}_t = \mathbf{z}_t - \boldsymbol{\mu}_t$ are the zero unconditional mean stable deviations, and $\mathbf{c} = \sum_{i=1}^{P}\boldsymbol{\Phi}_0\boldsymbol{\Phi}_i\bar{\mathbf{z}}_{T+1-i}$. Moreover, let $\mathbf{r}_z$ be the $N_q\times 1$ vector of quantile movements we want to impose. They can be conveniently expressed in deviations from the constant part of equation (B1), ie $\mathbf{r}_z = \mathbf{R}_z\mathbf{c}+\bar{\mathbf{q}}$. $\mathbf{R}_z = [\mathbf{I}_{N_q} \quad \mathbf{0}_{N_q,N_y+1}]$ is a $N_q\times N$ selection matrix, which picks the first N_q elements of $\mathbf{c}$, whereas $\bar{\mathbf{q}}$ represents the various movements imposed. $\bar{\mathbf{q}}$ varies across experiments, and is always depicted in the first Panel of each Figure. For the mean shock, $\bar{\mathbf{q}}$ is simply a vector of ones, implying a rightward movement of all the quantiles, that doesn't affect the shape of the distribution. The two restrictions can be written as:

$$\mathbf{R}_z\bar{\mathbf{z}}_{T+1} = \mathbf{r}_z, \quad \mathbf{R}_u\mathbf{u}_{T+1} = \mathbf{0}_{N-N_e}$$

$\mathbf{R}_u = [\mathbf{0}_{N_y,N_e} \quad \mathbf{I}_{N_y}]$ selects the last N_y structural shocks. Although the second restriction is on the structural shock, there is one to one map with $\bar{\mathbf{z}}_{T+1}$, ie $\mathbf{R}_u\boldsymbol{\Phi}_0\bar{\mathbf{z}}_{T+1} = \mathbf{R}_u\mathbf{c}$. This allow us to rewrite both compactly as

$$\mathbf{R}\bar{\mathbf{z}}_{T+1} = \mathbf{r}$$

where $\mathbf{R} = [\mathbf{R}_z', \quad \boldsymbol{\Phi}_0'\mathbf{R}_u']'$, and $\mathbf{r} = [\mathbf{r}_z', \quad \mathbf{c}'\mathbf{R}_u']$. The resulting structural shock can be obtained as

$$\boldsymbol{\mu}_u = (\mathbf{R}\boldsymbol{\Phi}_0^{-1})^{+}(\mathbf{r}-\mathbf{R}\boldsymbol{\Phi}_0^{-1}\mathbf{c})$$

where the $\mathbf{X}^{+}$ is the Moore Pensore inverse of $\mathbf{X}$.

C Michigan survey on Consumer attitudes

- *Mnemonic:* MSC.
- *Population*: Cross section of the general public.
- *Organization*: Survey research center, University of Michigan
- *N. of respondents*: 566 (mean), 480 (min), 1459 (max).
- *Type*: Short rotating Panel. For each month, an independent cross-section sample of households is drawn. The respondents chosen in this drawing are then reinterviewed six months later. The total sample for any one survey is normally made up of two- thirds of new respondents, and one-third being interviewed for the second time.
- *Timing*: variable, usually the fourth week of the month.
- *Forecast Horizon*: One year ahead (from January 1978), one and five year ahead (from April 1990).
- *Questions (for inflation expectations):*
  1. During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?
  2. By about what percent do you expect prices to go (up/down), on average, during the next 12 months?
- *Questions (for the Consumer Sentiment construction):*
  1. We are interested in how people are getting along financially these days. Would you say that you (and your family living there) are better off or worse off financially than you were a year ago?
  2. Now looking ahead–do you think that a year from now you (and your family living there) will be better off financially, or worse off, or just about the same as now?
  3. Now turning to business conditions in the country as a whole–do you think that during the next twelve months we'll have good times financially, or bad times, or what?
  4. Looking ahead, which would you say is more likely–that in the country as a whole we'll have continuous good times during the next five years or so, or that we will have periods of widespread unemployment or depression, or what?
  5. About the big things people buy for their homes–such as furniture, a refrigerator, stove, television, and things like that. Generally speaking, do you think now is a good or bad time for people to buy major household items?

D Additional results

Here we show some additional results that could be of interest.

Figure 12 shows a 3-D plot graphical representation of the time series of the Kernel smoothed inflation expectations distribution from the Michigan Survey of Consumers.

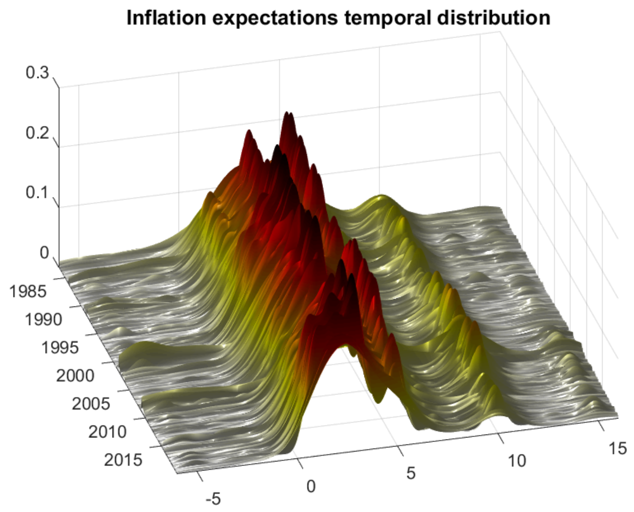

Figure 12: Kernel smoothed inflation expectations distribution from the Michigan Survey of Consumers. Monthly data from January 1983 to December 2019.

The next Figure 13 shows the estimated trends of the macroeconomic variables and the associated cyclical variations, i.e., deviations from the trend.

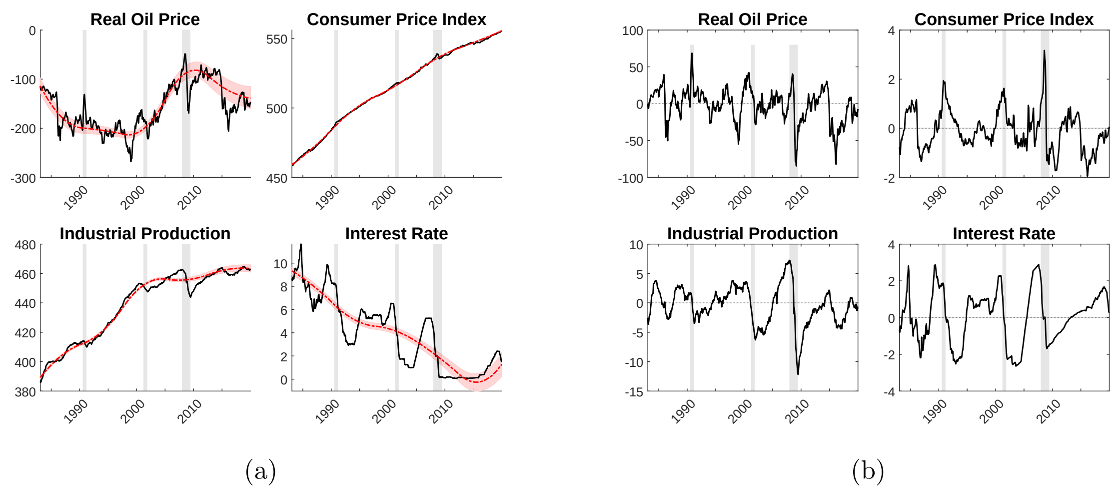

Figure 13: Panel (a): Estimated posterior long-run component for the four macroeconomic variables. Panel (b): Associated gaps.

Figure 14 shows the contribution to the FEVD of the thirteen quantiles of the inflation expectation distribution of the macroeconomic variables, including oil price shocks. Comparing it with 4 in the main text, that excludes oil prices, we can see that oil price shocks are an important source of variation. However, the overall variance explained by the macroeconomic variables – including oil shocks – remains low, so that exogenous shocks to expectations remain by far the most important determinant of consumers' beliefs.

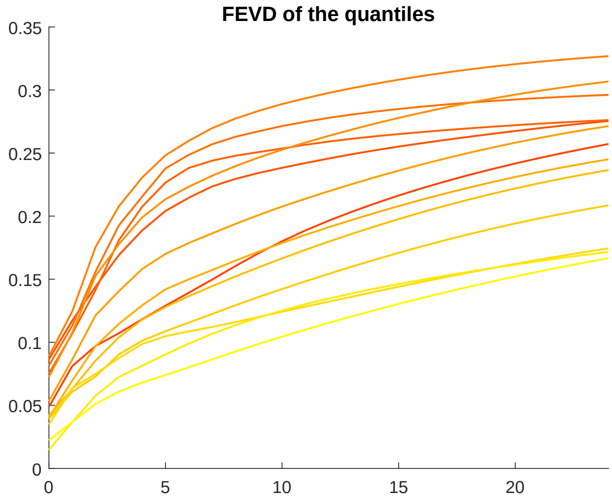

Figure 14: Panel (a): Contribution of the three macroeconomic variables and real oil to the FEVD of the 13 Quantiles. Different colors indicate the location of the quantiles in the distribution. Starting from the left tail (red), to the right tail (yellow). Panel (b): FEVDs of the macroeconomic variables at different horizons (x-axis).

The next two Figures shows the IRFs to an asymmetric variance and kurtosis shock. As mentioned in the main text, Figures 15 shows that there is no difference between the symmetric and asymmetric shocks in case of variance shocks. Figure 16 shows that the responses to a kurtosis asymmetric shocks are similar to the case of a tail shock commented above. A left kurtosis shock is definitely contractionary, while a right kurtosis shock is inflationary with no effects on IP.

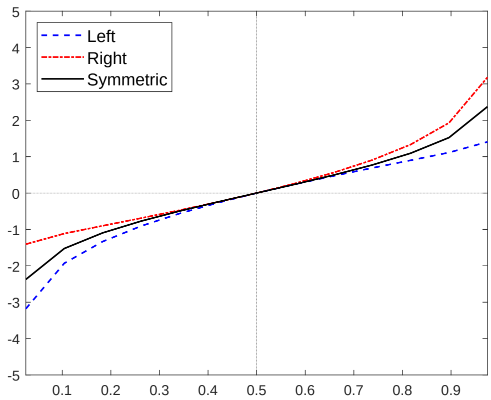

(a)

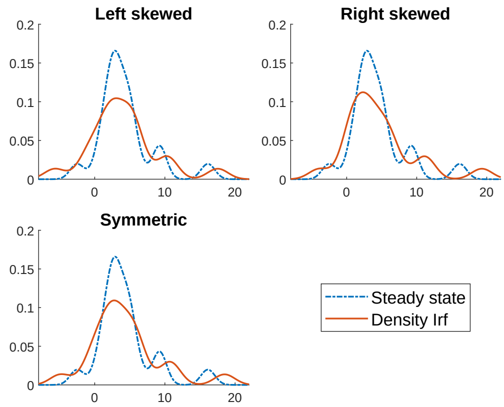

(b)

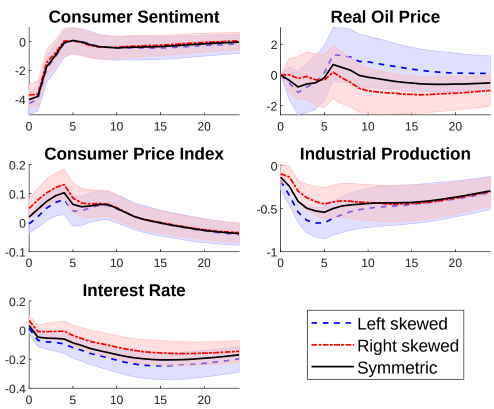

(c)

Figure 15: Panel (a): Left variance shock (dashed blue line), right variance shock (dashed-dotted red line), and symmetric variance shock (black line). Panel (b): Perturbed distributions (solid orange lines), and SS distributions (dashed-dotted blue lines), reflecting the three shocks depicted in Panel (a).
Panel (c): IRFs of Macroeconomic variables for the three variance shocks: left variance shock (dashed blue line), right variance shock (dashed-dotted red line), and symmetric variance shock (black line). Confidence bands of the symmetric variance shock are omitted to ease interpretation, given the overlap of the IRFs.

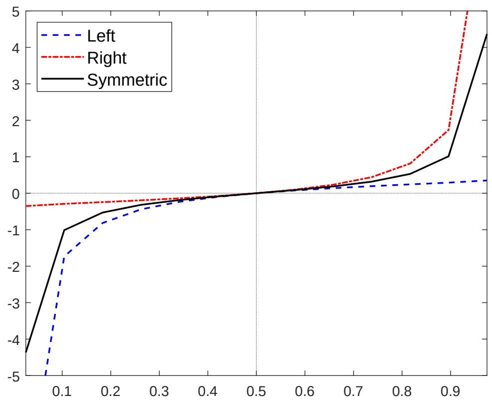

(a)

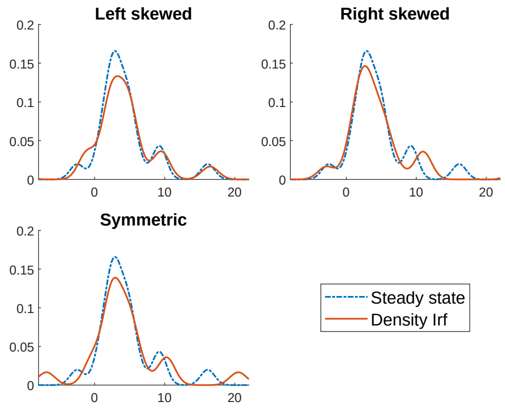

(b)

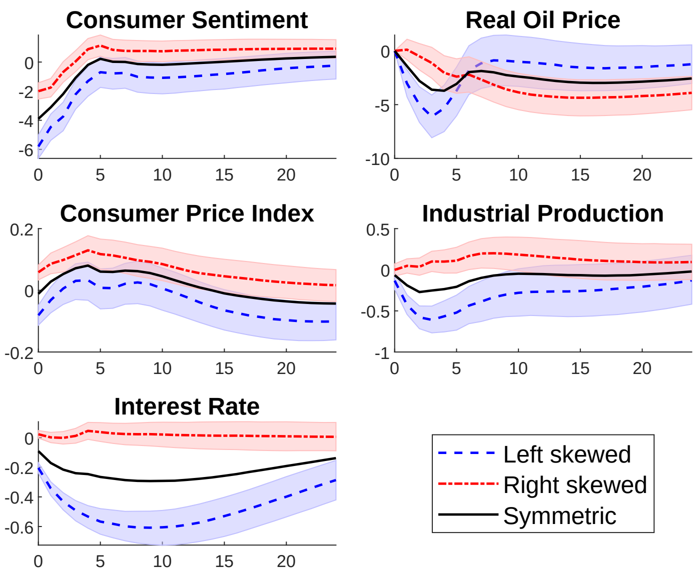

(c)

Figure 16: Panel (a): Left kurtosis shock (dashed blue line), right kurtosis shock (dashed-dotted red line), and symmetric kurtosis shock (black line). Panel (b): Perturbed distributions (solid orange lines), and SS distributions (dashed-dotted blue lines), reflecting the three shocks depicted in Panel (a). Panel (c): IRFs of Macroeconomic variables for the three kurtosis shocks: left kurtosis shock (dashed blue line), right kurtosis shock (dashed-dotted red line), and symmetric kurtosis shock (solid black line). Confidence bands of the symmetric kurtosis shock are omitted to ease interpretation, given the overlap of the IRFs.

The following Table presents the full specification of the regression results in Table 2.

$$\mathrm{P}\left[\mathrm{E}_{i,t}(\pi_{t,t+12}) \in A\right] = \mu_t + \mathbf{x}_{i,t}^{'}\boldsymbol{\beta} + \epsilon_{i,t}$$

Variable	$A = (-\infty, q_3]$	$A = (q_1, q_3]$	$A = (q_2, q_3]$
Birth cohort (*lin. effect*)	0.0007***	0.0005***	0.0002***
	(0.00003)	(3.5^{-5})	(2.5^{-5})
Birth cohort (*sq. effect*)	0.0083***	0.0050**	0.0161***
	(0.0021)	(0.0021)	(0.0014)
Gender effect (*Male = 1*)	−0.0065***	−0.0019	0.0064***
	(0.0030)	(0.0029)	(0.0020)
Regional effect (*South = 1*)	−0.0113**	−0.0077**	0.0082***
	(0.0033)	(0.0032)	(0.0022)
Regional effect (*West = 1*)	−0.0093***	−0.0017	0.0085***
	(0.0031)	(0.0031)	(0.0021)
Regional effect (*Midwest = 1*)	−0.0113**	−0.0131***	0.0130***
	(0.0044)	(0.0043)	(0.0027)
Education (*Diploma = 1*)	−0.0051**	−0.0051**	0.0073***
	(0.0023)	(0.0023)	(0.0015)
Education (*Degree = 1*)	0.0109***	0.011***	0.0068***
	(0.001)	(0.001)	(0.0007)
Income quartiles	0.0194***	0.0211***	0.0076***
	(0.0018)	(0.0017)	(0.0011)
Sentiment	0.0322***	0.0303***	0.0090***
	(0.0013)	(0.012)	(0.0008)
R^2	0.02708	0.02674	0.00854
Observations	180226	180226	180226
Effects	Time	Time	Time
Controls	Yes	Yes	Yes
VCV Robust	Yes	Yes	Yes

Table 3: Full specification of the pooled OLS regression on inflation expectations. The dependent variables take value of 1 if individual i has an inflation expectation that falls within the set A, 0 otherwise. $q_1 = -2.51$, $q_2 = 0.63$, and $q_3 = 1.41$ represent the first three empirical quantiles employed in the study.