UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

Heidelberg University

Alfred Weber Institute of Economics

Dissertation:

Essays on Inflation and Inflation Expectations: The Role of Media and Regime-Dependence

A Thesis Submitted to Obtain the Degree of:

Doctor of Philosophy (Ph.D.)

Presented by:

Sarah Arndt

Supervisor:

Prof. Dr. Zeno Enders

October, 2025

Dissertation

Submitted to the

Alfred Weber Institute of Economics

Heidelberg University, Germany

for the degree of

Doctor of Philosophy (Ph.D.)

Presented by

Sarah Arndt

Supervisor

Prof. Dr. Zeno Enders

Subject

Essays on Inflation and Inflation Expectations:
The Role of Media and Regime-Dependence

Acknowledgements

First and foremost, I would like to thank my supervisor and co-author, Zeno Enders. Thank you for giving me the freedom to follow and explore my own ideas and for never discouraging even the raw and undeveloped ones. I am very grateful for your constructive feedback, which always pushed me in the right direction and helped me improve my work. I also truly appreciate your patience with my strong-willed nature and with the fact that my habit of speaking up sometimes ruffled feathers in the department. You never held it against me, and instead encouraged me to keep questioning and refining the status quo.

Special thanks go to my co-author, Farah Tohme. The last mile is always the hardest, but working with you turned it into one of the most rewarding parts of this journey. I cherish your positive nature, creativity, and critical thinking, which pushed and improved our project, opened up many new research avenues, and gave me the energy to keep going.

I would like to thank my colleagues at the Alfred-Weber-Institut, especially Monika Lülf and Barbara Neef, for their constant administrative support, Christoph Schineller and Oliver Zeh for their IT support, and, last but not least, my dear fellow PhD students and postdocs for their moral support.

I am deeply grateful to my friends Hanna, Hélène, Caro, Lora, Klara, Hendrik, Prachi, Alisha, and Annelie, who patiently listened whenever I was struggling and never failed to lift my spirits. Knowing I can always count on your support is the greatest gift.

This thesis would not have been possible without the support of my family, to whom I am deeply indebted. First, I want to thank my dear mother for her unwavering support, on which I can always rely. You are not only an outstanding mom, but also my mentor and coach from whom I will never stop learning. You taught me to take risks and follow my passion, and your advice has helped me overcome many difficult situations. This thesis is dedicated to you, as well as to my grandmother, who told me already back in high school to hurry up with getting a PhD so she could live to see it. I am grateful I managed to do just that. The third pillar of this family bond is my brother and his wife, with whom I always enjoy discussing new AI-related ideas, and whose enthusiasm and insights continually inspire me.

"Es ist nicht deine Schuld, dass die Welt ist, wie sie ist.
Es wär nur deine Schuld, wenn sie so bleibt."

— Die Ärzte, *Deine Schuld* (2003)

Contents

List of Figures

List of Tables

1. Introduction

At the heart of this dissertation lies the study of one macroeconomic variable: inflation, i.e., *"the general increase in prices and fall in the purchasing value of money"*, as defined in the Oxford English Dictionary (Oxford University Press, nd). Prices are formed on the market when a supplier and a buyer meet, but in reality, this process is much more complex than the simple intersection of two linear functions, as one can learn in the introductory economics lecture. Beyond this very simple abstraction, prices and the change in prices reflect highly non-linear aspects such as market participants' beliefs and expectations, personal preferences and valuations, private information and bargaining power, as well as external factors that are outside of their control. Therefore, this thesis goes beyond the supply and demand diagram and explores what factors affect firms in their price-setting behavior and consumers' inflation expectations, which ultimately might affect their spending willingness.

In seeking to understand these macroeconomic dynamics, economists do not simply uncover an objective truth; rather, they illuminate the aspects of reality they consider most salient, inevitably projecting their own theoretical frameworks onto the data. In this sense, models became the primary tool through which economists formalize their understanding of the world. So in essence, models are *opinions dressed up as math*, as Marc-Uwe Kling (2019) stated in his novel *Qualityland*. This is not meant as criticism of the economic sciences; it is rather intended to highlight that this thesis also presents the views of what mechanisms my co-authors and I believe to be in place and that play a role in the development of price changes that we measure.

In this sense, this thesis contributes to two modelling directions. The first is to model agents' information set using newspaper articles. Journalists act as intermediaries, observing factual events which they translate into text, sometimes adding background explanations to enhance accessibility for people who consume their news. As journalists are also human, their reporting is usually subject to some sort of bias and hence may influence the reader's set of beliefs. By analyzing this type of data, one can, on average, construct a proxy of what households know and believe. This can then serve as an alternative to information-provision survey experiments, with the advantage of capturing the natural information environment to which households are exposed over time, rather than relying on artificial, one-off information treatments. Unlike survey experiments, which typically deliver selected facts in a controlled setting and at a single point in time, using newspaper articles allows for the construction of rich, high-frequency time series that reflect the evolving narratives, framing, and sentiment of real-world news coverage. This approach thereby improves external validity, as it more closely mirrors the context in which households actually form their expectations and beliefs.

As a backdrop, one major criticism that emerges when using newspaper articles is the general impression that nowadays fewer and fewer people rely on print media. Instead, they consume other traditional media like television or radio, but also non-traditional sources such as various social media platforms. Systematically analyzing the latter seems almost impossible, as the feed on these platforms varies by user and changes whenever the applications are reopened or reloaded. Acknowledging this fact, one argument in favor of using newspaper articles is that, despite the growing prevalence of digital and social media, traditional news outlets continue to play a central role in shaping public discourse and setting the broader informational agenda. Much of the content that circulates on social media platforms, such as trending stories and public debates, often originates from or is amplified by established journalistic sources. So even if not everybody is reading newspapers, a substantial share of their information will reach the wider public eventually.

The second modeling avenue to which this thesis contributes is connected to the recent advances in large language models (LLMs) brought to a new level by the generative transformer architectures developed by OpenAI. These models are capable of processing and interpreting vast amounts of unstructured text data, mimicking the way humans read and extract meaning from text. Trained on a massive corpus of human-written texts of all kinds, these language models are able to recognize context, nuance, and implicit meaning in natural language. For this reason, I interpret LLMs as models of human cognition, capable of understanding, interpreting, and generating texts much like humans do. This point of departure opens up a wide range of possibilities: it allows us to automatically and systematically process large quantities of text and extract information such as sentiment or concrete signals, which we can then use to classify articles and transform qualitative text data into time series suitable for empirical analysis.

But this approach has to be considered with caution, too. The type of LLMs used in this thesis work by predicting the most likely next sequence of text. Therefore, the generated text is never purely deterministic and running the same prompt twice might produce two different results. In this sense, LLMs resemble humans, as two different people might interpret a text in two different ways. On the other hand, if the task is to classify a large quantity of text, humans might become tired or bored over time and therefore start to only skim texts, resulting in increasing false classifications. In comparison, the advantage of the machine is its consistency in its abilities and its speed.

This thesis is structured into three self-contained parts. The first of the following three chapters focuses on the price setter, i.e., firms, and analyzes how the inflation environment—specifically, overall inflation volatility—affects their price-setting behavior. Motivated by the latest high-inflation period, this question has become particularly relevant, as firms and policymakers alike have had to adapt to rapidly changing price dynamics. In an empirical exploration, we find that cost shocks are transmitted faster to consumer prices when there is higher volatility in price changes compared to more stable times. Then, in a theoretical model, we aim to explain this observation by the ability of firms to invest in price flexibility.

The second part of this thesis shifts the focus to the price takers, i.e., households or consumers, and examines how media, particularly newspapers, influence their expectation formation process and the broader macroeconomic implications. Chapter 3 investigates how the tone towards inflation varies across different types of German newspapers and how this, in turn, shapes the inflation expectations of their respective readerships. In addition to household characteristics such as income and education, the choice of newspaper emerges as an important factor contributing to heterogeneity in household inflation expectations.

Chapter 4 then builds on the previous one by employing similar methods of text analysis with LLMs. We explore a new concept of *fiscal narratives*, which we measure from German newspaper texts. Analyzing these narratives in the media allows us to continuously track changes in the public debate on fiscal policy and to quantify how shifts in the framing of fiscal issues may influence the beliefs and expectations of households over time. Going beyond the microdata level, we are able to show that narratives in the media can also have macroeconomic effects by shaping the transmission of government spending shocks.

Taken together, the three chapters illustrate how inflation dynamics cannot be fully understood without accounting for the information that firms and households actually receive—and how modern text-based methods allow us to measure that information with far greater precision than was possible even a few years ago. By bringing newspaper content and large language models into the macroeconomic toolkit, the thesis not only sheds light on recent bouts of high and volatile inflation but also paves the way for future work that can trace, in real time, how evolving narratives feed back into prices and policy.

2. The Transmission of Supply Shocks in Different Inflation Regimes

Authors

Sarah Arndt & Zeno Enders [1]

Abstract

We show that the impact of supply on consumer prices is state-dependent. First, we let the data determine two inflation regimes and find that they are characterized by high and low inflation volatility. We then identify supply shocks using instrumental variables based on outliers in the producer price series. Such shocks exhibit a more substantial and more persistent effect on downstream prices during periods of elevated inflation volatility compared to phases of more stable consumer price growth. Exogenously differentiating regimes by the level of inflation or the shock size does not reveal state dependency. The evidence supports a model in which producers optimally invest in price flexibility. This model predicts that stricter inflation targeting lowers inflation volatility in two ways: it reduces price flexibility and, consequently, the pass-through of all shocks to inflation on top of the standard channel that affects demand.

Keywords

Inflation regimes, supply shocks, monetary policy, cost pass-through, producer prices

2.1 Introduction

Policymakers have, particularly during times of rising inflation, voiced the suspicion that the reaction of inflation to external shocks is not stable over time but depends on the level or volatility of inflation itself.[2] Such changing dynamics would be particularly significant for central banking, impacting inflation forecasts and the expected outcomes of monetary policy actions. Specifically, inflation projections often hinge on assumptions regarding the speed and extent to which changes in producer prices are transmitted to consumer prices. These considerations are crucial when central banks aim to contain price

[1] **Acknowledgements:** We thank Hervé Le Bihan, Christoph Große Steffen, Matthias Meier, and workshop participants at the HeiTüHo Workshop, the Banque de France, the ECB, the German Council of Economics Experts, and the T2M conference for helpful discussions. Part of this research was conducted while Enders was a visiting scholar at the Banque de France, whose hospitality is gratefully acknowledged.

[2] Philip Lane, Member of the Executive Board of the ECB, writes on November 25, 2022: "Our corporate contacts started [towards the end of 2021] expressing more concern about the persistence of input cost pressures, raising their price expectations for 2022 (also in view of rising energy prices). [...] Since the beginning of this year, many contacts also told us that prices would be increased more frequently." (Lane, 2022) More frequent price changes would alter the nature of the inflation process profoundly, as regards, e.g., the strength and speed of cost pass-through to inflation.

pressures generated by supply shocks.[3] Relying on theory for this question is difficult, as alternative models of nominal rigidities, such as menu costs or Calvo pricing, yield different predictions for the pass-through of supply shocks to consumer prices. Consequently, identifying changing inflation dynamics also informs us about the validity of certain model assumptions.

We investigate this issue empirically by analyzing whether and when inflation dynamics undergo general changes. Using US data, we uncover two regimes by estimating a Markov-switching process based on inflation dynamics. Crucially, we do not restrict the regimes to depend on some exogenous indicator, such as an inflation threshold, but let the inflation process itself endogenously determine them. It turns out that inflation volatility (quick changes in inflation rates) plays a more significant role in determining the regimes than its level. More precisely, if annualized monthly inflation changes by more than 5.2 pp. (as in April, May, and July 2022), the economy is likely to be in a high volatility regime.[4]

In a second step, we investigate state-dependent causal effects of a shock to producer prices—provided by the Bureau of Labor Statistics—on downstream price growth. Starting in 1948, we estimate how supply shocks to the crude material PPI dynamically affect consumer prices. We also investigate the effects on intermediate stages of the production process. We rely on PPI data as we are interested in more broadly defined supply shocks instead of price movements of a single input factor, which generalizes the results. Given that, e.g., crude materials display a much larger variance compared to consumer prices, PPI price processes are noisier. We, therefore, use movements in the crude-material PPI series that exceed normal fluctuations in input prices and move material prices and production in different directions as instruments for supply shocks.

Our results show that in periods of high inflation volatility, downstream prices, including the consumer price index (CPI), react much more strongly to cost shocks on impact and in subsequent months. In this regime, prices are arguably more flexible and, hence, react more promptly to shocks. We validate our results for general supply shocks by estimating the responses to a specific one, i.e., oil-supply shocks as identified by Baumeister and Hamilton (2019). Again, the CPI exhibits a swifter and more pronounced reaction in the high-volatility state.[5]

To emphasize the critical role of inflation volatility in determining regimes, we explore whether similar state dependencies emerge when departing from the endogenous determination of regimes via the Markov-switching model. Our long sample—879 months—allows us to disentangle periods of large shocks, high inflation, and high inflation volatility. These episodes are correlated, but not identical and are of different natures. Specifically, we repeat our analysis but condition regimes exogenously on the level of CPI inflation or the size of the shocks. Both separations fail to generate a state dependency that comes close

[3]See Sinn (2021) for an early warning of the 2021/22 surge in inflation based on rising producer prices and the implications for monetary policy.

[4]Here and the following, we use the words state and regime interchangeably.

[5]Our findings square well with the observation in Borio et al. (2021) that 'salient,' i.e., large and positive, sectoral price movements displayed a lower pass-through to core PCE inflation during the great moderation, compared to previous periods.

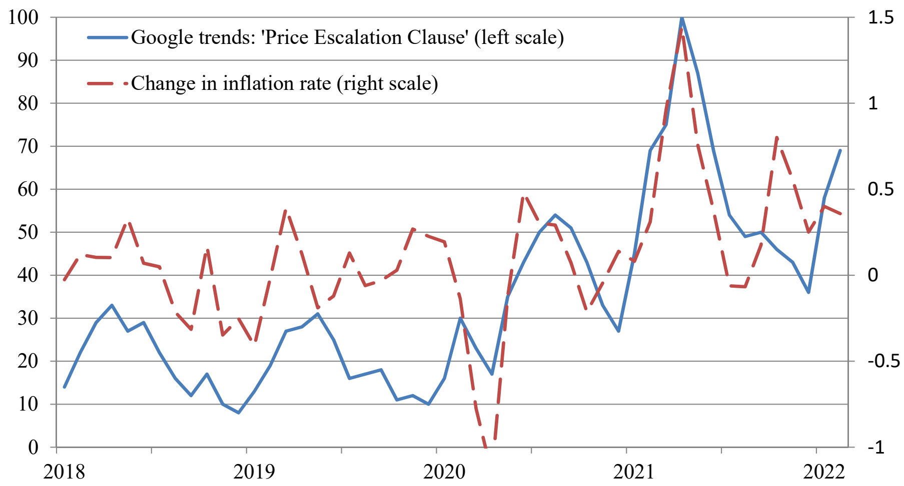

Figure 2.1: Escalation clauses

Notes: Index for Google searches of 'Price escalation clause' (left axis) and monthly change in annualized s.a. CPI inflation rate in percentage points (right scale).

to the one induced by inflation volatility.

Our findings can be explained by firms' quicker price adjustments when facing higher price volatility in their sales markets. This explanation is supported by anecdotal evidence from the 2021/22 surge in inflation. Figure 2.1 depicts Google searches for the term 'Price escalation clause' alongside the change in the CPI inflation rate. If agreed upon in contracts between seller and buyer, these clauses automatically adjust sales prices based on changes in the seller's input costs.[6] That is, widespread use of these clauses implies a much faster price reaction to upstream cost changes, significantly altering inflation dynamics.

Interest in this kind of clause is, as visible in the figure, correlated to the *change* in the inflation rate, peaking in the spring of 2021. This coincided with a swift global rise in input prices due to several factors, among them strained global supply chains. Survey evidence corroborates this observation, as 34% of sampled German firms in the Bundesbank Online Panel reported using price escalation clauses from 2021 onward, compared to only 17% before 2021. Regarding economic theory, our results, therefore, speak in favor of models in which prices react quicker to shocks in the face of higher inflation volatility.

We propose a model based on Devereux (2006) in which price setters can invest in the flexibility of their prices.[7] The crucial difference to other models of state-dependent pricing,

[6] The use of price escalation clauses is not just a recent phenomenon in the US; articles dating back to the 1940s already mention these clauses. For example, Mack (1946) describes different variations and provides advice for buyers facing escalation clauses.

[7] We build our theoretical explanation on Devereux (2006) since his model setup captures the essential determinants for a firm's decision to invest in price flexibility in the most parsimonious way. Moreover, it represents a straightforward implementation of price-escalation clauses in a theoretical framework. Alternatively, but in a very similar spirit, observation costs in a menu cost model as in Álvarez et al. (2018) would also predict that higher volatility leads to more frequent price reviews and, hence, a higher cost pass-through. Rational-inattention models work in a similar way (Mackowiak and Wiederholt, 2009).

such as menu-cost models, is the assumption that firms have an influence on price-setting costs if, in anticipation, they take adequate measures, such as using price-escalation clauses in new contracts. In the presence of strategic complementarities in price setting, the payoff of being able to react quickly to new developments is higher in times of elevated inflation volatility. This increased incentive to invest in price flexibility explains our finding of a more substantial pass-through of cost shocks during periods of volatile inflation. The model predicts a 'double dividend' to inflation targeting in terms of reducing inflation volatility, as it leads to a lower pass-through of shocks to inflation through the traditional direct channel of altering demand, but also indirectly via reducing optimal price flexibility.[8] In contrast, monetary policy that is more accommodating in the face of supply shocks tends to increase price flexibility.

In a dynamic extension of the model, we follow Kimura and Kurozumi (2010) and allow firms to choose an optimal price-setting frequency, based on overall volatility and price-setting costs. The predicted inflation responses to supply shocks in high and low-volatility regimes are reasonably close to our empirical findings. In line with the predictions of the analytical model, we find that stricter inflation targeting dampens the inflation response to cost-push shocks via an endogenous reduction of the price-setting frequency, on top of the standard demand channel. This effect is particularly strong for the high-volatility case. We also compare our empirical results to predictions of prominent alternative pricing models. In menu-cost models, as developed by, e.g., Golosov and Lucas (2007), the shape of CPI responses depends strongly on the shock size, which we do not find in our data for cost-push shocks. Furthermore, standard Calvo price setting would not predict any state dependency at all.

Despite the important implications, surprisingly little research has focused on the pass-through of shocks to consumer prices in different inflation regimes until recently. Given the policy relevance of this question, most existing research was conducted in policy institutions. By using Granger-Causality tests, Weinhagen (2002, 2016) demonstrates that upstream changes in prices explain price changes at each stage of production in the BLS PPI data, while more downstream price changes do not Granger-cause price changes. Bobeica et al. (2020, 2021) concentrate on the pass-through of labor costs to output prices, considering two regimes that depend on whether the level and volatility of inflation are above or below their historical means.

Their findings, based on a Cholesky decomposition to identify labor cost shocks, indicate a quicker and more substantial pass-through in the high-inflation regime. Similarly, the Bank for International Settlements (2022) investigates the pass-through of relative price changes, oil price shocks, and exchange-rate movements into consumer prices, finding them to be dampened in periods of inflation below 5% (see also Borio et al., 2021, 2023). De Santis and Tornese (2023) find a stronger transmission of energy supply shocks on consumer prices in high-inflation regimes, too, while Ascari and Haber (2022) estimate more substantial price effects of monetary policy shocks, as identified by Romer and Romer (2004), in

[8] See also Kimura and Kurozumi (2010) as well as Paciello and Wiederholt (2014) for a related theoretical mechanism in a context of rational inattention.

high-inflation regimes and for large shocks.

Using micro data, Vavra (2014) shows that price changes become more dispersed during recessions and that this dispersion is high when more products are changing prices. Our model aligns with these observations, as recessions typically feature higher volatility and higher volatility increases the share of firms that are able to change prices.[9]

Our approach differs from the above studies in that we analyze the effects of general supply shocks, derived by a novel identification scheme, on prices in later stages of production. Importantly, when identifying different inflation regimes, we do not impose a threshold of inflation or its volatility but let the inflation process itself determine the regimes. By doing so, we uncover the significance of inflation volatility in determining the regimes, a factor that has not been considered so far.[10] Using a very long time series allows us to disentangle the effects of high inflation vs. high inflation volatility. This is not an easy task, as, e.g., inflation surges are generally not one-time spikes but lead to a prolonged period of inflation movements, see Blanco et al. (2025). Our sample, however, features periods of high volatility during times of higher and lower inflation levels. Moreover, in our sample, some large shocks trigger high-volatility phases, others do not.

We also contribute to the literature on the general pass-through of cost shocks.[11] A large part of this literature centers on the exchange-rate pass-through (see, e.g., Taylor, 2000; Campa and Goldberg, 2005; International Monetary Fund, 2006; Auer and Schoenle, 2016; Álvarez et al., 2017; Enders et al., 2018; Bonadio et al., 2019). A recurrent finding is a falling exchange-rate pass-through over time until recently, in line with our result that lower inflation volatility is associated with less frequent price adjustments. Amiti et al. (2019) and Muehlegger and Sweeney (2022) consider cost shocks more broadly and find strong strategic complementarities in price setting, an important element in our explanation of the role of CPI inflation volatility in price setting.[12]

The remainder of this paper is organized as follows. Section 2.2 outlines our methodology, including shock identification. Section 2.3 presents the results, with robustness checks discussed in Section 2.3.5. Section 2.5 develops the model, and Section 2.6 concludes.

[9] Berger and Vavra (2019) also find evidence for time-varying responsiveness of prices to shocks. Vavra (2014) favors a theoretical explanation based on a menu-cost model featuring shocks to the volatility of idiosyncratic firm productivity, see also Hall (2023). Higher volatility reduces the effect of aggregate demand shocks on output in this context. As described above, our empirical results for supply shocks do not support menu-costs models.

[10] In fact, the empirical literature on state-dependent inflation dynamics typically focuses on the effect of the level of inflation without separating it from the impact of its volatility (see, e.g., Álvarez et al., 2019).

[11] Our paper is also related to studies on the price-setting behavior of firms. Given the vast number of significant contributions in this field, we cannot even give a partial overview of this literature here and thus focus on the most directly related studies.

[12] In a similar vein, using surveys, Blinder et al. (1998) and Fabiani et al. (2005) find that firms hesitate to change prices due to the fear of losing customers to competitors. The importance of competitors' prices is further underlined by Dedola et al. (2022), who, employing micro data, ascertain that the pass-through of import cost shocks is lower for larger firms than for smaller ones, suggesting a role for strategic complementarities. Similarly, Gödl-Hanisch and Menkhoff (2023), also using micro data, show that the pass-through of individual cost shocks undershoots that of aggregate shocks by 40%, likely an effect of strategic complementarity. Moreover, they find a more pronounced pass-through for firms that are uncertain about their future business situation, aligning with our result of a higher pass-through in volatile times.

2.2 Methodology

2.2.1 A Markov-switching model to detect inflation regimes

We detect inflation regimes by employing a Markov-switching autoregressive model (MS-AR) based on log differences of US CPI data. This type of model was introduced by Hamilton (1989). The basic modeling idea is that there are different states s_t of the AR model characterized by regime-specific model coefficients and error variances. A discrete first-order Markov process governs the transition between these states. In our setting, we restrict the model to have two states. The Markov process can then be described by the following transition matrix:

$$P = \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix}, \quad \text{where} \quad p_{i,j} = Pr(s_{t+1} = j | s_t = i),$$

where inflation dynamics are allowed to differ across states 1 and 2:

$$\Delta CPI_t = \begin{cases} \nu_1 + A_{1,1}\Delta CPI_{t-1} + \cdots + A_{1,4}\Delta CPI_{t-4} + e_{1,t}, & \text{if} \quad s_t = 1 \\ \\ \nu_2 + A_{2,1}\Delta CPI_{t-1} + \cdots + A_{2,4}\Delta CPI_{t-4} + e_{2,t}, & \text{if} \quad s_t = 2. \end{cases} \tag{2.1}$$

We explain ΔCPI_t (seasonally adjusted CPI data in monthly log differences) by an intercept ν_m and autoregressive terms of four lags, which all switch between $m = \{1, 2\}$ states, just like the variance of the residual term $e_{m,t}$.[13] We choose a rather small number of regimes and lags to keep the model as parsimonious as possible and thus to increase the reliability of the estimates. In this way we also reduce computational cost significantly.

We estimate the model parameters and the hidden Markov chain with the expectation maximization (EM) algorithm.[14] We then obtain the filtered state probabilities $Pr(State_t)$, which we use for constructing the state indicator H_t (Chauvet and Hamilton, 2006). When the filtered probability of being in State 2 is greater than 0.5 in period t, H_t is assigned the value of 1, and 0 otherwise. Correspondingly, the indicator for being in State 1 is $1 - H_t$.[15]

2.2.2 State-dependent local projections

We follow the local projection instrumental variable (LP-IV) approach of Stock and Watson (2018) to construct the impulse responses. This method consists of a first-stage regression (2.2) in which the endogenous variable x_t is regressed on the instrument Z_t, and a second stage (2.3) that regresses the response variable y_t on the fitted values of the first

[13] Since we use monthly data, we also estimated an MS-AR including four lags plus the 12th lag. We did not observe significant differences in the timing of the resulting regimes. The identified regimes are generally not sensitive to the lag length.

[14] For further explanation of the EM algorithm, see Hamilton (1990).

[15] Our main results remain unchanged if we assign periods to State 2 if the filtered probability is above 0.4 or 0.7, where in the latter case we have to reduce the number of lags to 8, as we would otherwise end up with too few outliers in State 2, see below.

stage, $\hat{x}_t$, and a set of (lagged) control variables W_t:

$$x_t = \mu_1 + \beta_1 Z_t + \sum_{l=1}^{n} \delta_{1,l} W_{t-l} + \epsilon_t \tag{2.2}$$

$$y_{t+h} = \mu_{2,h} + \beta_{LPIV,h} \hat{x}_t + \sum_{l=1}^{n} \delta_{2,l} W_{t-l} + u_{t+h}. \tag{2.3}$$

The coefficients $\hat{\beta}_{LPIV,h}$ then represent the impulse responses at each projection horizon h. $\hat{\mu}_1$ and $\hat{\mu}_2$ denote the intercepts, ϵ_t and u_t the error terms.

Adding to this core model, we interact the fitted values $\hat{x}_t$ and the controls W_t with a state indicator H_t taking the value 0 in State 1, and 1 in State 2. Modifying the second-stage equation (2.3) in this way allows us to estimate state-dependent impulse response functions (IRFs):

$$\begin{aligned} y_{t+h} =& \mu_{2,h} + (1 - H_t)(\beta^1_{LPIV,h} \hat{x}_t + \sum_{l=1}^{n} \delta^1_{2,l} W_{t-l}) \\ &+ H_t(\beta^2_{LPIV,h} \hat{x}_t + \sum_{l=1}^{n} \delta^2_{2,l} W_{t-l}) + u_{t+h}. \end{aligned} \tag{2.4}$$

The coefficients $\hat{\beta}^1_{LPIV,h}$ and $\hat{\beta}^2_{LPIV,h}$ form the impulse responses at horizon h in states 1 and 2 respectively. Estimation of equation (2.4) is done via ordinary least squares regression for each projection horizon h separately.

The sample we use to estimate our baseline model (2.4) for the United States is in monthly frequency and spans from October 1948 to December 2021. The endogenous variable x_t is the log difference of the crude materials producer price index (referred to as Crude PPI) of the Bureau of Labor Statistics' stage-of-processing (SOP) system. For the response y_t in the baseline model, we use log differences of the US CPI. In alternative setups, we also employ the SOP-PPI data for intermediate materials, supplies, and components (Intermediate PPI), and finished goods (Finished PPI) or the industrial production SOP data for crude goods (Crude IP) as dependent variables. Appendix 2.7.1 provides more details on the PPI and IP data.

2.2.3 Shock identification

To identify the causal effect of a producer price shock on consumer price inflation, we identify the effects of unexpected and unusual price movements, filtering out smaller ups and downs over time. Given the relatively high frequency of our data set (monthly), this approach makes us more confident that we identify actual shocks. To do so, we introduce a new identification approach and argue that outliers in time series data, which are often due to rare and unforeseen events, are correlated with the exogenous shocks that we wish to identify.[16] Specifically, we instrument producer prices with a variable based

[16] Li et al. (2022) also follow a data-driven approach for shock identification as they identify shocks of Bitcoin and crude oil returns via the empirical quantiles of the two series. Kapetanios and Tzavalis

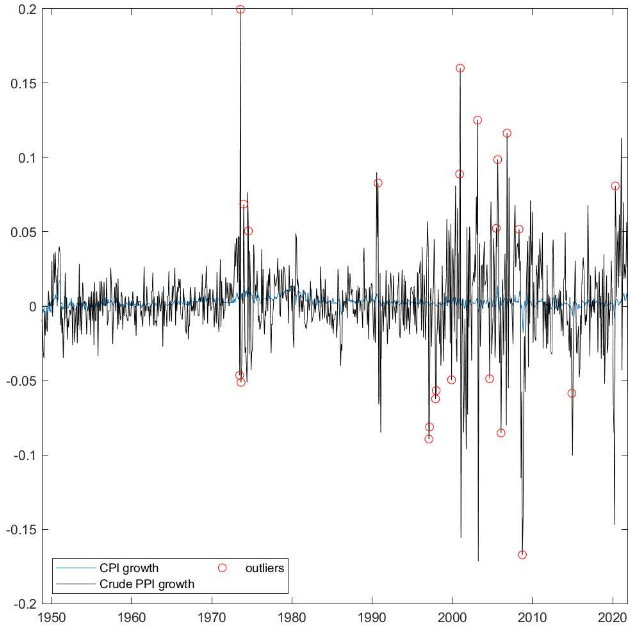

Figure 2.2: Crude PPI growth and outliers

Notes: Monthly growth rates (black) of the Crude PPI series, respectively, against monthly CPI growth (blue). Red circles mark the outliers generated with the iForest algorithm that survive the restriction described in equation (2.5).

on data outliers in the respective PPI series and assume that outliers in the PPI series are correlated with structural producer price shocks but uncorrelated with other shocks. The outlier-based instrument, hence, satisfies the LP-IV relevance and contemporaneous exogeneity condition of Stock and Watson (2018).[17]

To ensure that demand shocks are not the cause of the observed outliers, we only consider those outliers for which the materials industrial-production index IP^M from the board of governors does not move contemporaneously in the same direction as the Crude PPI.[18] That is, we construct the outlier-based instrument Z_t in the following way:

$$Z_t = \begin{cases} 1, & outlier > 0 \quad \& \quad \Delta IP^M < 0 \\ -1, & outlier < 0 \quad \& \quad \Delta IP^M \geq 0 \\ 0, & else. \end{cases} \tag{2.5}$$

Z_t takes the value of 1 when there is a positive outlier in the PPI series and no positive movement in the IP series in period t. In case of a negative outlier and no negative change in the corresponding IP series, $Z_t = -1$, and $Z_t = 0$ if no anomaly is detected. To

(2010) show that well-known oil price shock events coincide with periods in which they find an outlier in their oil price data.

[17]Those are: i) Z_t must be relevant, i.e., the shock of interest $\eta_{j,t}$ must be correlated with the instrument: $E[\eta_{j,t}Z_t] \neq 0$, ii) Z_t must be contemporaneously exogenous to all other shocks $\eta_{-j,t}$: $E[\eta_{-j,t}Z_t] = 0$ and iii), Z_t must be exogenous to all shocks at all leads and lags: $E[\eta_{t+i}Z_t] = 0, \forall i \neq 0$.

[18]We use this IP index as it corresponds closely to the Crude PPI and is available for our whole sample, starting in 1948.

ensure that Z_t satisfies the third LP-IV condition (exogeneity to all shocks at all leads and lags), we follow Stock and Watson (2018) and include 12 lags of Z_t, y_t, $\Delta \log IP_t^M$, and the growth of the log of the Intermediate PPI, as controls in regressions (2.2) and (2.4). Furthermore, we include lags of Z_t as controls to correct for a possible correlation between the instrument and past values of the shock of interest. By including lags of the materials IP series as a monthly proxy for activity, we correct for any correlation between Z_t and earlier developments.[19] Controlling for lags of CPI and Intermediate PPI growth rules out the possibility that the instrument Z_t is correlated with a shock to consumer prices or the producer prices of the previous stage. This, in addition to the restriction on ΔIP^M, further ensures that the dynamic effect we measure is not driven by a previous hike in demand leading to an increase in downstream prices first, followed by increasing upstream prices thereafter.

We detect outliers in the producer price indices using the isolation forest algorithm (iForest) proposed by Liu et al. (2012).[20] Instead of first defining normal instances in the data, the iForest directly detects anomalies through two quantitative properties: i) anomalies are the minority, and ii) they have attribute values different from those of normal instances. When setting the proportion of outliers in the PPI series (transformed to log differences) to 0.08, the iForest algorithm detects 71 outliers.[21] Figure 2.2 shows the Crude PPI series and the detected outliers at which the materials IP index does not move in the same direction. The outliers coincide with periods when there were prominent events on the supply side that led to large movements in Crude PPI inflation. Visible are the oil-price shock in 1973, which led to a spike in crude-material prices with a subsequent adjustment and re-escalation, the tensions surrounding the gulf war in 1990, the Asian crisis in 1997—which caused a decline of Asian commodities demand (exogenous to the US) and an appreciation of the dollar—OPEC production cuts in early 2001, supply adjustments and a dollar appreciation following the financial crisis in 2008, and supply chain disruptions in 2021.

These shocks occur in phases of low and high inflation volatility. That is, a single outlier does not necessarily move the economy to a high-volatility regime, which might happen for a series of large and/or more frequent smaller shocks. For example, as shown below, turbulent oil prices in the 1970s induced switches to a high-volatility regime, while the Asian crisis did not.

2.3 Empirical results

We now turn to the results of the baseline specification. We first describe the differences in the identified regimes, then the effects of shocks to producer prices on consumer prices and industrial production in these regimes. In sections 2.3.3.-2.3.5 we conduct robustness

[19] If Crude IP is the dependent variable, we correspondingly control for Crude IP.

[20] Specifically, we use the implementation in the Scikit-learn Python package by Pedregosa et al. (2011). For further explanations of the algorithm, see Liu et al. (2012).

[21] We choose 0.08 as lower values result in too few shocks and consequently weak instruments. Higher values might identify price movements that are not connected to clear supply shocks. We, therefore, prefer this rather conservative value. In any case, we also obtain state-dependent responses of inflation to shocks to Crude PPI for values of, e.g., 0.04 or 0.16.

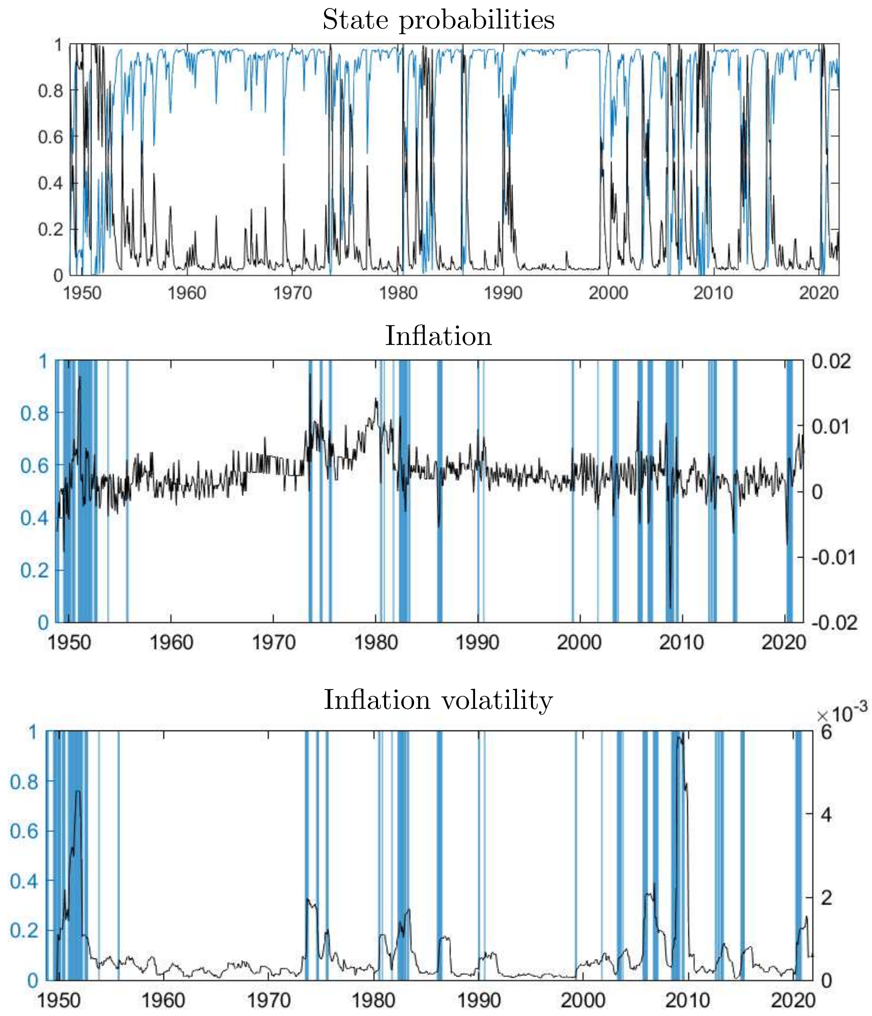

Figure 2.3: Identified regimes

Notes: Top panel: filtered state probabilities estimated from model (2.1); blue line: State 1, black line: State 2. Middle panel: monthly growth of CPI (black line); white areas: State 1, blue areas: State 2. Bottom panel: inflation volatility (black lines, variance of monthly CPI growth over a rolling window of 12 months); white areas: State 1, blue areas: State 2.

checks regarding alternative state-dependencies, based on either the level of inflation or the size of the shock. Alternatively, we rely on oil-supply shocks as identified by Baumeister and Hamilton (2019) to measure supply shocks. We also check whether the sign of the shock makes a difference, change the starting date, check the responses of the interest rate and other stages of production, and control for the exchange rate. The conclusion remains the same: it is the prevailing volatility that has a significant impact on the transmission of supply shocks.

2.3.1 Identified regimes

Figure 2.3 shows the filtered state probabilities, estimated with the methodology described in Section 2.2.1, and the resulting state indicator H_t in comparison with monthly growth rates of CPI and inflation volatility. We measure inflation volatility by the variance of monthly CPI growth over a rolling window of 12 months. As is visible, the inflation regime is in State 2 whenever there are sudden swings in monthly CPI growth and generally increased volatility. Specifically, the correlation between the state indicator and a volatility indicator variable vol_t—which takes the value 1 if the absolute change in the CPI is above its average and zero otherwise—is 30% and significant.[22]

Parameters	State 1	State 2
Probability to stay in regime	0.97	0.87
Avg. state duration in months	33	7.7
Std. dev. of monthly Δ CPI in %	0.27	0.56
Mean size outliers crude	0.06	0.07
CPI autocorrelation lag 1	0.75	0.48
CPI autocorrelation lag 2	0.65	0.21
Mean of monthly Δ CPI in %	0.34	0.24

Variables	β	p-value	Variables	β	p-value
constant	-0.30	0.00	vol_{t-5}	0.02	0.36
vol_t	0.46	0.00	vol_{t-6}	0.02	0.37
vol_{t-1}	0.30	0.00	vol_{t-7}	0.05	0.04
vol_{t-2}	0.13	0.00	vol_{t-8}	0.03	0.14
vol_{t-3}	0.12	0.00	vol_{t-9}	0.03	0.24
vol_{t-4}	0.08	0.00	vol_{t-10}	0.00	0.94
R^2	0.69		*Adj.* R^2	0.68	
Obs.	589				

Table 2.1: Regime characteristics and determinants

Notes: Upper panel: characteristics of the two regimes. All statistics in percent. Lower panel: regression of filtered state probabilities on exogenous volatility indicator and its lags, maximizing R^2.

In the upper panel of Table 2.1, we report descriptive statistics for the inflation regimes. The states are relatively persistent: The probability of staying in State 1 when being in the same state (i.e., p_{11}) is 0.97, and 0.87 for State 2 (p_{22}). This translates to an average state duration of 33 periods for State 1 and 7.7 periods for State 2. Comparing the standard deviation of monthly inflation within each state we find an average of 0.27 in State 1 and more than double (0.56) in State 2. This higher volatility is only to a very low degree driven by larger outliers in the SOP PPI data, as we find similar values for their mean values at the different stages of production across regimes. Instead, the regime-dependent autocorrelation of monthly CPI growth seems to contribute more to the state differences. We calculate this autocorrelation up to two lags, considering only those regime realizations that consist of at least three consecutive periods. In State 1, we find a value of 0.75 for the first lag and 0.65 for the second, in contrast to 0.48 and 0.21 for lag one and two in State 2. Interestingly, the overall mean of monthly CPI growth is 0.34% in State 1 and

[22] Using European micro data from 11 countries over the period 2005-19, Gautier et al. (2024) find an increased frequency of price setting at the end of the 2000s in the period during and after the financial crisis, in line with our theoretical interpretation of more flexible prices in State 2. Similarly, Dedola et al. (2023), making use of the same micro data, argue that recent evidence suggests that the return of higher and more volatile inflation seems to be associated with higher frequencies of price changes. Furthermore, Galeone and Gros (2023) find core inflation behavior to have shifted in the 2022/23 period as regards it magnitude, its rate of change, and its stickiness, as well as its responsiveness to energy prices.

only 0.24% in State 2. This highlights that not the overall level of inflation but rather its volatility characterizes the different inflation regimes.

We further demonstrate the regime dependence on inflation volatility by regressing the Markov filtered state probabilities $Pr(State_t)$ on the volatility indicator vol_t in the following way:

$$Pr(State_t) = c + \sum_{i=0}^{t=10} vol_{t-i}. \tag{2.6}$$

The contemporaneous indicator and the first four lags are significant at the 5% level.[23] Alternatively, we define the volatility indicator variable vol such that the R^2 of the mentioned regression is maximized, reaching 0.69, and find a threshold for the absolute value of the monthly change in CPI growth of 0.43 pp., or 5.28 pp. in annualized terms. That is, the optimized indicator variable takes the value of 1 if the absolute change in monthly inflation is above this threshold and zero otherwise. This value corresponds to approximately the 90th percentile of our sample; it was reached in, e.g., April 2022 (change in monthly inflation: -0.6 pp.), May 2022 (0.5 pp.), and July 2022 (-1.22 pp.). The correlation between the Markov state probabilities and this indicator is 0.65 and significant. The lower panel of Table 2.1 reports the resulting coefficients from repeating regression (2.6) with the optimized threshold. If the current monthly absolute change in CPI growth is above 0.43 pp., the likelihood of being in State 2 increases by 46 pp. (significant at the 1% level), ceteris paribus. The first four lags are also significant at the 1% level with decreasing coefficients.

Results are very similar if we include the contemporaneous values of the monthly VIX index, the growth rate of industrial production, and trend inflation (obtained by HP-filtering monthly inflation rates): the contemporaneous value and the first four lags of the volatility remain significant at the 1% level, while the adjusted R^2 increases to 0.71. The optimal threshold for the indicator is still 0.43 pp. of the change in CPI growth and the correlation of the Markov state probabilities with the indicator remains at 0.65. To sum up, if annualized monthly inflation changes by more than 5.2 pp., the inflation regime is likely to switch to State 2. Furthermore, the longer inflation is volatile, the higher the likelihood of reaching State 2.

2.3.2 Effects of supply shocks in different volatility regimes

The left panel of Figure 2.4 shows the state-dependent responses of monthly CPI to a unit shock to Crude PPI over a horizon of 12 months. We estimate regression (2.4) by setting y_t equal to the changes in the CPI and report the cumulated responses. They are significantly different from each other in states 1 and 2 over almost the whole horizon considered. Specifically, in State 2—the one associated with higher volatility in monthly CPI growth—CPI reacts faster and stronger, compared to State 1. That is, we find clear evidence for state dependency of the CPI response to supply shocks, where the transmission of producer price

[23] That is, observing a higher-than-average absolute change in the CPI increases the likelihood to be in State 2, resulting from the Markov-switching model, by 17 pp. If, additionally, the last four monthly absolute changes were also above average, the likelihood is 46 pp. higher.

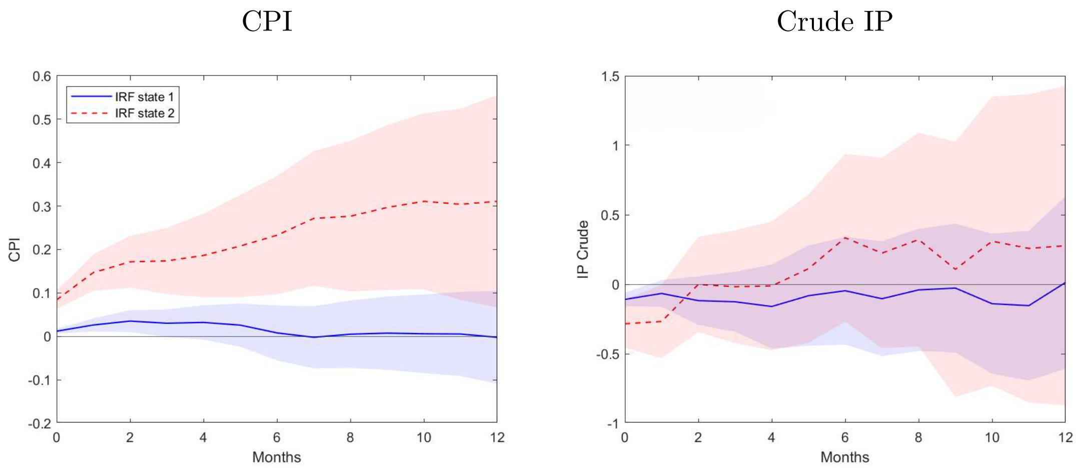

Figure 2.4: Baseline results

Notes: Impulse responses in Regime 1 (low volatility, solid blue lines) and Regime 2 (high volatility, dashed red lines) of CPI to a shock to Crude PPI (left) and corresponding industrial production response (right). Horizontal axes denote months. Shaded areas represent 68% confidence intervals.

shocks to consumer prices is stronger and quicker during a high-inflation-volatility regime than during times of more tranquil inflation.[24] The shaded areas represent 68% confidence bands. We construct them with Eicker-Huber-White (EHW) heteroskedasticity-robust standard errors as suggested by Montiel Olea and Plagborg-Møller (2021).[25]

In Appendix 2.7.2 we check several econometric issues, among them potentially weak instruments.[26] Furthermore, the left panel of Figure 2.7-1 in Appendix 2.7.3 demonstrates that the different CPI responses in the two regimes are not due to a more expansionary monetary policy reaction in State 2.[27] Lastly, the right panel of Figure 2.7-1 shows that the exchange rate appreciates more in the high-volatility regime, such that regime differences are not due to an exchange-rate depreciation that raises PPIs and the CPI alike.

We also calculate the effect of a shock to Crude PPI on industrial production of crude goods.[28] The right panel of Figure 2.4 depicts the results. As discussed in Section 2.2.3, to identify supply shocks we restrict industrial production to decrease in the period of a contractionary PPI shock. In the high-volatility regime, this effect is somewhat more pronounced, but the difference between regimes is much smaller compared to the CPI response and statistically not different from each other throughout.

2.3.3 Alternative regimes: inflation level and shock size

As stated in Section 2.2.1, the Markov-switching model indicated to separate regimes by their inflation volatility and not by the level of inflation itself. To further demonstrate

[24] Inflation remains mostly higher in State 2 up to a horizon of 23 months and falls thereafter.

[25] They show that when augmenting the local projection with lags of the response variable, EHW standard errors produce favorable results without the need to further correct for serial correlation in the regression residuals. In line with this argument, we include 12 lags of y_t in the local projection regressions.

[26] Using the test of Lewis and Mertens (2022), we show that none of our instruments is weak, see the left panel of Figure 2.7-1.

[27] Specifically, monetary policy reacts more strongly to a shock to Crude PPI in State 2 than in State 1, in line with the larger CPI response.

[28] Given the availability of the Crude IP series, we move the starting date to January 1967.

Above/below-average values of:

inflation level | inflation volatility

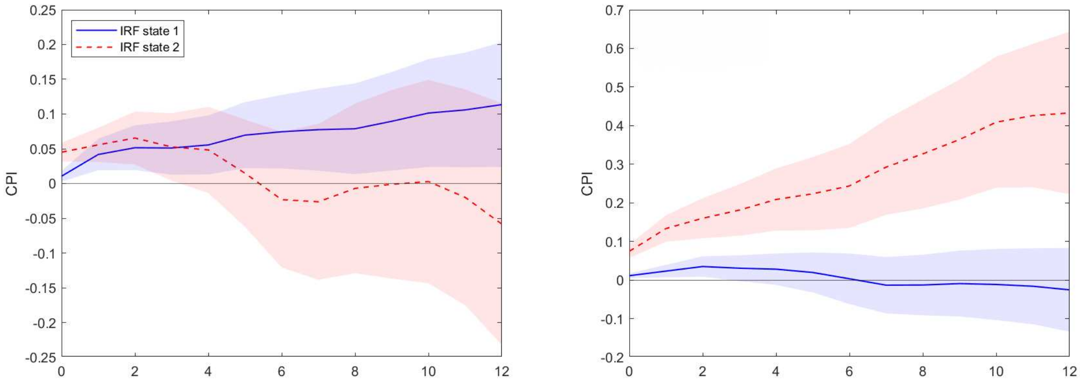

Figure 2.5: Inflation level vs. volatility

Notes: Impulse responses in Regime 1 (low state, solid blue lines) and Regime 2 (high state, dashed red lines) of CPI to a shock to Crude PPI. Left: State 1/2 if level CPI inflation is below/above average. Right: inflation volatility, calculated as in Figure 2.3, below/above average. Horizontal axes denote months. Shaded areas represent 68% confidence intervals.

that it is this dependency that causes impulse responses to differ across regimes, we no longer consider regimes as they were found by our Markov-switching model. Instead, we investigate whether alternative regime definitions based on the inflation level or the shock size also result in a state-dependent transmission of supply shocks. To this end, we split regimes such that we are in State 1 whenever inflation is below its average value and in State 2 if it is above the average. The left panel of Figure 2.5 shows the results for regimes below (blue solid lines) and above (red dashed lines) the average inflation level. No clear state dependency is visible. In particular, while the impact response of CPI inflation after a shock to Crude PPI is slightly higher than in State 2, it is below State 1 in the following periods. To corroborate this finding, we redo the same analysis for different cut-offs of the inflation level for defining states 1 and 2. Comparing the reactions to supply shocks in states where inflation is above or below its 65th, 70th, 80th, or 90th percentile shows very similar responses.[29]

We then verify that this approach yields state-dependent effects similar to our baseline results when we use inflation volatility, i.e., the change in CPI inflation, to exogenously separate regimes (rather than endogenously, as in our baseline). The right panel of Figure 2.5 shows the results. State 1 corresponds to a below-average inflation volatility, defined as in Figure 2.3. Blue solid lines depict the respective responses, while red dashed lines show the responses in State 2 (inflation volatility above its average). The state dependency is indeed similar to our baseline Figure 2.4, if not stronger. Supply shocks to Crude PPI are transmitted more quickly and strongly to consumer prices if inflation volatility is above average.

Next, we turn to the effects of the shock size. In standard menu cost models without observation costs (such as Golosov and Lucas 2007), price-setting behavior depends on

[29]Results are available upon request.

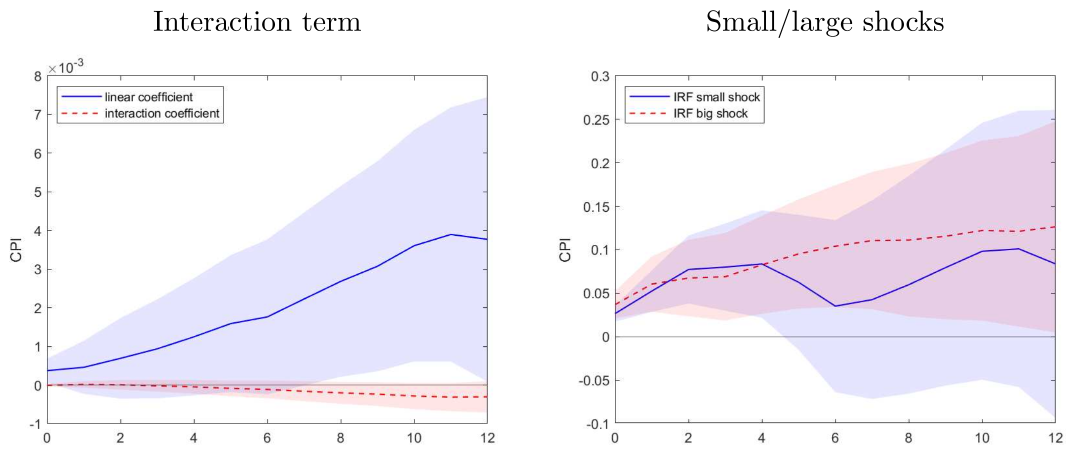

Figure 2.6: Effects of large vs. small shocks

Notes: Impulse responses of CPI to a shock to Crude PPI. Left: specification including linear (solid blue lines) and interaction term $|\hat{x}_t| \cdot \hat{x}_t$ (red dashed lines). Right: shock sizes below (solid blue lines) or above (dashed red lines) average. Horizontal axes denote months. Shaded areas represent 68% confidence intervals.

the size of contemporaneous shocks. A central result is that large input-price shocks have a relatively larger impact on consumer prices compared to smaller ones, see Ascari and Haber (2022). Given that periods of higher inflation volatility could be correlated to the average shock size in these periods, we check whether this correlation can explain the above findings. Figure 2.6 shows the reaction to small versus large shocks. We pursue two alternative strategies. In the left panel, we follow the approach of Ascari and Haber (2022) and include the term $|\hat{x}_t| \cdot \hat{x}_t$ in Model (2.3), in addition to the existing terms. That is, we measure the effect of the squared shock but conserve the sign of the shock. We do this independently of the regimes, as we are here interested in the effect of the shock size as an alternative explanation for our results. The effects of input-price shocks on the CPI via this interaction term and the linear coefficient are plotted by red dashed and blue solid lines, respectively. The interaction term is either insignificant or even negative, showing that large supply shocks do not automatically lead to a larger pass-through compared to smaller shocks.[30] If, however, several larger shocks (or a series of smaller shocks) result in higher CPI volatility, the shock transmission is profoundly altered, see above.

In the right panel of Figure 2.6, we conduct a similar exercise. Specifically, we separate the outliers, as identified in Section 2.2.3, depending on whether they are larger or smaller than the average. As in the previous exercise, we do not find a significant difference between the effects of relatively large vs. small shocks. That is, the influence of inflation volatility on the effect of supply shocks cannot be explained by the differential effects of the shock size.

[30]Given that Ascari and Haber (2022) consider the effects of monetary policy shocks instead of supply shocks, our results do not contradict their findings. For example, the effects of monetary policy decisions depend to a large degree on central bank communication and media coverage, influencing expectations, which might work quite differently depending on the size of the shock.

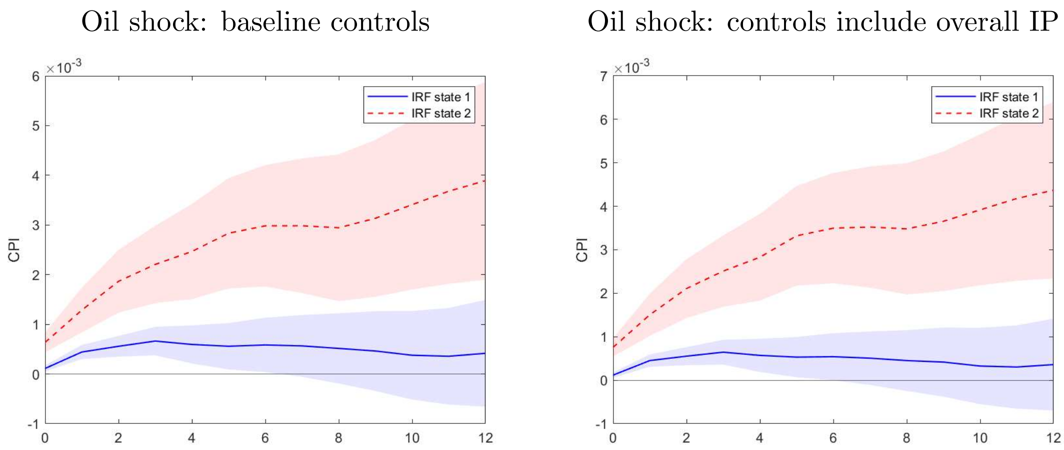

Figure 2.7: Effects of oil price shocks

Notes: Impulse responses in Regime 1 (low volatility, solid blue lines) and Regime 2 (high volatility, dashed red lines) of CPI to a contractionary oil price shock by Baumeister and Hamilton (2019). Left: baseline controls; right: controls with lags of overall IP. Horizontal axes denote months. Shaded areas represent 68% confidence intervals.

2.3.4 Alternative shocks: oil price shocks

We now turn to an alternative scheme for identifying supply shocks. Specifically, we exchange our identified shocks with oil-supply shocks, i.e., a series of supply shocks that are well established in the literature. We use the oil-supply shocks from Baumeister and Hamilton (2019), which range from February 1975 to December 2022. We again investigate possible differences in the CPI response in the two regimes identified in Section 2.3.1. The left panel of Figure 2.7 shows the results. In the right panel we include lags of overall industrial production as a control, another way to exclude demand shocks as the source of the responses. As is visible, the effects are similar to our more broad-based supply shocks of the baseline specification. Specifically, the effects of a supply shock are stronger on impact and thereafter in the high-volatility State 2.

2.3.5 Further Robustness

In this section we explore the robustness of our results with regard to the sign of the shock, different samples, regression setups, and identification schemes.

First, we analyze potential asymmetries between positive and negative shocks. We first create an instrument containing only the positive outliers and then a second one with only negative outliers. We estimate both directions of the shock at the same time to avoid potential biases by truncated variables (Garzon and Hierro, 2021):

$$y_{t+h} = \mu + \beta_h^+ \hat{x}_t^+ + \beta_h^- \hat{x}_t^- + \sum_{l=1}^{n} \delta_{2S,l,1}^T W_{t-l} + u_{t+h}. \tag{2.7}$$

In Model (2.7), $\hat{\beta}_h^+$ and $\hat{\beta}_h^-$ denote the positive and negative impulse responses, respectively. $\hat{x}_t^+$ and $\hat{x}_t^-$ are the fitted values from a regression of the dependent variable x_t (Crude PPI) on the positive or negative instrument and lagged controls W_t, which are the same as

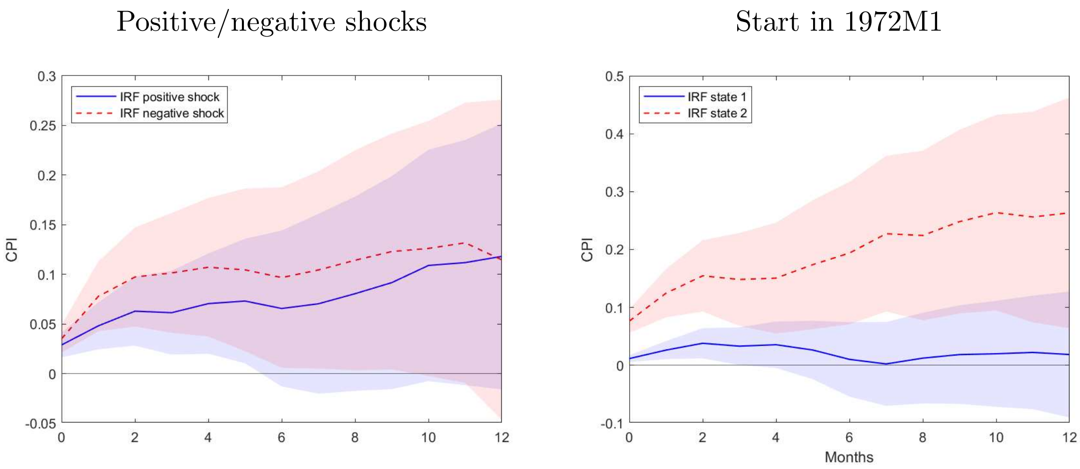

Figure 2.8: Robustness I

Notes: Left: Impulse responses to positive (solid blue lines) and negative (dashed red lines, positive values) of CPI to a shock to Crude PPI. Right: Impulse responses in Regime 1 (low volatility, solid blue lines) and Regime 2 (high volatility, dashed red lines) of CPI to a shock to Crude PPI, starting in 1972M1. Horizontal axes denote months. Shaded areas represent 68% confidence intervals.

employed in Model (2.4). The left panel of Figure 2.8 reports the resulting CPI responses to positive (solid blue lines) or negative (red dashed lines, positive values for ease of comparison) shocks to Crude PPI. The point estimates are fairly similar and confidence intervals overlap at all horizons. That is, the direction of the shock hardly changes the shape of the responses. An uneven distribution of positive versus negative shocks is, therefore, not responsible for the documented state dependency.

Second, we move the sample start to 1972, after the peg of the dollar to gold was cut and towards the end of regulated oil prices in the US. Results are shown in the right panel of Figure 2.8 and are similar to the baseline estimates. Third, we change the lag length to 8 lags. Figure 2.9 shows the cases for a starting date in 1948M10 (left panel) and 1972M10 (right panel). Again, results do not change much.

Fourth, we test different specifications of the local projections to further demonstrate that we do not pick up demand shocks in our analysis. In particular, we include the contemporaneous value of IP^M (in addition to its lags) in the regression. The left panel of Figure 2.10 displays the results, which are similar to the baseline. The right panel of Figure 2.10 depicts the case in which we do not only restrict IP^M to have the opposite sign as Crude PPI, but both IP^M and overall industrial production.[31] Again, results change only mildly.

Lastly, we explore the possibility that the identified regimes depend on the dependent variable. Specifically, as shown by Gonçalves et al. (2024), if a shock affects the response variable y_t, it could also alter the state indicator H_t, if this depends on y_t. This might affect the state-dependent LP estimands and thus generate a bias in the impulse response. Nonetheless, in our baseline we assume that a one-time unit shock will not induce an alternation of the states as the regimes we estimate a relatively high persistence of 33

[31] We also include overall IP in the controls in this specification.

Eight lags | Eight lags, start in 1972

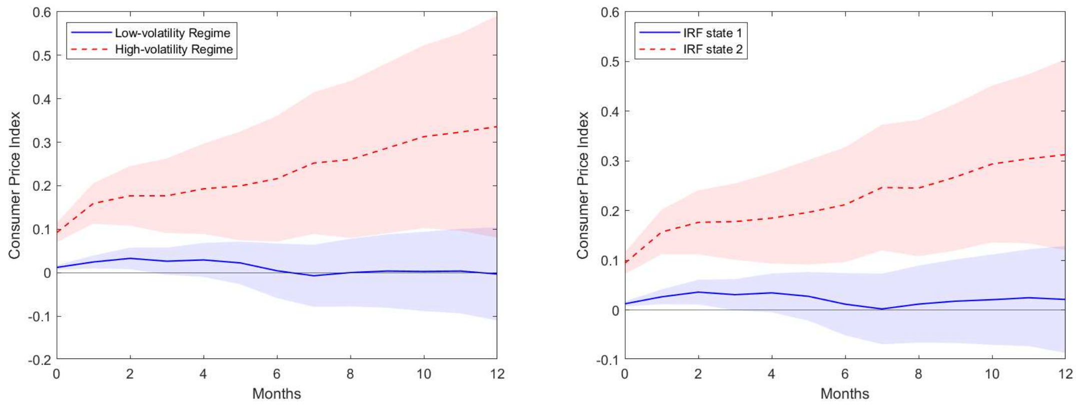

Figure 2.9: Robustness II

Notes: Impulse responses in Regime 1 (low volatility, solid blue lines) and Regime 2 (high volatility, dashed red lines) of CPI to a shock to Crude PPI. Left: 8 lags, start in 1948M10. Right: 8 lags, start in 1972M1. Horizontal axes denote months. Shaded areas represent 68% confidence intervals.

Contemp. IP^M in controls | Sign restrictions on IP^M and IP

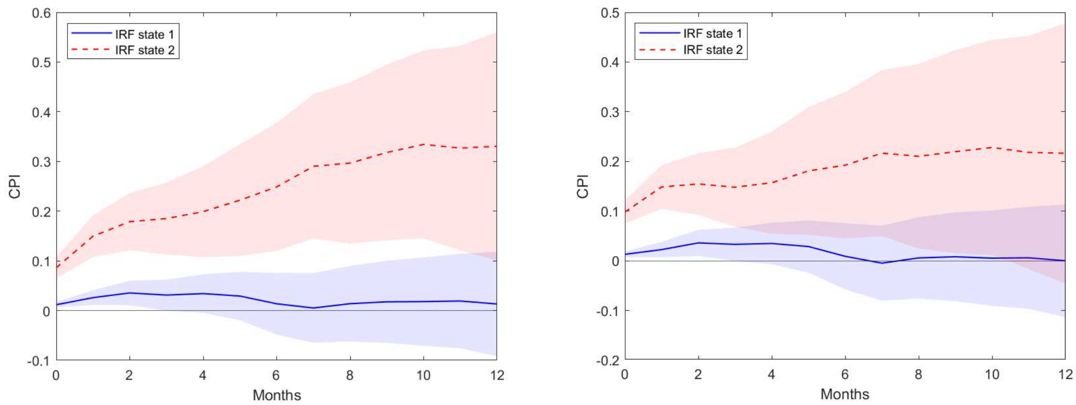

Figure 2.10: Robustness III

Notes: Impulse responses in Regime 1 (low volatility, solid blue lines) and Regime 2 (high volatility, dashed red lines) of CPI to a shock to Crude PPI. Left: contemporaneous value of IP^M included in controls. Right: sign restrictions on IP^M and overall IP employed. Horizontal axes denote months. Shaded areas represent 68% confidence intervals.

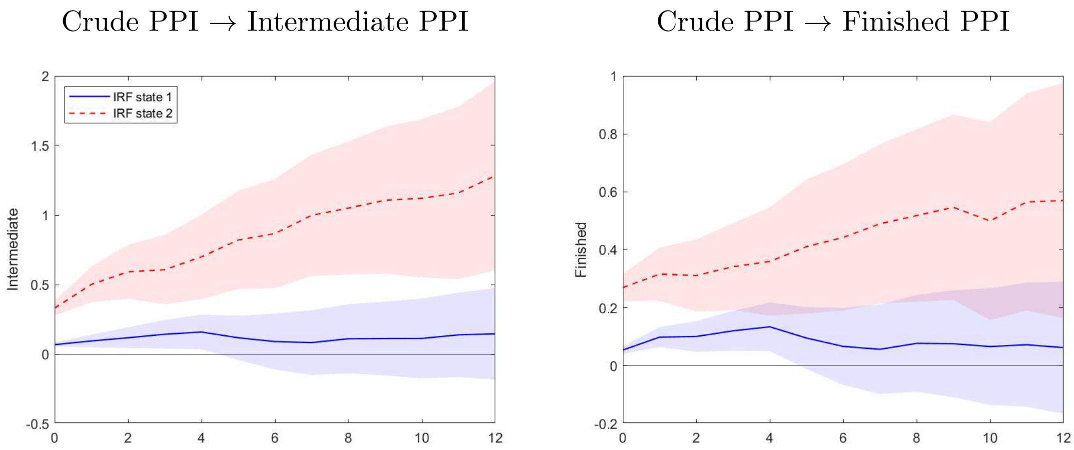

Figure 2.11: Effects of shocks to Crude PPI on intermediate stages
Notes: Impulse responses in Regime 1 (low volatility, solid blue lines) and Regime 2 (high volatility, dashed red lines) of Intermediate PPI (left panel) and Finished PPI (right panel) to a shock to Crude PPI. Horizontal axes denote months. Shaded areas represent 68% confidence intervals.

months in State 1 and almost 8 months in State 2.[32] In a robustness check, we follow Ramey and Zubairy (2018) and Cloyne et al. (2023) by lagging the indicator variable in regression (2.4). Results remain similar to our baseline. We also regress the state indicator variable on the contemporaneous and three lags of the fitted values of equation (2.2). None of the coefficients turns out to be significant.[33]

2.4 Effects on intermediate stages of processing

Next, we analyze the effect of a producer price shock on the prices of products located downstream in the stages of processing system. That is, we check how a shock to Crude PPI impacts Intermediate and Finished PPI by setting the response variable y_t in (2.4) equal to Intermediate (left panel of Figure 2.11) or Finished PPI (right panel). Note that neither PPI includes imports. We add the corresponding industrial production data in the controls, which moves, due to data availability, the starting date to 1972M1.[34] We leave the rest of Model (2.4) unchanged.

We again see a significantly differing response between states 1 and 2 on impact and in the following periods. The observation is quantitatively larger for Intermediate PPI than for Finished PPI, where, in turn, the effect is larger than for the CPI. This effect is as expected since at each stage of processing further inputs, such as labor, are added to the input materials.

[32] Furthermore, given that lags of inflation volatility are important in determining the volatility regime (see above), a one-time shock is not likely to induce a regime switch.

[33] All results are available upon request.

[34] We equate the industrial production index for primary & semifinished processing with Intermediate PPI and that of finished processing with Finished PPI.

2.5 Implications for theory

We now turn to potential theoretical explanations for our empirical finding of a stronger and quicker transmission of input prices to consumer prices in times of high inflation volatility. Our preferred theory assumes that firms are able to invest in price flexibility. In Section 2.5.1, we first rely on the mechanism developed by Devereux (2006) in a one-period model to derive analytical results. Observation costs as in Álvarez et al. (2018) or models of rational inattention (Mackowiak and Wiederholt, 2009) could also account for our evidence. The main intuition is the same across these models: depending on current observations, firms change their future price-setting behavior. Yet, Devereux's mechanism is much simpler while leading to very similar conclusions. It can also be seen as a direct implementation of price escalation clauses into a standard pricing model. Our version is kept deliberately simple since we aim to derive analytical results and to develop an intuition that could be used in several larger models.[35] We then move on to derive quantitative predictions in an infinite-period version, relying on the mechanism proposed by Kimura and Kurozumi (2010) and others. Here, firms invest in price flexibility by choosing the probability of being able to set prices in future periods, where higher flexibility is associated with larger costs.

We discard explanations based on menu costs or Calvo pricing with a fixed Calvo parameter for the following reasons. In standard menu cost models without observation costs (such as Golosov and Lucas 2007), price-setting behavior depends on the size of contemporaneous shocks. A central result is that large input-price shocks have a larger impact on consumer prices than smaller ones, see Ascari and Haber (2022). This prediction can be tested in our data, see Figure 2.6 for the reaction to small versus large shocks. As discussed in Section 2.3.3, we do not find a significant difference between the effects of large and small shocks.

Calvo pricing with a fixed Calvo parameter, on the other hand, would predict a constant impact of cost changes on inflation and is, therefore, clearly unable to replicate a state-dependent pass-through. In the following sections, we, therefore, explore a different class of models that can replicate our empirical findings.

2.5.1 Analytical model

We now sketch our preferred theory in a one-period model. We deviate from the original model in Devereux (2006) by introducing raw input material and a reaction function for the central bank—the model then features demand, supply, and monetary policy shocks—as well as simplifying the model by reducing it to a closed-economy setup and assuming pre-set wages. The following description of the model setup largely follows Devereux (2006), where more detailed derivations can be found. We introduce more significant changes to the original model in Section 2.5.1.2 and list the corresponding calculations in Appendix 2.7.4. Model predictions are derived in Section 2.5.1.3.

[35] See Khalil and Lewis (2024) for a quantitative version of the model in Devereux (2006) that includes endogenous entry and exit of firms.

2.5.1.1 Setup

Households maximize a utility function

$$U_t = \sum_{t=0}^{\infty} \log C_t - \frac{L_t^{1+\zeta}}{1+\zeta},$$

subject to the budget constraint $C_t P_t + B_t = (1+i_t)B_{t-1} + W_t L_t + C_{R,t} R_t + \Pi_t$, with $L_t = \int_0^1 L_{j,t}$; C_t is consumption, $L_{j,t}$ is hours worked at firm j, R_t is the aggregate input of raw materials, $C_{R,t}$ their price, Π_t are profits or losses (including price setting costs) from firms, and B_t are nominal bonds that pay $1+i_t$ in period $t+1$.[36] W_t is the wage, which is equal for all firms. Consumption bundles are composed of infinitely many varieties of goods:

$$C_t = \left(\int_0^1 C_{j,t}^{(\varepsilon-1)/\varepsilon} dj \right)^{\varepsilon/(\varepsilon-1)},$$

where $\varepsilon > 1$ is the elasticity of substitution between differentiated goods and market clearing implies $Y_{j,t} = C_{j,t} \forall j, t$. The aggregate price index is then

$$P_t = \left(\int_0^1 P_{j,t}^{1-\varepsilon} dj \right)^{\frac{1}{1-\varepsilon}}.$$

This setup gives rise to a standard demand function

$$Y_{j,t} = \left(\frac{P_{j,t}}{P_t} \right)^{-\varepsilon} Y,$$

with $P_{j,t}$ as the output price of firm j and P_t denotes the overall price level. Y_t represents total demand in the economy. As we will consider only one period in this version of the model, we drop time indexes in the remainder of this section.

Now consider firm j that produces according to

$$Y_j = (I_j - D_j \Phi(j))^{\alpha}, \tag{2.8}$$

where $I_j = R_j^{\gamma} L_j^{1-\gamma}$ represents firm j's usage of a combined input factor consisting of raw material R_j and employment L_j.[37] $\Phi(j)$ is a firm-specific cost of price flexibility. The parameter $0 < \alpha < 1$ measures the degree of decreasing returns to scale. The indicator variable D_j equals one if the firm chooses to have ex-post flexible prices in the period under consideration and zero if it decides to forego the opportunity of setting prices after observing this period's shock realizations. In our context, we interpret this cost as, e.g., using price-escalation clauses, which might require price discounts to clients and/or additional legal advice. Similarly, preserving price flexibility by using contracts that cover only short periods instead of fixing prices for longer may cause costs, such as lower negotiable output prices and more frequent contracting costs.

[36] As we model a closed economy, we assume that raw materials are available with unlimited supply at an fixed price $C_{R,t} = C_{RR,t} W_t$, where the relative price $C_{RR,t}$ of raw material to labor is exogenously given.

[37] We fix capital by fixing it at unity, as we are mainly interested in the short-term decisions of firms.

A related, but more complex, mechanism relies on 'observation costs,' proposed by Álvarez et al. (2018). In our model, $\Phi(j)$ would then be a shortcut to costs arising from a closer market observation. These costs would induce firms to monitor economic developments more thoroughly in times of higher volatility, while the model of Devereux (2006) relies on higher investments in price flexibility. Both models predict that current observed volatility raises the responsiveness of prices to future shocks, which will be crucial for accounting for our findings. That is, even large supply shocks transmit to consumer prices only to a low degree if they happen in tranquil times. This prediction differentiates these models from other approaches, such as menu cost models without observation costs, discussed above.

The price MC for one unit of the input factor I consists of the wage W, which is set in advance and is therefore fixed in this one-period model, and of the price of the raw material C_R. The latter is stochastic, and so are Y and P, as seen from the firm's perspective. As usual, minimized costs for one unit of I are then

$$MC = \frac{C_R^{\gamma} W^{1-\gamma}}{\gamma^{\gamma}(1-\gamma)^{(1-\gamma)}}. \tag{2.9}$$

We refer to unexpected movements in the costs of raw materials as supply shocks. Expected discounted profits of the firm are

$$E\Gamma\left[P_j\left(\frac{P_j}{P}\right)^{-\varepsilon} Y - MC\left(\left(\frac{P_j}{P}\right)Y\right)^{\frac{1}{\alpha}} - MCD_j\Phi(j)\right],$$

where E is the expectational operator and $\Gamma = 1/(PY)$ is the stochastic discount factor of the firm, corresponding to the marginal utility of one dollar of a hypothetical household with log utility. If the firm chooses to pay the (known, idiosyncratic) costs $\Phi(j)$, it can adjust its price after observing MC, Y, and P; otherwise, it sets its price based on expectations regarding these variables. The optimal price for firms that have chosen to invest in price flexibility is

$$P_j^1 = \delta\left[MC^{\alpha}(\hat{Y})^{1-\alpha}\right]^{\omega}, \tag{2.10}$$

where $\delta = \{\varepsilon/[\alpha(\varepsilon-1)]\}^{\alpha\omega}$ and $\omega = 1/[\alpha+\varepsilon(1-\alpha)]$. Furthermore, $\hat{Y} = P^{\varepsilon}Y$ is the part of a firm's demand that is independent of its price. Firms that chose to set their price in advance do this according to

$$P_j^0 = \delta\frac{E\left[\Gamma MC(\hat{Y})^{\frac{1}{\alpha}}\right]^{\alpha\omega}}{E\left[\Gamma\hat{Y}\right]^{\alpha\omega}}. \tag{2.11}$$

Expected profits under optimal price setting then depend on the choice to invest in price flexibility in the following way

$$\begin{aligned} V^1(\Theta) &= \Psi E\Gamma(MC^{\alpha(1-\varepsilon)}\hat{Y})^{\omega} \\ V^0(\Theta) &= \Psi(E\Gamma MC\hat{Y}^{1/\alpha})^{(1-\varepsilon)\alpha\omega}(E\Gamma\hat{Y})^{\varepsilon\omega}, \end{aligned}$$

where $V^1(\Theta)$ are profits for $D_j = 1$ and $V^0(\Theta)$ for $D_j = 0$. The parameter Ψ equals $\delta^{1-\varepsilon} - \delta^{-(\varepsilon/\alpha)}$ and $\Theta = \{C, Y, P\}$. The firm chooses ex-post price flexibility whenever the difference in expected profits for $D_j = 1$ and $D_j = 0$ is higher than the discounted costs of investing in price flexibility, i.e., if $V^1(\Theta) - V^0(\Theta) \geq \Phi(j) E\,\Gamma\, MC$, or

$$\Delta(\Theta) = \frac{V^1(\Theta) - V^0(\Theta)}{E\,\Gamma\, MC} \geq \Phi(j). \tag{2.12}$$

$\Delta(\Theta)$ is the discounted gain from investing in price flexibility, normalized by the cost of the combined input factor. This equation can be solved by taking a second-order approximation around the mean value $E \ln \Theta$, see Devereux (2006) for details:

$$\Delta(\Theta) \approx \frac{\Omega\alpha}{2} Var\left(\ln MC + \frac{1-\alpha}{\alpha} \ln \hat{Y}\right) = \frac{\Omega\alpha}{2}\left[\sigma^2_{mc} + \left(\frac{1-\alpha}{\alpha}\right)^2 \sigma^2_{\hat{y}} + 2\frac{1-\alpha}{\alpha}\sigma_{mc,\hat{y}}\right] > 0, \tag{2.13}$$

where lower-case letters stand for percentage deviations from the stochastic steady state, such as $mc = \ln MC - E \ln MC$. Furthermore, $\Omega = [V(\exp(E \ln \Theta))/\exp(E(\ln \Gamma + \ln MC))]\varepsilon(\varepsilon - 1)\omega^2 > 0$, where $V(\exp(E \ln \Theta))$ are profits evaluated at the mean $E \ln \Theta$ and $\sigma^2_c, \sigma^2_{\hat{y}}, \sigma_{mc,\hat{y}} > 0$ are the variances of input costs and market demand, as well as their covariance. Given expression (2.9), the cost variance σ^2_{mc} depends on the variance of (the log of) raw material costs in the following way: $\sigma^2_{mc} = \gamma^2 \sigma^2_{c_R}$. Equations (2.12) and (2.13) deliver an important insight in line with our empirical findings: higher volatility $\sigma^2_{\hat{y}}$ of market demand $\hat{Y} = P^{\varepsilon} Y$, which itself depends on price volatility, increases the incentives for firms to invest in price flexibility.

2.5.1.2 Closing the model

We now close the model, leading to several differences to Devereux (2006). Assume that there is a unit mass of firms. We then rank firms according to their cost of investing in price flexibility. The firm with the index $j = 0$ has the lowest costs $\Phi(0) = 0$ and the one with $j = 1$ the highest. We also assume that $\Phi(j)$ is uniformly distributed and differentiable. Denote the index of the firm that is indifferent to whether to invest in price flexibility or not as z. That is, z is the measure of firms that do invest. The resulting value of z is determined by the following conditions

$$\Delta(\Theta) = \Phi(z), \quad 0 \leq z < 1, \tag{2.14}$$

$$\Delta(\Theta) > \Phi(1), \quad z = 1. \tag{2.15}$$

The overall price index for a given value of z is then

$$P = \left[z(P^1)^{1-\varepsilon} + (1-z)(P^0)^{1-\varepsilon}\right]^{\frac{1}{1-\varepsilon}}. \tag{2.16}$$

Nominal demand is determined by the money supply in the following way

$$YP = \frac{M}{\chi}, \tag{2.17}$$

where χ features i.i.d. shocks to velocity and has an expected value of unity.[38] We refer to these shocks as the demand shock from now on. Inserting equation (2.17) into the optimal prices of firms (2.10) and (2.11), while observing that all firms that can adjust set the same prices, results in

$$P^1 = \delta \left[MC^\alpha P^{(1-\alpha)(\varepsilon-1)} (M\nu/\chi)^{1-\alpha} \right]^\omega \tag{2.18}$$

$$P^0 = \delta \frac{E\left[\Gamma MC \left(P^{\varepsilon-1}(M\nu/\chi)^{1-\alpha}\right)^{\frac{1}{\alpha}}\right]^{\alpha\omega}}{E\left[P^{\varepsilon-1}\right]^{\alpha\omega}}. \tag{2.19}$$

The central bank sets the change in the nominal money supply based on current inflation:

$$\frac{M}{M_{-1}} = \left(\frac{P}{P_{-1}}\right)^{-\phi} \nu, \tag{2.20}$$

where we normalize the previous period's values of the money stock and the price level to unity $M_{-1} = P_{-1} = 1$. We assume that the central bank does not react to higher inflation by increasing the money supply overproportionally, i.e., $\phi \geq -1$. Stricter inflation targeting corresponds to a higher value of ϕ. The variable ν with an expected value of unity may stand for monetary policy shocks, but also for systematic deviations from a rule that focuses on inflation only. In particular, we allow for a positive correlation between ν and the supply shock, which represents a monetary policy strategy that is relatively more accommodating in case of supply shocks.[39] Theoretically, ν could also be linked to demand shocks. Given the debate in some policy circles surrounding lower reactions to inflation in case of supply shocks, we focus on a correlation with this kind of shock.[40]

To derive the expression for equation (2.13) in general equilibrium, we use the linearized price index (2.18) together with the linearized versions of equations (2.16) and (2.20), see Appendix 2.7.4. This yields

$$p = \frac{\varphi(z)\omega}{\Delta} \left[\alpha mc + (1-\alpha)(\hat{\nu} - \hat{\chi})\right], \tag{2.21}$$

with

$$\Delta = 1 - \varphi(z)\omega(1-\alpha)(\varepsilon - \phi - 1),$$

where $\hat{\chi} = \ln\chi - E\ln\chi$ and $\hat{\nu} = \ln\nu - E\ln\nu$. The parameter $\varphi(z)$ is given in the appendix and follows $\varphi(0) = 0, \varphi(1) = 1, \varphi'(z) > 0, \varphi''(z) > 0$. Using equation (2.21) we derive—again in the appendix—the variance of $\ln MC + \frac{1-\alpha}{\alpha}\ln\hat{Y}$ and use this in equation (2.13)

[38] These shocks can be derived from shocks to households' preference for holding money, see Devereux (2006).

[39] The functional form would be $\nu = (MC/MC_{-1})^{-\phi_{mc}}\tilde{\nu}$, with $\tilde{\nu}$ being 'pure' monetary policy shocks.

[40] See, e.g., Fabio Panetta, member of the executive board of the ECB, who stated: "Bad inflation reflects negative supply shocks that raise prices and depress economic activity, which monetary policy should look through." (Panetta, 2022)

to arrive at equations (2.14) and (2.15) in general equilibrium as

$$\frac{\Omega\alpha}{2\Delta^2}\left[\sigma_{mc}^2+\left(\frac{1-\alpha}{\alpha}\right)^2(\sigma_{\hat{\chi}}^2+\sigma_{\hat{\nu}}^2)+2\frac{1-\alpha}{\alpha}\sigma_{mc,\hat{\nu}}\right]=\Phi(z) \qquad 0\leq z<1 \tag{2.22}$$

$$\frac{\Omega\alpha}{2\Delta^2}\left[\sigma_{mc}^2+\left(\frac{1-\alpha}{\alpha}\right)^2(\sigma_{\hat{\chi}}^2+\sigma_{\hat{\nu}}^2)+2\frac{1-\alpha}{\alpha}\sigma_{mc,\hat{\nu}}\right]>\Phi(1) \qquad z=1, \tag{2.23}$$

The covariance $\sigma_{mc,\hat{\nu}}$ corresponds to $-\phi_{mc}\sigma_{mc}$, see footnote 39.

2.5.1.3 Model predictions

Equations (2.22) and (2.23) then determine the equilibrium value of z, depending on the variances and covariances of the three shocks. As shown in the appendix, there can be one or three equilibria. However, in case of multiple equilibria, one is unstable. In the following, we focus on the description of the stable equilibrium in which the economy is not already at full price flexibility (i.e., $z<1$).[41] We first assert the relation between price flexibility and the pass-through of shocks to inflation. Given that the derivative of the term $\varphi(z)\omega/\Delta$ in the expression for the price index (2.21) with respect to z is positive, we directly obtain the following lemma.

Lemma 1 (Effect of price flexibility)**.** *A higher price flexibility (a higher z) translates into a larger pass-through of shocks to inflation.*

The following proposition then follows from equation (2.22).[42]

Proposition 1 (Effects of shock volatilities)**.** *Higher volatility of the shocks to the costs of raw materials ($\sigma_{c_R}^2$), demand ($\sigma_{\hat{\chi}}^2$), and/or the money supply ($\sigma_{\hat{\nu}}^2$, for a given covariance with input costs) raises price flexibility (z) and hence the pass-through of shocks to inflation.*

We also obtain the following corollary, which is linked to our empirical findings.

Corollary 1 (Relation to inflation volatility)**.** *Any change in the shock volatilities $\sigma_{c_R}^2, \sigma_{\hat{\chi}}^2$ and/or monetary policy variables ($\sigma_{\hat{\nu}}^2, \sigma_{mc,\hat{\nu}}$, and ϕ) that increases inflation volatility raises price flexibility and hence the pass-through of all shocks to inflation.*

Intuitively, higher variances of costs and/or demand make the possibility of a price adjustment after observing shock realizations more valuable (Proposition 1). This effect also works via the level of inflation volatility: If the prices of competitors are fluctuating strongly, it pays off to invest in the ability to change prices after observing the resulting demand. Higher price flexibility, in turn, increases the response of inflation to shocks. This aligns with our empirical result: higher inflation volatility leads to a larger pass-through of cost shocks to inflation (Corollary 1).

[41] If all firms have already invested in price flexibility, changes in parameter values can reduce price flexibility but can obviously not increase it any further.

[42] Proofs for the propositions and the corollary are given in Appendix 2.7.4.

Technically, equation (2.21) implies a larger shock pass-through if more firms have invested in price flexibility (how many firms are able to adjust their price after observing the shocks) and if monetary policy is less aggressive in fighting inflation (by how much do the adjusters adjust). The latter, direct effect of monetary policy on demand is standard in the literature. In particular, a higher value of ϕ raises Δ and corresponds to stricter inflation targeting. In the extreme, ϕ approaches infinity, which fixes the price level at its previous level. Additionally, the impact of ϕ on the variances of the price level and hence demand changes the firms' incentives to invest in price flexibility (see above), which entails an indirect influence of monetary policy via $\varphi(z)$. Regarding the effects of monetary policy, we can derive the following result.

Proposition 2 (Effects of monetary policy)**.** *Stricter inflation targeting (a higher ϕ) reduces the response of inflation to all shocks in two ways: directly by reacting to the change in inflation and indirectly by reducing price flexibility. In contrast, an accommodating monetary policy stance towards supply shocks (raising $Cov(mc, \hat{\nu})$) increases price flexibility (z) and thereby the pass-through of all shocks to inflation.*

Regarding the last part of the proposition, note that contractionary supply shocks increase costs and the general price level simultaneously. Seen from the perspective of an individual firm under strategic complementarity in pricing, both developments create an incentive to raise prices.[43] Similar reasoning applies to expansionary demand shocks, which increase demand and the price level. Firms are thus more likely to invest in price flexibility if the correlation of shocks with the price level is high. By dampening the price response, monetary policy can reduce this incentive.[44] A more accommodating policy, overall or just in case of supply shocks, counteracts this reasoning and leads—ceteris paribus—to a higher price flexibility and therefore a higher pass-through of shocks to inflation.

Despite this clear result, two caveats are in order. First, one argument for a muted monetary policy reaction to supply shocks is their transitory nature in combination with lags in the transmission of policy actions. Given that we consider a quite stylized model, we do not capture this notion here. Second, we are only interested in the connection between shocks and inflation and, hence, do not conduct a proper welfare analysis.

2.5.2 Dynamic model

To obtain quantitative predictions beyond those of the analytical one-period version above, we now move on to a numerical simulation of the infinite-period version. Here, we follow Kimura and Kurozumi (2010), which is based on concepts from Devereux and Yetman (2002), and let firms choose their individual degree of price flexibility (their Calvo parameter θ_j) once, given the parameters and shock variances. That is, they can set the probability of being able to adjust prices. As above, they pay the costs Φ whenever firms get the opportunity to do so, such that higher flexibility entails larger costs.

[43]Strategic complementarity is the standard case in this kind of model and is given by assuming $\alpha < 1$.
[44]Naturally, lower volatility achieved by reducing monetary policy shocks has the same effect.

Specifically, we introduce the above structure of raw material inputs to production into the New Keynesian framework of Kimura and Kurozumi (2010) and use the resulting model to analyze different inflation regimes. Thus, we retain the setup of the analytical model in Section 2.5.1 but assume an infinite planning horizon and allow the wage to be set in each period. For simplicity, we assume constant returns to scale, $\alpha = 1$, and constant costs of being able to adjust prices, $\Phi(j) = \Phi$.[45]

Furthermore, for ease of notation, we define this cost, expressed in prices of the aggregate output good, as $F_t \equiv \Phi MC_t/P_t$.

The firm's profit maximization is equivalent to minimizing its loss in profit from not being able to reset its price. Up to second order, this loss is proportional to (see Walsh, 2003)

$$\mathcal{L}_t(\theta_t, \theta) = F_t + \min_{p_{j,t}} E_t \sum_{k=0}^{\infty} (\beta\theta_j)^k (p_{j,t} - p^*_{j,t+k})^2 + \beta(1-\theta_j) \sum_{k+1}^{\infty} (\beta\theta_j)^{k-1} E_t \mathcal{L}_{t+k}(\theta_j, \theta),$$

where lower-case letters refer to variables linearized around the flexible-price steady-state, β is the firms' discount factor, and $p^*_{j,t}$ is the price the firm would set if no nominal rigidities were present, which is $p^*_{j,t} = mc_t = \gamma c_{R,t} + (1-\gamma)w_t$. The wage is determined from households' optimization problem and is given by

$$w_t - p_t = \sigma c_t + \zeta l_t = \xi y_t + \zeta\gamma c_{RR,t},$$

where $c_{RR,t}$ is the relative price of raw materials to labor $c_{R,t} - w_t$ and $\xi = \sigma + \zeta/[1 + (1-\theta)\Phi/Y)]$, with Y denoting output in steady state.[46] The desired price is therefore equal to

$$p^*_{j,t} = p_t + \gamma c_{RR,t} + w_t = p_t + \xi x_t, \tag{2.24}$$

with x_t as the output gap. The optimal price that results from this minimization is

$$p^0_{j,t} = (1-\beta\theta_j) E_t \sum_{k=0}^{\infty} (\beta\theta_j)^k p^*_{j,t+k} = (1-\beta\theta_j) E_t \sum_{k+0}^{\infty} (\beta\theta_j)^k (p_t + \xi x_{t+k}). \tag{2.25}$$

Following Kimura and Kurozumi (2010), we assume that the firm chooses its individual Calvo parameter θ_j to minimize the unconditional expected loss in profit due to rigid prices, which is

$$E\mathcal{L}_t(\theta_j, \theta) = \frac{1-\beta\theta_j}{1-\beta} \left[F + E \sum_{k=0}^{\infty} (\beta\theta_j)^k (p^0_{j,t} - p^*_{j,t+k})^2 \right],$$

where E is the unconditional expectations operator and F the unconditional expectation of F_t. That is, firms may decide for higher price flexibility (a lower θ_j) if they reckon that it pays off to be able to respond quickly to changing conditions. This is associated with higher costs, as they have to pay the price-setting costs F more often in this case. The

[45] That is, fluctuations in aggregate demand affect costs via the wage rather than through decreasing returns to scale, with similar implications.

[46] This expression is derived from the linearization of aggregate labor demand, given the production function (2.8), in which D_j equals unity if firm j can set its price and $\alpha = 1$.

first-order condition is then, using (2.24) and (2.25),

$$F + \sum_{k=0}^{\infty} (\beta\theta_j)^{k-1}[(k+1)\beta\theta_j - k]V\left[\sum_{h=1}^{k} \pi_{t+h} + \xi x_t - \tilde{\mathcal{L}}_t(\theta_j, \theta)\right] = 0, \tag{2.26}$$

where V is the unconditional variance and

$$\tilde{\mathcal{L}}_t(\theta_j, \theta) = \sum_{h=1}^{\infty} (\beta\theta_j)^h E_t \pi_{t+h} + \gamma(1-\beta\theta_j) \sum_{h=0}^{\infty} (\beta\theta_j)^h E_t x_{t+h}.$$

We reach an equilibrium if the optimal $\theta_j = \theta$ for each firm j, which yields the standard New Keynesian Phillips Curve

$$\pi_t = \beta E_t \pi_{t+1} + \frac{\gamma(1-\theta)(1-\theta\beta)}{\theta} x_t.$$

On the demand side, household optimization results in the dynamic IS equation

$$x_t = E_t x_{t+1} - (i_t - E_t \pi_{t+1} - r_t^*)/\sigma,$$

where r_t^* is the natural rate of interest, which is given by

$$r_t^n = -\frac{\sigma\gamma(1+\zeta)}{\xi} E_t \Delta c_{RR,t+1}.$$

Lastly, we assume a Taylor rule for the interest-rate decisions of the central bank

$$i_t = \phi_\pi \pi_t + \phi_x x_t.$$

To obtain a numerical solution, we search for a θ that, once the model is solved for this value and the equilibrium paths are inserted into (2.26), fulfills this equation for $\theta_j = \theta$.

2.5.2.1 Calibration and model predictions

The calibration of the model equals that of Kimura and Kurozumi (2010), where applicable. That is, we set $\beta = 0.99, \sigma = 1.86, \zeta = 1, \phi_\pi = 1.5, \phi_y = 0.5$ (on an annual basis), and $\rho^n = 0.83$, assuming an AR(1) process for $c_{RR,t}$ and hence r_t^n. Instead of employing their assumed variance of the natural-rate shock across both volatility regimes, however, we set this variance differently in each regime. In particular, we choose values such that the model generates the observed standard deviation of CPI inflation in each regime (0.27% and 0.56%, respectively). The resulting standard deviation of innovations to r_t^n is 0.2% in the low-volatility regime and 0.4% in the high-volatility regime. We then simulate a shock to the real price of raw materials that raises the nominal costs of raw materials by 1%, as in our empirical estimations.

Figure 2.12 displays the response of CPI inflation after such a shock in the two regimes. The blue solid line in both panels represents the low-volatility scenario with a resulting Calvo-parameter of 0.7, while the red dashed line shows the high-volatility case with an endogenous Calvo-parameter of 0.53. Considering the stylized nature of the—three-equation New Keynesian—model, we deem the fit to the corresponding responses in Figure 2.4 as

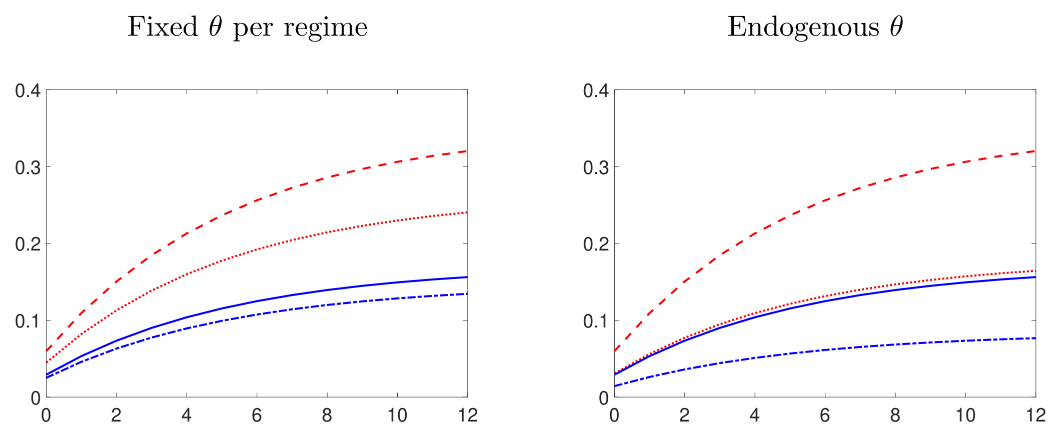

Figure 2.12: Theoretical responses with counterfactuals
Notes: Both panels: Impulse responses in Regime 1 (low volatility, blue solid lines) and Regime 2 (high volatility, red dashed lines) of CPI to a unit shock to the price of crude material $c_{R,t}$. Left: Responses in Regime 1 (blue dashed-dotted line) and Regime 2 (red dotted line) for stricter inflation targeting with unchanged, regime-specific price-setting frequency. Right: Responses in Regime 1 (blue dashed-dotted line) and Regime 2 (red dotted line) for stricter inflation targeting with endogenous, regime-specific price-setting frequency. Horizontal axes denote months.

a success. In particular and in line with Proposition 1, we obtain a higher inflation on impact and in the following periods in the high-volatility regime, induced by the higher share of price-adjusting firms. We also conduct two hypothetical scenarios in which the central bank adheres to stricter inflation targeting by increasing its reaction coefficient ϕ from 1.5 to 2. The left panel of Figure 2.12 shows the responses if we leave the Calvo-parameter unchanged for each regime, i.e., at 0.7 and 0.53, respectively. We thereby isolate the traditional monetary-policy channel that reduces inflation by dampening demand.

While the stronger reaction already achieves a lower inflation response to the cost-push shock for given values of θ, the effect is magnified once we allow for an endogenous re-adjustment of the price-setting frequency, as also discussed by Kimura and Kurozumi (2010). The right panel displays the corresponding responses that result from optimally chosen Calvo-parameters, based on the shock variances in both regimes (which are unchanged) and the new value for ϕ. Specifically, we obtain values of $\theta = 0.79$ in the low-volatility regime and $\theta = 0.67$ in the high-volatility regime. By comparing both panels, we find that the dampening of the inflation response is particularly successful in the high-volatility regime. In particular, the response in the high-volatility regime for $\phi = 2$ is similar to that in the low-volatility regime for $\phi = 1.5$. In short, stricter inflation targeting pays off double in terms of reducing inflation fluctuations, as predicted by Proposition 2.

2.6 Conclusion

We examine the impact of producer price shocks on consumer price inflation in the United States, taking into account different inflation regimes. Employing a Markov-switching model, we identify two distinct regimes and use the filtered state probabilities to construct a regime indicator. It turns out that the regimes are characterized by different inflation volatilities.We then interact a local projections model with the indicator and estimate responses with Stock and Watson (2018)'s LP-IV approach, using data outliers in the Crude PPI series as instruments.

We find that the impulse responses of the CPI following a producer price shock are indeed regime-dependent. If a producer price shock occurs during the high volatility regime, the increase in consumer prices is more pronounced on impact and takes longer to decay than in times of stable and low inflation. This distinction is not observable when considering different levels of inflation or shock sizes.

The main policy implication we draw from our results for inflation-targeting central banks is that they should pay close attention to the current and potential future inflation regimes when assessing the impact of current developments. If these developments lead to high CPI volatility, the economy may transition to a regime where cost shocks are passed on to consumer prices more rapidly and to a larger extent. This could result in persistently higher CPI inflation volatility. Put differently, a stricter monetary policy stabilizes inflation not only directly, but also indirectly by reducing price flexibility.

2.7 Appendix

2.7.1 Data description

Seasonally adjusted data on the CPI and the three producer price indices were obtained from the US Bureau of Labor Statistics (BLS). Until 2014, the BLS used the stage of processing (SOP) aggregation system to report producer prices. Afterward, the BLS switched to the Final Demand-Intermediate Demand (FD-ID) system. Table 2.7-1 reports the SOP and the corresponding FD-ID codes as well as the respective variable names.

The BLS defines crude materials as unprocessed goods and intermediate materials as processed goods that businesses purchase as inputs for their production. Products included in the Crude PPI enter the market for the first time and will undergo processing when purchased. Intermediate materials are already processed to some degree but need further processing before becoming a finished good. Finished goods comprise commodities used for personal consumption or that businesses use as capital investment. Government purchases and exports are excluded from the SOP system.

Seasonally adjusted data on the stages of processing industrial production indices and overall industrial production were retrieved from the Federal Reserve Board (FRB). The indices are classified into raw materials, primary & semifinished processing, and finished processing, and are available since 1972, or 1947 in the case of IP Materials.

2.7.2 Econometric checks

Our instrumental variable consists of few non-zero data points and can thus be characterized as a *sparse instrument*. Giacomini et al. (2022) argue that sparse instruments, often constructed from narrative restrictions, are likely to be weak instruments. We test the relevance of our IV by applying the robust test for weak instruments with multiple endogenous regressors proposed by Lewis and Mertens (2022). We interact the instrument and Crude PPI (our endogenous regressor) with the state indicator H_t and use the same set of controls as in our respective local projection specifications. Following Lewis and Mertens (2022), the test rejects weak instruments if the test statistic lies above the critical value. For our baseline specification, this is the case at all horizons and for all three stages of processing PPIs, as can be seen in Figure 2.7-1.

SOP Code	Title	FD-ID Code	Title
SOP1000	Crude materials	ID62	Unprocessed goods for intermediate demand
SOP2000	Intermediate materials, supplies and components	ID61	Processed goods for intermediate demand
SOP3000	Finished goods	FD49207	Finished goods

Table 2.7-1: Variable description of Crude (SOP1000), Intermediate (SOP2000), and Finished (SOP3000) PPI

Notes: More information available on https://www.bls.gov/ppi/fd-id/ppi-stage-of-processing-to-final-demand-intermediate-demand-aggregation-system-index-concordance-table.htm.

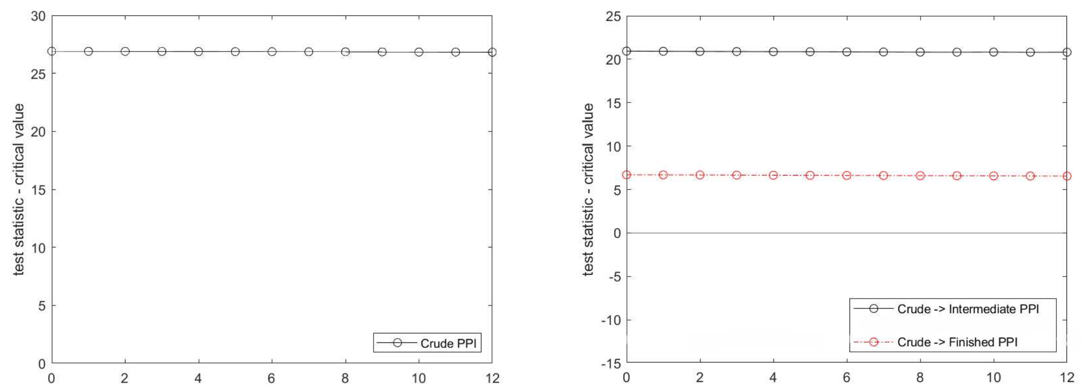

Figure 2.7-1: Weak instrument tests
Notes: Left panel: Results of the Lewis and Mertens (2022)-test for weak instruments: difference of test statistic and critical value for baseline results (Figure 2.4). Right panel: same statistic between stages of processing (Figure 2.11). Horizontal axes denote months.

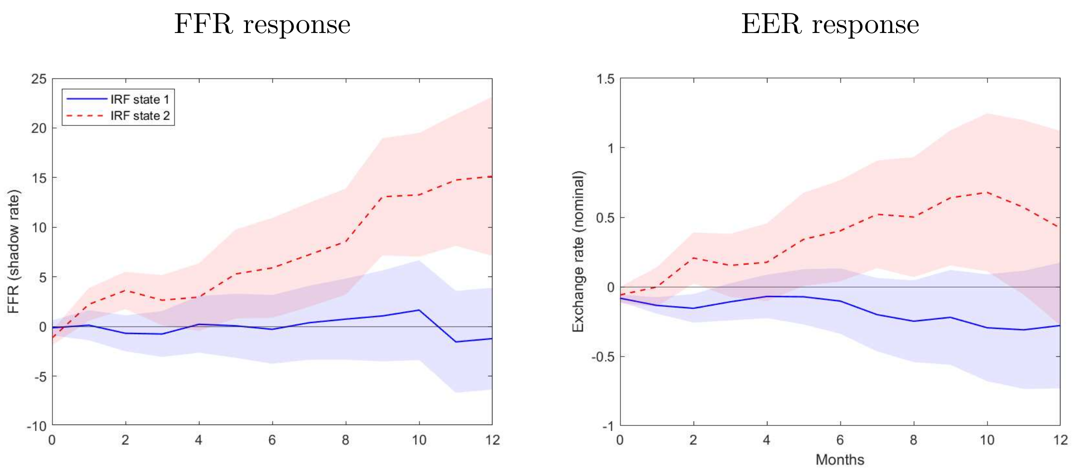

Figure 2.7-1: Impulse responses of shadow rate and nominal effective exchange rate
Notes: Right panel: higher values correspond to an appreciation, in Regime 1 (low volatility, solid blue lines) and Regime 2 (high volatility, dashed red lines) to shocks to PPI. Horizontal axes denote months. Shaded areas represent 68% confidence intervals.

2.7.3 Alternative channels

The left panel of Figure 2.7-1 shows the response of the shadow rate—the updated series from Krippner (2013)—to shocks to Crude PPI. As visible, the monetary policy reaction is not responsible for the observed state dependency of CPI responses. Monetary policy reacts more to shocks to Crude PPI in State 2 than in State 1, in line with the stronger inflation response. That is, if anything, monetary policy dampens the further course of inflation.

The right panel shows the response of the nominal broad effective exchange rate (EER), provided by the BIS. We reduced the lag number to 8 since the exchange rate is only available from 1994 onward. The exchange rate appreciates more in the high-volatility regime, in line with the stronger interest-rate response. That is, the stronger inflation reaction in the high-volatility regime cannot be explained by a depreciation that leads to rising PPIs at all stages of production and the CPI. Similarly, the responses of the Intermediate PPI and the Finished PPI in Figure 2.11, which exclude imports, further demonstrate that our results are not driven by the exchange-rate response.

2.7.4 Model derivations and proofs

Derivation of equation (2.21)**.**

The linearized price index is, see Devereux (2006),

$$p = \varphi(z)p^1,$$

with

$$\varphi(z) = \frac{z\exp(E\ln P^1(1-\varepsilon))}{z\exp(E\ln P^1(1-\varepsilon)) + (1-z)\exp(E\ln P^0(1-\varepsilon))}.$$

The linearized price (2.18) of flexible firms p^1 reads as

$$p^1 = \alpha\omega mc + (1-\alpha)(\varepsilon-\phi-1)\omega p + (1-\alpha)\omega(\hat{\nu}-\hat{\chi})$$

such that (2.21) results.

Derivation of equations (2.22) **and** (2.23)**.**

Given the expression (2.21) for the price index, we obtain $\hat{y}$ as

$$\begin{aligned}\hat{y} &= \frac{(\varepsilon-\phi-1)\varphi(z)\omega}{\Delta}\left[\alpha mc + (1-\alpha)(\hat{\nu}-\hat{\chi})\right] + \hat{\nu}-\hat{\chi}\\ &= \frac{(\varepsilon-\phi-1)\varphi(z)\omega\alpha}{\Delta}mc + \frac{1}{\Delta}\left(\hat{\nu}-\hat{\chi}\right).\end{aligned}$$

We therefore get the following

$$mc + \frac{1-\alpha}{\alpha}\hat{y} = \frac{1}{\Delta}\left[mc + \frac{1-\alpha}{\alpha}\left(\hat{\nu}-\hat{\chi}\right)\right].$$

The resulting variance is then

$$Var\left(mc + \frac{1-\alpha}{\alpha}\hat{y}\right) = \frac{1}{\Delta^2}\left[\sigma^2_{mc} + \left(\frac{1-\alpha}{\alpha}\right)^2(\sigma^2_{\hat{\nu}} + \sigma^2_{\hat{\chi}}) + 2\frac{1-\alpha}{\alpha}\sigma_{mc,\hat{\nu}}\right],$$

which can be used in equation (2.13), together with equation (2.12), to derive conditions (2.22) and (2.23).

Proof of Proposition 1.

Note that

$$\begin{aligned}\Delta &= \frac{\alpha + \varepsilon(1-\alpha) - \varphi(z)(1-\alpha)(\varepsilon-\phi-1)}{\alpha+\varepsilon(1-\alpha)}\\ &= \frac{\alpha - \varphi(z)(1-\alpha)(\phi-1) + \varepsilon(1-\alpha)(1-\varphi(z))}{\alpha+\varepsilon(1-\alpha)} > 0,\end{aligned}$$

which holds since $\phi < 1$. Furthermore, $\Delta = 1$ at $z = 0$, such that the left-hand-side of inequality (2.22) is positive at $z = 0$. At this point, the right-hand-side $\Phi(0) = 0$ (there is a firm that has zero costs of investing in price flexibility). Moreover, $\Phi'(z) > 0$. The sign of the slope of the left-hand-side is determined by

$$\frac{\partial \Delta^{-2}}{\partial z} = 2\Delta^{-3}\omega(1-\alpha)(\varepsilon - \phi - 1)\varphi'(z).$$

This expression is positive if $\phi > 1 - \varepsilon$ and vice versa. A positive slope corresponds to strategic complementarity in the choice of flexibility: the more firms choose to invest in price flexibility, the more it pays off for an individual firm to also do so. A negative slope corresponds to strategic substitutability in the choice of flexibility, see Devereux (2006). We hence get a unique equilibrium value for z if $\phi \leq 1-\varepsilon$. Note that the second derivative of Δ^{-2} with respect to z can only be negative if the first derivative is also negative. For $\phi > 1-\varepsilon$, we have therefore three possibilities: a) one unique equilibrium at $0 < z < 1$, b) one unique equilibrium at $z = 1$, or c) three equilibria, one for a low value of $0 < z < 1$, one at an intermediate value of $0 < z < 1$, and one at $z = 1$. All considered equilibria are stable—except for the intermediate one in the case of three equilibria—as for lower z the benefit of investing in price flexibility (left-hand-side of inequality (2.22)) is higher than the costs $\Phi(z)$. We therefore disregard the intermediate equilibrium in the case of three equilibria. If we are already at the corner solution, z can obviously not rise any further. Since the left-hand-side of inequality (2.22), for any given value of z, is increasing in $\sigma^2_{mc}, \sigma_{\hat{\chi}}$, and $\sigma^2_{\hat{\nu}}$, and its slope is, for interior solutions, larger than that of the right-hand-side, Proposition 1 obtains. ■

Proof of Corollary 1.

The volatility of the price level (2.21) is

$$\begin{aligned}
\sigma_p^2 &= \left(\frac{\varphi(z)\omega}{\Delta}\right)^2 \left[\alpha^2\sigma^2_{mc} + (1-\alpha)^2(\sigma^2_{\hat{\nu}} + \sigma\hat{\chi}^2) + \alpha(1-\alpha)\sigma_{mc,\hat{\nu}}\right] \\
&= \frac{2\alpha(\varphi(z)\omega)^2}{\Omega}\Delta(\Theta).
\end{aligned}$$

The corollary directly follows from this. ■

Proof of Proposition 2.

The direct effect of a changing ϕ is visible when taking the derivative with respect to ϕ of the term in the price index (2.21) that multiplies all shocks:

$$\frac{\partial \varphi^2(z)\omega\Delta^{-1}}{\partial \phi} = \varphi(z)\omega\Delta^{-2}\omega(1-\alpha) > 0.$$

Reducing ϕ (stricter inflation targeting) hence decreases the effect of shocks on inflation for a given value of z. The indirect effect of changing ϕ on z depends on the following derivatives (remember that $\Delta > 0$ from the proof of Proposition 1):

$$\frac{\partial \Delta^{-2}}{\partial \phi} = 2\Delta^{-3}\varphi(z)\omega(1-\alpha) \geq 0$$
$$\frac{\partial \Delta(\Theta)}{\partial \sigma_{mc,\hat{\nu}}} = \frac{\Omega(1-\alpha)}{\Delta^2} > 0,$$

where the first derivative determines the sign of $\partial\Delta(\Theta)/\partial\phi$ and $\Delta(\Theta)$ is the left-hand-side of inequality (2.22). Proposition 2 follows directly from these derivatives. ■

3. Different Newspapers—Different Inflation Expectations

Author

Sarah Arndt[1]

Abstract

I investigate how inflation signals from different types of newspapers influence household inflation expectations in Germany. Using text data and the large language model GPT-3.5-turbo-1106, I construct newspaper-specific indices and find heterogeneity in their informativeness based on the genre—tabloid versus broadsheet. Local projections reveal that tabloid sentiment shows an immediate decrease following a monetary policy shock, whereas responses from broadsheet newspapers are smaller and delayed. Household expectations are also affected differently by positive shocks to newspaper sentiment (causing the reporting to suggest increasing inflation), depending on the type of newspaper. These findings underscore the differentiated impact of media on inflation expectations across various segments of society, providing valuable insights for policymakers to tailor communication strategies effectively.

Keywords

Inflation expectations, text mining, forecasting, monetary policy, LLM, ChatGPT

3.1 Introduction

Households observe various inflation signals ranging from grocery prices to newspaper articles and social media posts. They use these sources of information to form their expectations about future inflation. Similar to their heterogeneity in information sources, households' inflation expectations vary persistently across different groups (see Figure 3.1-1), as pointed out by Meeks and Monti (2023). While we can explain a share of these variations by socio-economic characteristics, the reason why these various groups persistently differ in their expectations remains not yet entirely understood. In this paper, I explore the role of newspapers in potentially contributing to this observed heterogeneity. I show that households are systematically affected by the newspapers they read, and that some newspapers reflect the expectations of their respective audiences. Specifically, I compare the different signals towards inflation emitted by the yellow press, represented by Germany's largest newspaper in terms of sold units—*Die Bild*, and more reputable German daily newspapers like the *Frankfurter Allgemeine Zeitung* (FAZ) and *Die Süddeutsche*

[1]**Acknowledgements:** I extend my gratitude to the FAZ and SZ archives for granting access to their online databases, which allowed me to retrieve the necessary newspaper articles. Their cooperation is deeply appreciated. I also thank Zeno Enders, Timo Dimitriadis, and the participants of the 1st NEAR Conference, the MB seminar at the ifo Institute, and the AWI internal seminar for their valuable discussions and comments.

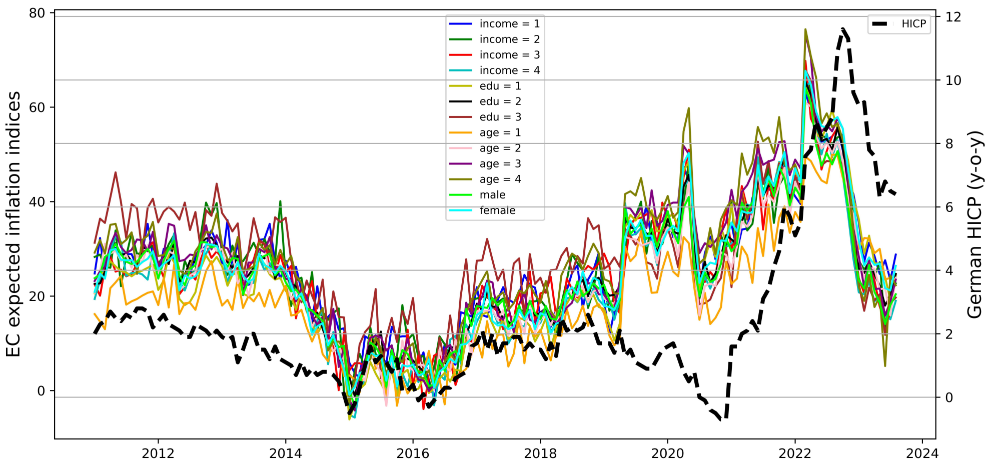

Figure 3.1-1: EC survey expectation indices for different socioeconomic groups (left-axis)
Notes: Right axis: monthly year-on-year growth rates of German HICP (black dashed line). Sample span: January 2011 to August 2023.

Zeitung (SZ), which are the two most popular representatives. I collect inflation-related articles from these three newspapers and analyze their sentiment towards inflation with the large language model gpt-3.5-turbo-1106 developed by OpenAI. I aggregate the raw sentiment scores per article to overall sentiment indices per newspaper by taking the monthly averages.

When following the reporting of these three newspapers, one can observe significant differences. Even though inflation was below target in 2020, by the end of that year, the Bild newspaper suddenly increased its reporting on inflation, as shown by the blue shaded area in Figure 3.2-2, representing the absolute weekly number of articles related to the topic of inflation. This press is known for its polarizing style of reporting and published headlines such as *"The inflation monster is back!"* (May 8, 2021), even though inflation was barely exceeding 2% back then, and one month later: *"Inflation is eating away our savings"* (June 2, 2021), fueling households' fear of inflation. There is the potential risk that public opinion had been distorted by this sudden increase of dramatized media coverage.

To test whether different newspaper sentiment indices can account for a share in the heterogeneity of inflation expectations, I compare the newspaper indices with household survey expectations on inflation, differentiated by income and education of the survey participants. The analysis reveals heterogeneity in contemporaneous correlations, suggesting that expectations derived from the yellow press resonate more with lower-education households, while the opposite holds true for FAZ and SZ.

In a second step, I analyze the forecasting power of the three news indices by targeting different household subcategories' survey expectations through an out-of-sample forecasting exercise. The results reveal heterogeneity among the indices, suggesting that the Bild index is more suitable for predicting inflation expectations among lower-education survey participants. Conversely, the FAZ and SZ indices outperform the Bild index in correlation with overall inflation in Germany. In general, all newspaper indices carry some degree of informativeness regarding household survey measures on inflation, as including them in a forecasting model reduces the root mean squared error in most cases. This makes these indices an attractive alternative measure of household inflation expectations and perceptions due to their high-frequency, real-time availability.

Building on this analysis, I investigate the effect of a euro zone monetary policy shock on the three newspaper inflation expectation indices. While FAZ and SZ display a more muted and similar response, the Bild index reacts with an immediate reduction in overall sentiment in the days following a shock. This observation suggests that either monetary policy is more effective in tempering Bild's negative sentiment towards inflation, or it could be attributed to Bild's higher volatility in reporting on inflation, which peaks around monetary policy meetings.

Lastly, I explore another potential source of heterogeneity by comparing the responses of household inflation expectations to shocks in the sentiment of the three newspapers. I define a sentiment shock as an increase in the newspaper indices that cannot be explained by current and past economic conditions (e.g., inflation and industrial production), monetary policy, and other newspapers' sentiment. The results show that survey participants with higher income and education levels are most responsive to shocks in FAZ and SZ sentiment. This indicates that some of the heterogeneous informativeness of the indicators can be attributed to their varying ability to influence households' inflation expectations.

So far, as Meeks and Monti (2023) highlight, the existing literature and policy discussions have primarily focused on aggregate household inflation expectations. Previous studies that derive inflation expectation indices from text data, specifically posts on Twitter (now known as X), test the informativeness of their indicators only against aggregate household survey expectations (see, e.g., Angelico et al., 2022; Born et al., 2023). One key contribution of this paper is to differentiate survey expectations across various socioeconomic groups and test the informativeness of different text sources for these specific households. This approach sheds light on potential drivers of heterogeneity in household inflation expectations.

There is already a growing body of literature that employs newspaper text data and machine learning methods to extract sentiment and expectations, often used for forecasting macroeconomic variables—see, for example, Bybee (2023), Shapiro et al. (2022), Picault et al. (2022), or Ellingsen et al. (2022). However, these studies typically use a single or a few reputable newspapers as representative of all news data and overlook the fact that different genres of newspapers might emit varying signals that reach audiences differently. In a similar vein to this paper, Menz and Poppitz (2013) explain socioeconomic differences in inflation expectations in Germany through heterogeneity in news consumption. They rely on manually labeled (into positive and negative) newspaper articles, including Bild (but not FAZ and SZ), and television news from a sample spanning 1999 to 2010. Unlike my study, their aim is not to extract inflation expectation sentiments from the news data but to test the epidemiological model introduced by Carroll (2003) to see if tone and news coverage bring household expectations closer to those of professional forecasters.

Furthermore, this paper contributes to the literature on inflation expectation elicitation by offering a real-time proxy to measure heterogeneous household expectations. For an overview of recent literature on inflation expectations, readers are referred to D'Acunto et al. (2023). Weber et al. (2023) find that households are better informed about inflation and more attentive during times of rising prices. Additionally, Andre et al. (2023) emphasize the importance of media in shaping and creating macroeconomic narratives that ultimately affect household expectations, while Conrad et al. (2022) show that households that follow traditional media have lower inflation expectations.

The paper is structured in the following way: Section 3.2 describes the newspaper and survey data set. Section 3.3 outlines the construction of the newspaper inflation expectations indices, followed by an analysis of their informativeness in Section 3.4. I present the effects of a monetary policy shock on the newspaper indices in Section 3.5 and the responses of survey expectations after a shock in news sentiment in Section 3.6. Finally, Section 3.7 provides some robustness checks and Section 3.8 concludes.

3.2 Data sources

3.2.1 Newspaper data

In Germany, there is only one popular daily nationwide yellow press newspaper—*Die Bild* (Bild), which recorded 1,146,616 sold units of its Sunday edition in the third quarter of 2023 and about 500 million page visits each month in 2023 on its online news portal *Bild.de* (Informationsgemeinschaft zur Feststellung der Verbreitung von Werbeträgern e.V. (IVW), 2023). It is widely known for its polarizing style of reporting, characterized by exaggerations, neologisms, and an oversimplified language. In contrast, the two most popular nationwide daily newspapers considered quality news sources are the *Frankfurter Allgemeine Zeitung* (FAZ) and *Die Süddeutsche Zeitung* (SZ). Each recorded 454,319 and 498,730 sold units respectively in the third quarter of 2023, with around 55 and 65 million monthly page visits on their news portals *Süddeutsche.de* and *FAZ.net* in 2023 (IVW, 2023). I use articles from both FAZ and SZ to increase the sample of quality news articles and cover a larger segment of the German population.

I collect FAZ and SZ articles from their respective news archives using search queries that filter articles related to inflation in Germany. The queries employ German terms such as *central bank*, *high prices*, *deflation*, *inflation*, *price index*, *wage-price-spiral*, or *monetary policy*[2]. Both archives include print and online articles. Bild newspaper articles are obtained from the online database *LexisNexis*. Since Bild articles are only available until 2017 there, I additionally use the open-source platform *mediacloud.org* (hereafter mediacloud) to collect URLs directing to inflation-related Bild online articles. I scrape all collected web pages to obtain the article data. This procedure results in a total dataset of 16,157 FAZ, 13,365 SZ, and 16,002 Bild articles published from January 2010 to August 2023. Despite using specific search queries, not all collected articles are actually concerned

[2]Detailed search queries for all newspapers can be found in Appendix 3.9.1.

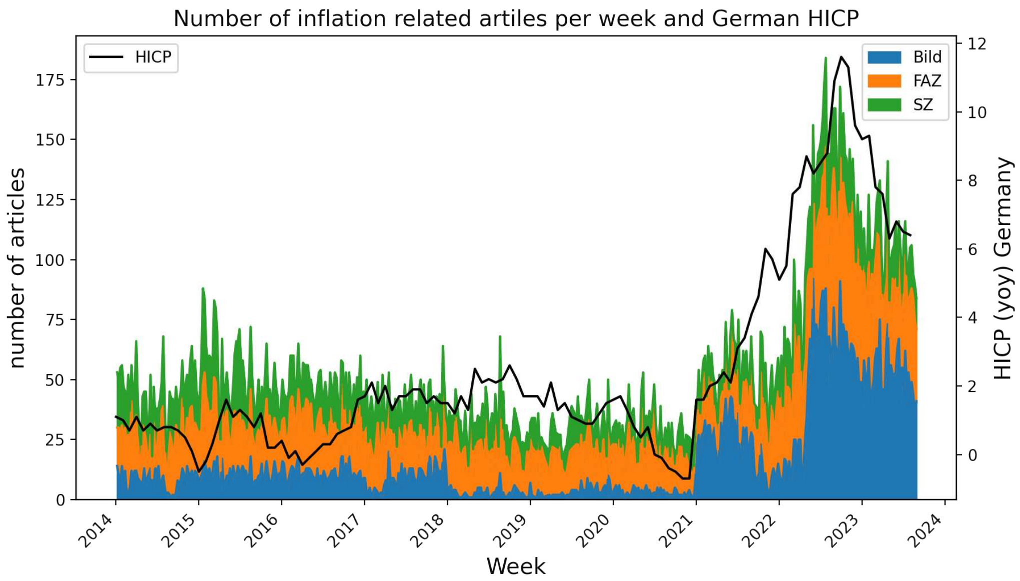

Figure 3.2-2: Weekly numbers of newspaper articles
Notes: Left axis: weekly absolute numbers of SZ (green), FAZ (orange), and Bild (blue) newspaper articles dealing with inflation in Germany (selection based on classification by ChatGPT), stacked on top of each other. Right axis: monthly year-on-year growth rates of German HICP (red line). Sample span: January 2014 to August 2023.

with inflation in Germany; this is particularly true for Bild articles. Therefore, Figure 3.2-2 depicts the distribution of articles over time that are genuinely related to inflation in Germany alongside the year-on-year growth rates of the German harmonized index of consumer prices (HICP). At the beginning of the sample period, there are very few Bild articles on inflation but the number of articles increases significantly starting from 2021. In contrast, a substantial number of inflation-related articles are available throughout the entire sample period for FAZ and SZ, with a noticeable increase starting in 2021 that coincides with rising year-on-year growth rates of German HICP.

3.2.2 Survey data

To assess the informativeness of the newspaper indicators regarding households' inflation expectations, I compare them to the European Commission's (EC) *Business and Consumer Survey* (BCS). The BCS is conducted on a monthly basis in all countries of the European Union. For Germany, the sample dates back to 1985 and includes around 2,000 households per survey wave. Participants are surveyed in the first two to three weeks of each month. To elicit the participants' inflation expectations, they are asked to provide answers to the question *"By comparison with the past 12 months, how do you expect that consumer prices will develop in the next 12 months? They will..."*, allowing the answers *"increase more rapidly"* (PP), *"increase at the same rate"* (P), *"increase at a slower rate"*, *"stay about the same"* (M), *"fall"* (MM), or *"don't know."* These qualitative answers are aggregated to an index, which is basically the difference in percentage shares of participants that think inflation will increase vs. it will decrease:

$$\mathbb{E}_t^{\pi_{t+12}} = PP + \frac{P}{2} - \frac{M}{2} - MM. \tag{3.1}$$

3.2.3 Readership data

The notion that different types of newspapers influence households' perspectives on inflation differently stems from the observation that the composition of their audiences varies significantly. Figure 3.2-3 illustrates the average number of readers per issue in millions, disaggregated by education and income level, age, and gender. Given that each newspaper has different levels of the average number of readers per issue (with Bild having the largest), I normalized the numbers to 1 million readers per issue for comparability. The data source is the *Media-Analyse Tageszeitungen 2022*, a computer-assisted self/web interviewing survey comprising approximately 40,000 interviews about participants' daily newspaper consumption. While the survey's household categories do not perfectly align with those of the BCS, further details on data matching can be found in Appendix 3.9.3.

In the bar charts in Figure 3.2-3, one can see that the composition of readers especially differs across education and income level. Most FAZ and SZ readers have attained the tertiary education level or further, while the largest share of Bild readers has only reached secondary education. Interestingly, all three newspapers have the largest readership among the highest income quartile, while for FAZ and SZ they outnumber the lower income groups by far. The distribution across age and gender is fairly even, although all three newspapers have more male than female readers.

3.3 Newspaper-based inflation expectation indices

3.3.1 Applying ChatGPT for sentiment analysis

I extract inflation sentiment signals from the newspaper texts with OpenAI's ChatGPT-3.5 model (precisely, *gpt-3.5-turbo-1106*), a large language model trained on data up to September 2021. Specifically, I feed each article to the prompt depicted in Figure 3.3-4 using OpenAI's Python API. ChatGPT is asked to provide a sentiment score for each article to capture if the article creates the impression that inflation is likely to in- or decrease. The prompt is designed in a way such that the generated answers are either integers on a scale from 1 (inflation decreases) to 10 (inflation increases) or the word 'no', if the respective article is actually not about inflation in Germany. I set the model's temperature parameter to 0, which makes it generate a deterministic output when given identical inputs (Wang et al., 2023). Specifying a system role additionally helps the model to fulfill this task correctly.

In comparison to simpler methods for text-based sentiment measurement, like the widely used dictionary approach (see for example Shapiro et al. (2022) or Angelico et al. (2022)), an advantage of ChatGPT is its advanced ability to understand human language, including nuances in the tone of the text. This is especially important when analyzing yellow press

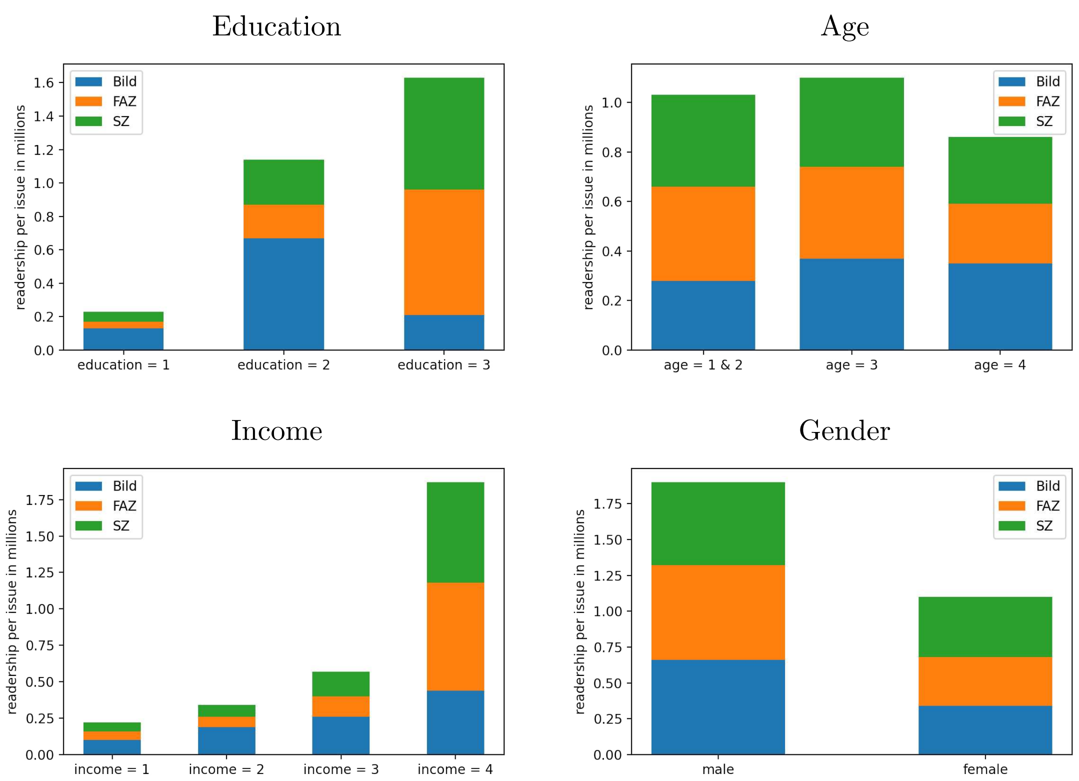

Figure 3.2-3: Readership characteristics

Notes: Readership per Bild (blue), FAZ (orange), or SZ (green) issue in 2022, normalized to 1 million readers per newspaper. Disaggregated per education and income level, age cohort and gender such that the data matches the EC survey household subcategories. Data source: agma (2022).

```
System: You are an AI that rates inflation sentiment from 1 to 10. You can
only respond with a number between 1 and 10.

User: Provide a sentiment score for inflation in this article on a scale of 1
to 10, where 1 means inflation is likely to decrease, and 10 means inflation is
likely to increase. If the article does not address inflation in Germany or the
EU, respond with the word 'no': “%title” + “%article”
```

Figure 3.3-4: Prompt for generating inflation expectation sentiments

Notes: Model used: ChatGPT-3.5. ”%title” and ”%article” are placeholders for single articles and their headlines, respectively.

German Headline	Translation	Date	Score
Kommentar; Schluss mit der **Inflations-Heuchelei**	Commentary; End to the **inflation-hypocrisy**.	10/12/2021	8
Ausgerechnet in der Inflations-Krise!; Dautschlands **oberster Geld-Hüter** gibt auf	Right in the inflation crisis!; Germany's **top money guardian** gives up	10/21/2021	8
Madame Inflation; **Luxus-Lagarde** macht Sparer und Rentner arm	**Madame Inflation**; **Luxury Lagarde** impoverishes savers and retirees	10/30/2021	9
Nach dem **Teuer-Schock**; So einfach schützen Sie Ihr Geld vor **Madame Inflation**!	After the **expensiveness-shock**: How to protect your money from **Madame inflation**!	11/18/2021	8
Zins-Hammer! EZB erhöht auf 1,25 Prozent	**Interest rate hammer**! ECB increases to 1.25 percent	09/08/2022	9

Table 3.3-1: Bild-typical language and sentiment scores

Notes: Examples of Bild-typical language and its use of neologisms in the context of inflation in Germany and the euro zone, alongside the respective inflation sentiment score provided by ChatGPT.

articles like those of the Bild newspaper, as they are characterized by a polemic language and the constant creation of new expressions. The highly adaptive nature of the Bild language makes creating a Bild-suited dictionary impractical as it would constantly need to be expanded. Table 3.3-1 presents examples of Bild expressions and tentative English translations thereof. As shown in Appendix 3.9.2, ChatGPT is indeed capable of understanding the expressions highlighted in Table 3.3-1 and capturing the headline's intended tone.

Besides ChatGPT's advanced ability to understand human language, the model additionally proves to perform well at analyzing the sentiments of texts. Belal et al. (2023) find that ChatGPT-3.5 significantly outperforms two lexicon-based algorithms (VADER and TextBlob) in analyzing the sentiment of Amazon product reviews and a soccer tweets data set. Comparing ChatGPT-3.5 to a fine-tuned BERT[3] model, Wang et al. (2023) find that ChatGPT scores equally well on zero-shot sentiment classification tasks and that it even beats the BERT model in recognizing sentiment polarity shifts.

3.3.2 Index construction

The raw data generated by ChatGPT consists of article-level sentiment scores from 1 to 10, or 'no' if the article does not cover German inflation. Since the frequency of articles is daily, it is feasible to aggregate the data into indices at any higher frequency. In this analysis, I focus on monthly aggregates to align with the monthly frequency of the household expectation surveys. As the baseline aggregation method, I take the mean over all articles published that ChatGPT deems to be related to inflation in Germany, to obtain a monthly sentiment index. This is just one of many possible ways to aggregate the data, as pointed out by Angelico et al. (2022). For robustness, I also test other ways of aggregation, like taking the moving average over various window lengths, or computing the exponentially weighted mean with different smoothing parameters. The alternatives are very similar to

[3] The BERT (Bidirectional Encoder Representations from Transformers) model is another large language model originally developed by Google (Devlin et al., 2018). Its set of parameters can additionally be fine-tuned on a specific task or dataset.

the baseline index and yield comparable results in the subsequent forecasting exercise (see Appendix 3.9.4).

Figure 3.3-5 shows the mean index for each of the newspapers separately and the FAZ and SZ articles data combined. Comparing the newspaper indices to actual year-on-year growth of German HICP reveals some correlation over time, with the newspaper indicators beginning to align with inflation trends at the end of 2021. While the indices based on FAZ and SZ data are very well aligned with each other, the Bild inflation index differs significantly. The Bild inflation index reaches higher sentiment scores than the other two newspapers, which is especially visible in the years around 2019. Actual inflation in Germany was mostly below the ECB's target during this period, so it is surprising that Bild's sentiment is significantly higher than that of the other two newspapers. A closer examination of the articles from that time reveals the cause of this spike in sentiment: In 2019, the ECB's key interest rate hit the zero lower bound, prompting banks to introduce negative interest rates on deposits. Bild reported very negatively on this development, spreading the narrative that it imposed substantial costs on "normal people" due to the lack of interest on their savings.[4] Additionally, Bild articles exaggerated the severity of inflation. Any positive level of inflation, even if below 2%, was portrayed very negatively, as it further eroded the purchasing power of deposits. Conversely, deflation was depicted as highly beneficial.

Due to the scarcity of Bild articles at the beginning of the sample, the Bild sentiment index is often inconclusive for the period before 2014. Therefore, I truncate the sample at 2014 in the subsequent analyses.

3.4 Informativeness of the newspaper indices

3.4.1 Heterogeneous correlations

First, I compare the sample correlations between the inflation sentiment indices of the different newspapers and the EC household survey measures of expected inflation. The household expectation indices are expressed as weighted averages of the respective newspaper readership compositions (weights are the same as those depicted in Figure 3.2-3). This analysis reveals initial heterogeneity in the strength of the co-movement between each newspaper index and the household subgroups. Specifically, I calculate the sample correlation with a one-month lag for the newspaper indices, as households are surveyed in the first weeks of each month, ensuring that the indices reflect only the information available to households at that time. From this point, the analysis will focus on heterogeneity in readership characteristics related to income and education levels, as these dimensions exhibit the strongest differences across the various newspapers.

Figure 3.4-6 presents the resulting correlation heat maps between the three different newspaper indices (on the y-axis) and the readership-weighted EC survey's measures of inflation expectations, and German inflation (on the x-axis). Generally, when considering the full

[4] Mario Draghi, the ECB's president at that time, was dubbed "Graf Draghila" by Bild, a reference to Count Dracula (see, for example, articles from 02/01/2020, 10/24/2019, and 09/12/2019).

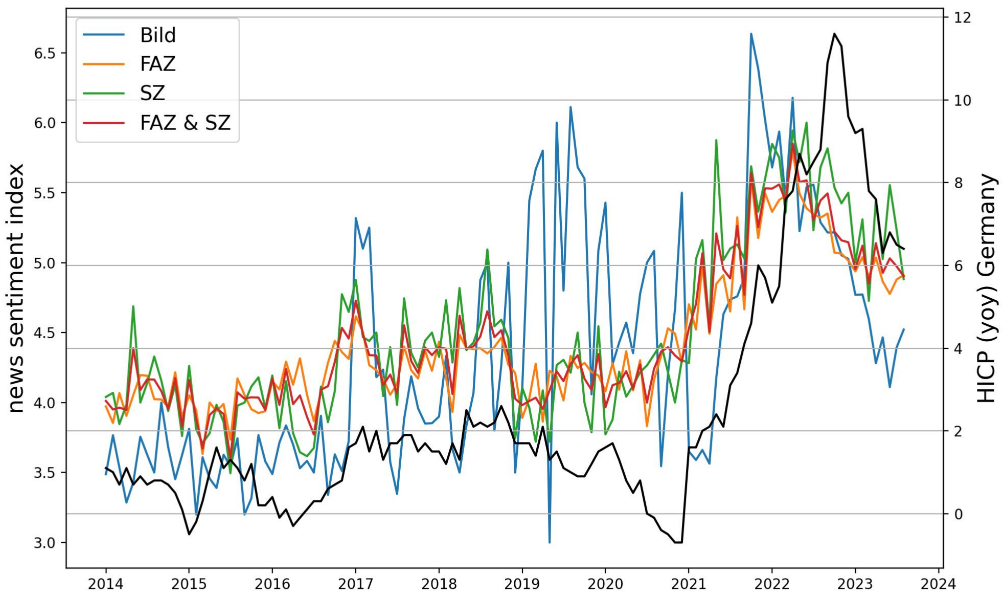

Figure 3.3-5: Monthly mean sentiment scores
Notes: Bild (blue), FAZ (orange), SZ (green) and both newspapers FAZ & SZ combined (red), against year-on-year inflation in Germany (black).

sample starting in January 2014, the Bild index achieves the highest sample correlation with household expectations weighted by their own readership's education levels (0.688), exceeding those weighted by the FAZ and SZ audience (0.622 and 0.637, respectively). While the Bild index is comparably weakly correlated with German HICP (last column in the heat map), FAZ and SZ indices attain high correlations with inflation, which is especially true for the second half of the sample (see Figure 3.4-7, panel b). There is more heterogeneity in the co-movements of the Bild sentiment index between the different household expectations, than for the FAZ and SZ indices, but overall the heterogeneity between different income groups is fairly constant.

To demonstrate that the high correlations between the Bild index and inflation expectations are not only driven by the recent post-COVID-19 high inflation period in Germany, I split the sample at 2020 and report the respective results in Figure 3.4-7. Panel a) shows that Bild's pre-2020 sample correlation is higher than during the recent high-inflation period depicted in panel b), while the opposite holds true for the FAZ and SZ indices. The high correlations of the latter with inflation itself are as well driven by the recent years.

3.4.2 Heterogeneous forecasting powers

In a second step, I run an out-of-sample forecasting exercise and compare the ability of the different newspaper indices to predict the survey expectations measures of the various groups of households. As a benchmark, I estimate an autoregressive model ($AR(p)$) for each newspaper-weighted survey expectation measure, determining the optimal lag order p for each model and forecast iteration using the BIC criterion. The in-sample starts in January 2014 and ends in December 2016, translating into a training sample length of 36 months. I then augment each benchmark model with one of the inflation expectation

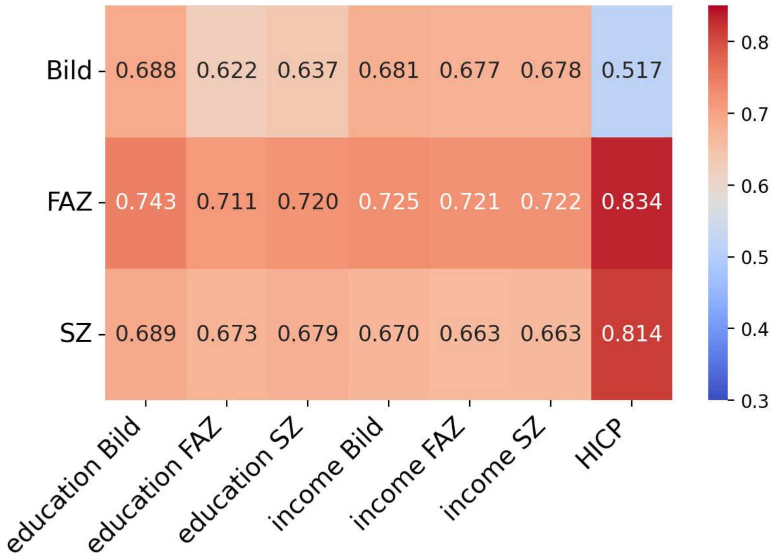

Figure 3.4-6: Correlations between newspaper indices and inflation expectations
Notes: Correlations between the newspaper indices of Bild, FAZ and SZ, and the readership-weighted EC household inflation expectations, as well as inflation in Germany (HICP), on a sample spanning from 2014M01 to 2023M08. The newspaper indices are lagged by one month.

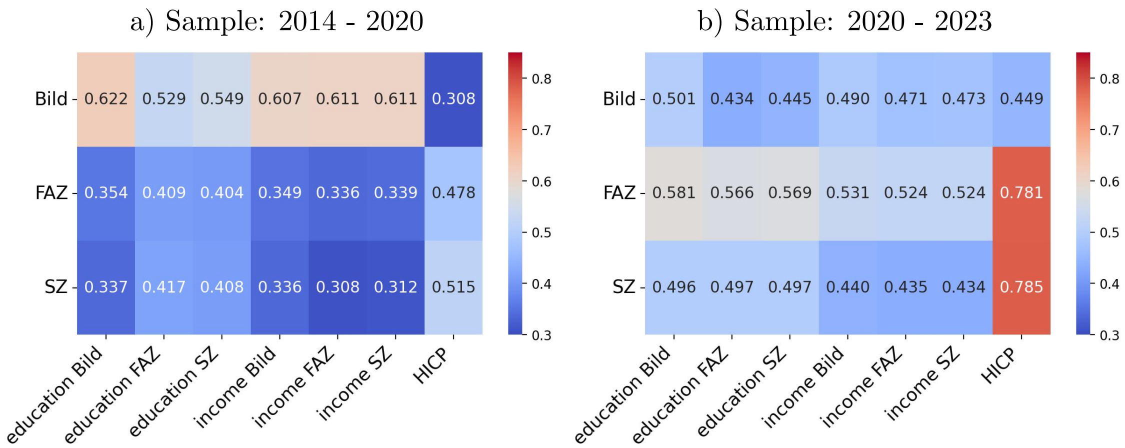

Figure 3.4-7: Correlations: sample splits
Notes: Panel a) shows the correlations between the newspaper indices of Bild, FAZ and SZ, and the readership-weighted EC household inflation expectations, as well as inflation in Germany (HICP), on a sample spanning from 2014M01 to 2020M01 and Panel b) from 2020M02 to 2023M08, respectively. The newspaper indices are lagged by one month.

	Education				Income			
h	Bild	FAZ	SZ	FAZ & SZ	Bild	FAZ	SZ	FAZ & SZ
1	1.02	1.01	1.02	1.01	1.02	1.01	1.01	1.01
2	1.01	0.93	0.94	0.93	1.03	1.02	1.01	1.02
3	1.03	0.92	0.93	0.93	1.05	1.04	1.04	1.04
4	1.04	0.91	0.93	0.92	1.07	1.07	1.07	1.07
5	1.05	0.91	0.93	0.92	1.08	1.09	1.09	1.09
6	1.06	0.92	0.94	0.93	1.1	1.11	1.11	1.11

Table 3.4-2: RMSE baseline results

Notes: This table presents the relative RMSE of the FAZ & SZ indices compared to the Bild index. Values above 1 indicate that adding the Bild index to the benchmark model outperforms the other newspaper indices. The left panel shows results for predicting households' inflation expectations, weighted by the education characteristics of the average reader of the respective newspapers (Bild, FAZ, SZ, and FAZ & SZ combined). The right panel displays the equivalent results weighted by income characteristics. RMSEs are calculated on a sample spanning from January 2017 to August 2023.

indices for each newspaper separately and also a combined version of FAZ and SZ data. To be conservative, as the BCS participants are surveyed in the first weeks of a month, all newspaper indices enter with one lag. The forecast horizon is six months and all six horizons are predicted recursively.

I compare the forecasting powers by calculating the root mean squared error (RMSE) for each model. Specifically, I divide the RMSE of the augmented model by that of the benchmark model; values below one indicate improved forecast accuracy when including the newspaper indices. Next, I divide the RMSE ratios of FAZ, SZ, and both newspapers combined by the RMSE ratio of the Bild newspaper to determine which newspaper type provides higher information gains when predicting household expectations—tabloids or serious newspapers. In this case, values above 1 indicate that the Bild index outperforms.

Table 3.4-2 reports the relative RMSE between FAZ and SZ combined[5], and the Bild newspaper over a forecasting horizon of one to six months. While there is almost no heterogeneity in forecasting power of the Bild index across different income groups, weighted by the respective newspaper readership, heterogeneity is much larger in education-weighted readership expectations. When predicting the inflation expectations of households with the education level of a typical Bild reader, the tabloid Bild outperforms the FAZ & SZ index, while the contrary holds true when predicting household's expectations with the average FAZ and SZ reader education level. I explain the lack of heterogeneity with respect to household income by the fact that for Bild, the largest share of readers is also found in the fourth income quartile. This might also suggest that the chosen matching between the EC survey income bins and the *Media-Analyse* income bins is rather imperfect and more granular data is needed for refinement.

To summarize, these findings suggest a pattern wherein the inflation expectation indicator from tabloids like Bild more accurately predicts expectations of households with lower education levels, which is in line with the readership characteristics depicted in Figure 3.2-

[5] Because the results for FAZ and SZ are very similar, I subsequently only report the ones when combining both newspapers (results for FAZ and SZ individually can be found in Appendix 3.9.5).

3. Conversely, the sentiment reflected in newspapers like FAZ and SZ tends to resonate more with more educated households, and aligns more closely with actual inflation figures.

3.5 Monetary policy shocks and news sentiments

To explore a possible source of the heterogeneity between the newspaper indices, I estimate the newspaper's sentiment reaction after a monetary policy shock. To do so, I use a local projections model as introduced by Jordà (2005):

$$ln(y_{t+h}) - ln(y_{t-1}) = \mu_h + \beta_h shock_t + \mathbf{\Delta}_h \sum_{l=0}^{n} W_{t-l} + u_{t+h}. \tag{3.2}$$

I estimate the model at daily frequency for each horizon h separately. The model includes a horizon-specific constant μ_h and a set of controls, W_t, consisting of the contemporaneous value and $n = 5$ lags (i.e., days) of the respective 3-day moving average newspaper sentiment index and German HICP. Similar to Born et al. (2023), I use a backward looking moving average of each newspaper's daily mean sentiment index to smooth out the days without any news reporting. Here, I use a window length of three days. The dependent variable $ln(y_{t+h}) - ln(y_{t-1})$ measures the cumulative percentage change of one of the three newspaper sentiment indices. Finally, the variable $shock_t$ is a shock series of ECB monetary policy surprises generated as in Jarociński and Karadi (2020) and provided online by Marek Jaroci'nski[6]. For estimating euro area monetary policy shocks they rely on the *Euro Area Monetary Policy Event-Study Database* developed by Altavilla et al. (2019) which contains intra-day asset price changes around ECB policy announcements. The shock series is available in daily frequency, even though at maximum one monetary policy decision is taken per month. As inflation is only available at monthly frequency, it enters the regression as a monthly value repeated for each day of the respective month. The model is estimated on a sample spanning from 2014M01 to 2023M08, due to the limited data availability of Bild articles at the beginning of the sample. The impulse responses are given by the β_h coefficients for each horizon separately and are reported with 95% Newey-West standard errors.

Figure 3.5-8 displays the cumulative impulse response functions over a horizon of 10 days. One period after the impact of a contractionary monetary policy shock, the Bild sentiment immediately starts to decline, while the SZ sentiment shows a significant decrease only after three and the FAZ sentiment does not decrease at all over the considered horizon. The responses of SZ and FAZ appear quite similar at the beginning, although the FAZ sentiment seems to lag behind SZ for a few days.

To examine these findings more closely, I restrict the sample to articles containing at least one of the German words for ECB, central bank, or key interest rate. This focuses the analysis on articles explicitly addressing ECB policy. Figure 3.5-9 shows that in this subsample, all three newspaper sentiment indices only begin to decrease significantly after three days, which can partly be attributed to the three-day rolling window moving average.

[6] `https://marekjarocinski.github.io/jkshocks/jkshocks.html`

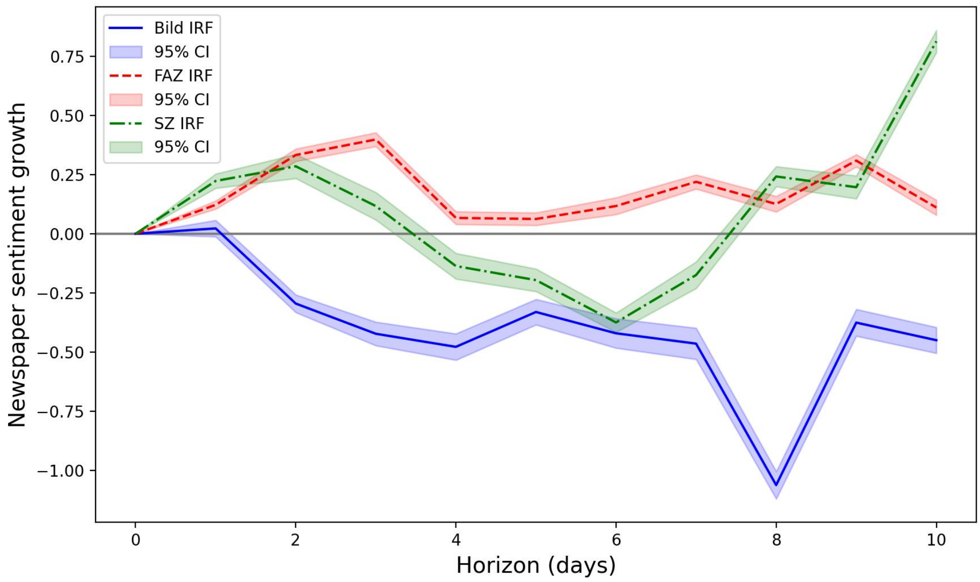

Figure 3.5-8: Monetary policy shock

Notes: Impulse responses of the three newspaper sentiment indices (Bild in blue, solid line, FAZ in red, dashed line, and SZ in green, dotted-dashed line) to a contractionary monetary policy shock on a 30-days horizon. Shaded areas are 95% confidence intervals computed with Newey-West standard errors.

This suggests that, prior to a monetary policy meeting, newspapers may write about it in a more negative tone, speculating and opining on the central bank's decision. The positive sentiment growth immediately after the shock reflects this pre-announcement sentiment, which then starts to decline as only post-shock newspaper data is considered. This shows, that a contractionary monetary policy is indeed effective in steering the newspaper's sentiment towards inflation downwards. While it seems to stay like this for the Bild sentiment, the persistence of the effect is smaller for FAZ and SZ.

3.6 Household's reaction to news sentiment shocks

In a second step, I investigate whether the heterogeneous informativeness of the three newspaper indicators stems from their potential ability to affect household inflation expectations and perceptions differently. To this end, I closely follow Shapiro et al. (2022)'s modeling strategy using local projections, but make some adaptations to this specific use case. Shapiro et al. (2022) estimate the effect of a structural news sentiment shock on measures of economic activity as the component of the news sentiment index that is *"orthogonal to current and 4 lags of economic activity as well as 4 lags of itself.*" In contrast, the dependent variable in my model is one of the four readership-weighted household inflation expectation measures by the EC survey. Likewise, I also identify the impulse responses of the expectations as the component of the newspaper sentiment indices that is orthogonal to the current value and $n = 6$ lags of the news sentiment index and indicators of economic activity, respectively. Specifically, I follow Lloyd and Manuel (2024)'s argumentation, in that one can directly identify the effect of a structural shock when appropriately controlling for confounding factors. I define a sentiment shock in one of the three newspapers

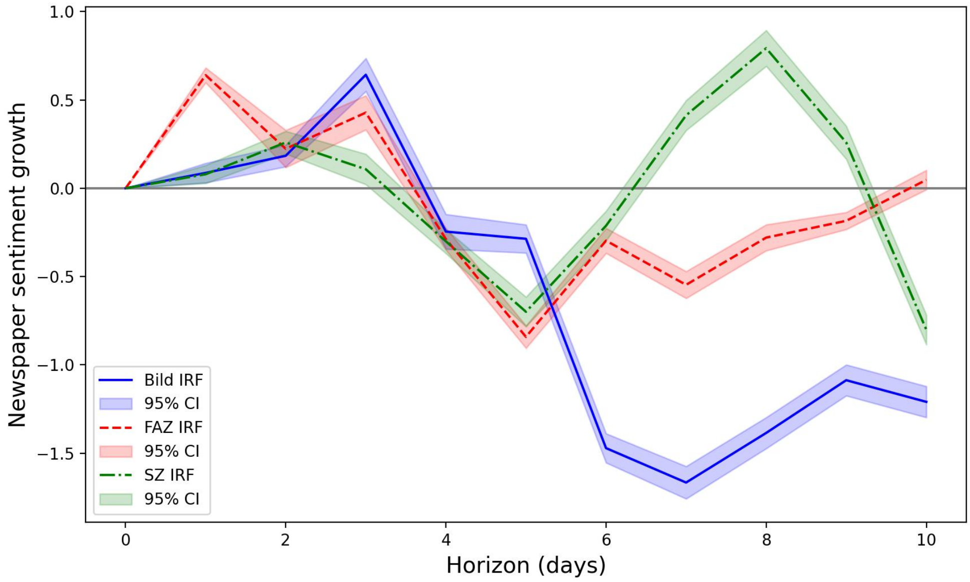

Figure 3.5-9: Monetary policy shock news sub-sample

Notes: Impulse responses of the three newspaper sentiment indices (Bild in blue, solid line, FAZ in red, dashed line, and SZ in green, dotted-dashed line) to a contractionary monetary policy shock on a 30-days horizon. Shaded areas are 95% confidence intervals computed with Newey-West standard errors. Articles are restricted to those containing at least one of the words "EZB", "Zentralbank", or "Leitzins".

as the part of the respective sentiment index that is orthogonal to current and past economic conditions (e.g., actual inflation is low, but the sentiment index suggests increasing inflation), that is not driven by monetary policy shocks, and that cannot be explained by other newspapers' sentiment. The responses I estimate result from a positive sentiment shock, which implies an increase in the newspaper's sentiment index, indicating inflation to be more likely to increase.

A separate regression is run for each forecast horizon $h = 1, \ldots, 6$ months:

$$y^i_{t+h} = \alpha^i_h + \beta^i_h NSI_t + \mathbf{\Gamma}^i_h \sum_{l=0}^{n} W_{t-l} + \sum_{l=1}^{n} \delta^i_{h,l} NSI_{t-l} + u^i_{t+h}, \tag{3.3}$$

where y^i is one of the $i = 1, \ldots, 13$ survey measures, α^i_h is a horizon-specific constant, and NSI denotes one of the four newspaper sentiment indices (Bild, FAZ, SZ, and FAZ & SZ combined). Both the survey and sentiment measures are expressed in first differences. The set of controls W_t is also based on Shapiro et al. (2022) and contains the contemporaneous value and 6 lags of the German industrial production index excluding energy and construction in log-differences (as a monthly proxy for output growth), year-on-year German HICP, a monthly measure of consumer confidence as a proxy for consumption, the Jarociński and Karadi (2020) ECB monetary policy shock series[7], as well as the FAZ & SZ newspaper sentiment index when estimating the response to a shock in the Bild index,

[7] As a robustness exercise, I alternatively include first differences of the ECB's main refinancing operations rate to control for monetary policy. The results, shown in Appendix 3.9.6, are not affected by this change.

and vice versa, including the Bild sentiment when estimating a response after a shock in FAZ and SZ sentiment. The first two series are obtained from the Federal Statistical Office of Germany, and the consumer confidence measure can be retrieved from Eurostat and is compiled by the EC. The impulse responses are then given by estimates of the β_h^i coefficients, which I report with 95% confidence bands constructed with Newey-West autocorrelation- and heteroskedasticity-robust standard errors. I estimate this model on 2014M01 to 2023M08 data, with the availability of Bild sentiment scores being the limiting factor.

Figures 3.6-10 and 3.6-11 show the resulting cumulative impulse response functions of readership-weighted expectations in response to shocks in Bild, and FAZ and SZ combined (because the individual results look almost identical) inflation sentiments, each figure representing one of the reader characteristics, income and education, respectively. The left panel in Figure 3.6-11 shows that the effect of a positive Bild sentiment shock on inflation expectations disaggregated by newspaper-characteristic income groups is not significantly different from zero, and it does not matter which newspaper's audience is targeted. In contrast, a shock to FAZ and SZ sentiment (right panel) does have a positive effect on households' expectations, even though, again, each newspaper's audience is equally affected. As mentioned earlier, one reason for this observation can be a mismatch of the survey income bins or similar reader-characteristic patterns.

When aggregating household expectations by each newspaper's average reader education level, as shown in Figure 3.6-10, there is again no significant response after a shock in Bild sentiment. In contrast, a positive shock in FAZ and SZ sentiment leads to an increase in inflation expectations. This effect is even stronger for the FAZ and SZ readership than for Bild readers.

The fact that a shock to Bild sentiment does not lead to a significant positive reaction in expectations might reflect the possibility that Bild reporting is primarily a reflection of the public's inflation expectations, instead of actively driving households' expectations. In contrast, more serious newspapers like FAZ and SZ seem to provide truly new information to their readers and therefore have a positive impact on inflation. As one would expect, they have a larger impact on the expectations of their typical readers, while households who are less likely to read FAZ and SZ are affected to a lesser extent.

3.7 Robustness checks

3.7.1 Consistency of ChatGPT

Even though I set ChatGPT's model temperature to zero, which is intended to make its responses more deterministic, there is no guarantee that the model will always produce the same output given the exact same input prompt. However, this setting significantly increases the likelihood of consistency. To demonstrate that this does not affect the overall informativeness of the newspaper sentiment indices generated by this zero-shot approach, I had the model reevaluate the articles' sentiments using a newer, recently released version of

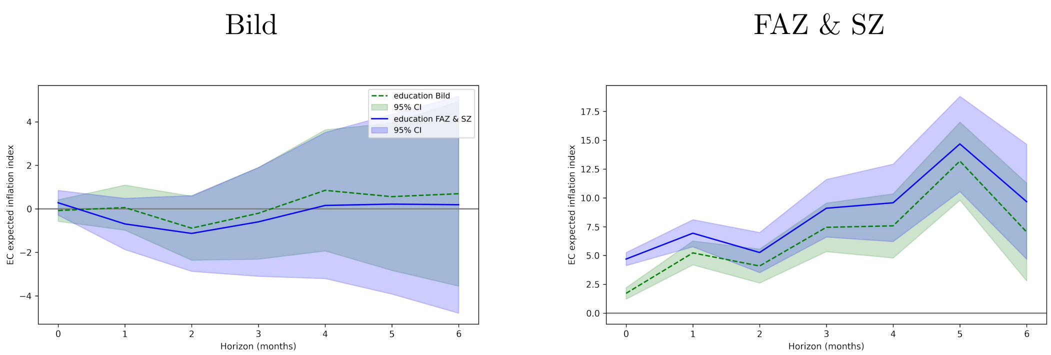

Figure 3.6-10: Newspaper sentiment shock by education level
Notes: Cumulative impulse responses of EC survey expectations disaggregated across newspaper-weighted readership education characteristics to a shock in newspaper sentiment. Shaded areas represent 95% confidence bands constructed with Newey-West standard errors.

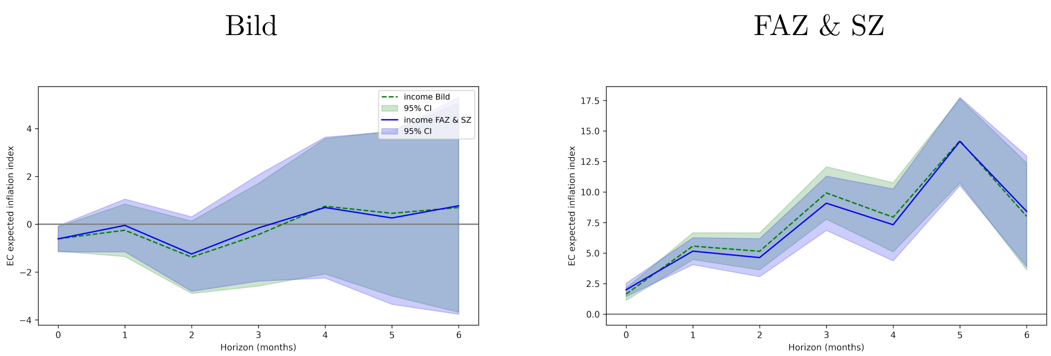

Figure 3.6-11: Newspaper sentiment shock by income level
Notes: Cumulative impulse responses of EC survey expectations disaggregated across newspaper-weighted readership income characteristics to a shock in newspaper sentiment. Shaded areas represent 95% confidence bands constructed with Newey-West standard errors.

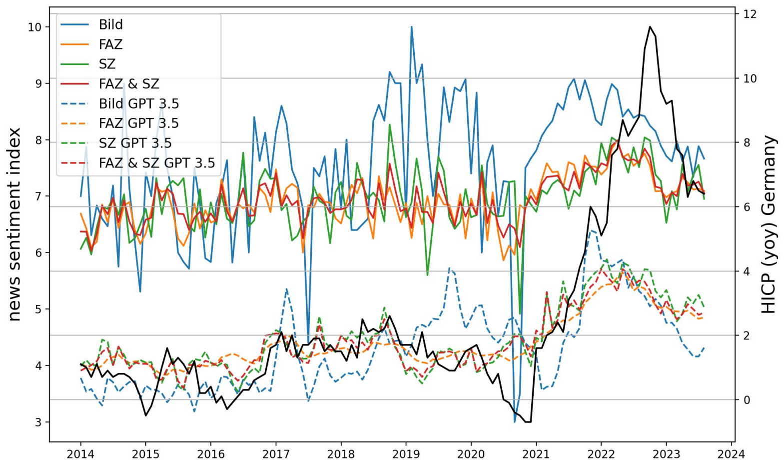

Figure 3.7-12: Robustness: ChatGPT-4o-mini
Notes: Newspaper indices generated with ChatGPT-3.5 (dashed lines) or ChatGPT-4o-mini (solid lines) against year-on-year inflation in Germany.

the model, gpt-4o-mini, for which the training sample ends in October 2023 and therefore covers my whole newspaper sample.

Despite the different, more advanced model and the longer training period, the final newspaper indices show a sizable positive correlation with the baseline indices (Bild: 0.54, FAZ: 0.71, and SZ: 0.55). Figure 3.7-12 shows the resulting indices constructed by taking the mean over the articles per month. The solid lines represent the ChatGPT-4 indices. The shape is still similar to the ChatGPT-3.5 version (dashed lines), FAZ and SZ are again closely aligned, but the main difference is the level of the indices. The average level when using gpt-4o-mini is much higher than in the previous version. One reason for this could be the inclusion of the latest high-inflation period in the LLM's inflation set. But other model-specific factors might also drive these differences.

Regardless of these level differences, the indices perform equally well when used to predict readership-weighted household expectations by education and income level. Table 3.7-3 shows the relative RMSE as constructed in section 3.4.2. The results display the same pattern of heterogeneous forecasting powers as the baseline model.

3.7.2 Sample splits

Since Weber et al. (2023) find that households pay more attention to news about inflation when inflation is high, it is possible that my results are influenced by the euro zone's recent high inflation period starting in 2021. To investigate this further, I report sample splits of the RMSE ratios with an out-of-sample period from 2017 until March 2020, and from April 2020 to August 2023, which coincides with the COVID-19 pandemic.

Table 3.7-4 presents the pre-inflation RMSE ratios on inflation expectations for the first sample split. With respect to education, there is even stronger heterogeneity between the individual newspapers. The Bild index fares worse in predicting the FAZ & SZ readership's expectations by education level as the RMSE ratios are smaller compared to the baseline.

	Education				**Income**			
h	Bild	FAZ	SZ	FAZ & SZ	Bild	FAZ	SZ	FAZ & SZ
1	1.0	1.0	1.0	1.0	1.0	1.01	1.0	1.0
2	1.01	0.98	0.99	0.98	1.01	1.01	1.01	1.01
3	1.02	0.98	0.99	0.98	1.03	1.04	1.04	1.04
4	1.03	0.95	0.96	0.96	1.03	1.04	1.04	1.04
5	1.02	0.94	0.95	0.95	1.03	1.04	1.04	1.04
6	1.03	0.94	0.95	0.94	1.04	1.05	1.05	1.05

Table 3.7-3: RMSE ChatGPT-4o-mini

Notes: This table presents the relative RMSE of the FAZ & SZ indices compared to the Bild index. Indices are constructed with gpt-4o-mini. Values above 1 indicate that adding the Bild index to the benchmark model outperforms the other newspaper indices. The left panel shows results for predicting households' inflation expectations, weighted by the education characteristics of the average reader of the respective newspapers (Bild, FAZ, SZ, and FAZ & SZ combined). The right panel displays the equivalent results weighted by income characteristics. RMSEs are calculated on a sample spanning from 2017M01 to 2023M08.

	Education				**Income**			
h	Bild	FAZ	SZ	FAZ & SZ	Bild	FAZ	SZ	FAZ & SZ
1	1.06	1.03	1.04	1.03	1.04	1.01	1.01	1.01
2	1.06	0.79	0.82	0.81	1.06	1.06	1.05	1.05
3	1.12	0.84	0.87	0.86	1.13	1.16	1.16	1.16
4	1.15	0.86	0.89	0.87	1.19	1.27	1.26	1.26
5	1.23	0.98	1.01	0.99	1.26	1.38	1.38	1.38
6	1.29	1.1	1.13	1.12	1.32	1.48	1.48	1.48

Table 3.7-4: RMSE on 2017M01-2020M03 sub-sample

Notes: This table presents the relative RMSE of the FAZ & SZ indices compared to the Bild index. Values above 1 indicate that adding the Bild index to the benchmark model outperforms the other newspaper indices. The left panel shows results for predicting households' inflation expectations, weighted by the education characteristics of the average reader of the respective newspapers (Bild, FAZ, SZ, and FAZ & SZ combined). The right panel displays the equivalent results weighted by income characteristics. RMSEs are calculated on a sample spanning from 2017M01 to 2020M03.

In contrast, as shown in Table 3.7-5, for the second half of the sample, the Bild index performs slightly worse in predicting the expectations of households with the education level of a typical Bild reader, but still better than when predicting those of FAZ & SZ readers. Therefore, the Bild index outperforms FAZ & SZ when predicting expectations weighted by the income distribution of Bild readers. Overall, the patterns of heterogeneity are similar to the baseline results for the whole sample, even though this analysis suggests that these results might be driven by the pre-2020 period.

3.8 Conclusion

In this paper, I extract inflation sentiments at the article level from three major German daily newspapers—Bild, FAZ, and SZ—using the large language model ChatGPT-3.5 to construct inflation indices. I then analyze the informativeness of these indices in relation to the inflation expectations of different socioeconomic groups, exploring their role in shaping those expectations. In terms of forecasting power, the Bild indices outperform in predicting

	Education				**Income**			
h	Bild	FAZ	SZ	FAZ & SZ	Bild	FAZ	SZ	FAZ & SZ
1	1.01	1.01	1.01	1.01	1.02	1.01	1.01	1.01
2	0.99	0.95	0.96	0.96	1.01	0.99	0.99	0.99
3	0.99	0.92	0.93	0.93	1.01	0.99	1.0	0.99
4	0.99	0.9	0.92	0.91	1.02	1.0	1.0	1.0
5	0.99	0.88	0.9	0.89	1.01	1.0	1.0	1.0
6	0.98	0.87	0.89	0.88	1.02	1.0	1.01	1.0

Table 3.7-5: RMSE on 2020M04-2023M08 sub-sample

Notes: This table presents the relative RMSE of the FAZ & SZ indices compared to the Bild index. Values above 1 indicate that adding the Bild index to the benchmark model outperforms the other newspaper indices. The left panel shows results for predicting households' inflation expectations, weighted by the education characteristics of the average reader of the respective newspapers (Bild, FAZ, SZ, and FAZ & SZ combined). The right panel displays the equivalent results weighted by income characteristics. RMSEs are calculated on a sample spanning from 2020M04 to 2023M08.

expectations among households that align with the income and education distribution of typical Bild readers, whereas the opposite is true for the FAZ and SZ indices. Furthermore, the FAZ and SZ sentiment indices exhibit a stronger co-movement with actual inflation in Germany.

Next, I employ local projections to estimate the impulse responses of the three newspaper indices to a monetary policy shock. Once again, heterogeneity in responses emerges: Bild sentiment shows an immediate and sizable decrease, suggesting a less negative tone towards the inflation environment. In contrast, the SZ index response is delayed by a few days and is more muted, while the growth of the FAZ index does not turn negative after the initial impact. This pattern may indicate that monetary policy is more effective in cooling the inflationary sentiment reflected in Bild and SZ, while FAZ sentiment remains largely unaffected.

When narrowing the focus to articles that explicitly mention monetary policy, all three newspapers' sentiments show a decline in response to a contractionary monetary policy shock. These findings suggest that monetary policy can effectively influence press sentiment, particularly in tabloid reporting. A key policy implication is that central banks should prioritize engaging with tabloids in their communication strategies. This could help counteract unfounded negative reporting, such as during the 2019 episode, and enhance the impact of monetary policy through more targeted messaging in tabloid media.

The final source of heterogeneity I examine is the varying responsiveness of different household types to shocks in the sentiment indices of the three newspapers. Consistent with my earlier findings, the expectations of more educated households are most influenced by shocks in FAZ and SZ sentiment, reflecting the characteristics of their respective readerships. For Bild, this heterogeneity is less pronounced, although there is still a weaker effect of FAZ and SZ sentiment shocks on the expectations of households with education levels similar to the average Bild reader.

Overall, this paper demonstrates that inflation sentiment indicators based on newspaper texts not only carry useful real-time information for forecasting but also impact different segments of society differently when shocked. They thus present themselves as valuable tools for short-run predictions of heterogeneous household inflation expectations. My results indicate that yellow press reporting is a better predictor of household inflation expectations during periods of stable inflation, while signals from more reputable newspapers serve as equally good proxies for inflation expectations during times of higher inflation, when they might reach a larger audience.

Looking ahead, I expect such alternative measures of inflation expectations to become more accurate and potent with advancements in text analytics using LLMs. Until then, this paper shows that text-based data generated by a simple zero-shot classification approach without any further pre-training carries valuable information for forecasting.

3.9 Appendix

3.9.1 Search queries

This section contains the detailed German versions of the search queries I use in the SZ and FAZ news archives, on LexisNexis and mediacloud to filter inflation related newspaper articles. On LexisNexis, Bild newspaper articles are available from 2010 to 2017. The Bild query uses a slightly different vocabulary such as *Zins-Hammer* (exaggeration for interest rate increase) or *Teuershock* (i.e., a shock that increases prices by a lot) which is unique to this newspaper.

SZ:

Inflation* OR "niedrige Preise" OR "hohe Preise" OR Disinflation OR Deflation OR Geldpolitik OR Lohn-Preis-Spirale OR Zentralb* OR Stagflation OR Energiepreis* OR Geldsteuerung OR EZB OR Kaufkraftverlust OR Notenb* OR Leitzins OR Kostenanstieg OR Preissenkung* OR Zinswende OR Preisexplosion OR Preisanstieg OR Lebenshaltungskosten OR Teuerung OR Preisindex AND Deutschland NOT Verlag* NOT Preisträger* NOT Preissieger* NOT *Film* NOT Oscar NOT Preisverleihung* NOT Preisgeld* NOT Billigung NOT Trainer* NOT Autor* NOT Fußball* NOT preisg* NOT Trump

FAZ:

(RESSORT: (Wirtschaft) OR RESSORT: (Politik) OR RESSORT: (Finanzmarkt) OR RESSORT: (Immobilienmarkt) NOT RESSORT: (Rhein-Main-Zeitung)) AND QUELLENANGABE: (Frankfurter Allgemeine Zeitung) AND (TEXT: (Inflation*) OR TEXT: (niedrige Preise) OR TEXT: (hohe Preise) OR TEXT:(Disinflation) OR TEXT: (Deflation) OR TEXT: (Geldpolitik) OR TEXT: (Lohn-Preis-Spirale) OR TEXT: (Zentralb*) OR TEXT: (Stagflation) OR TEXT: (Energiepreis*) OR TEXT: (Geldsteuerung) OR TEXT: (EZB) OR TEXT: (Kaufkraftverlust) OR TEXT: (Notenb*) OR TEXT: (Leitzins) OR TEXT: (Kostenanstieg) OR TEXT: (Preissenkung*) OR TEXT: (Zinswende) OR TEXT: (Preisexplosion) OR TEXT: (Preisanstieg) OR TEXT: (Lebenshaltungskosten) OR TEXT: (Teuerung) OR TEXT: (Preisindex) NOT (Verlag*)) AND ERSCHEINUNGSJAHR: [2010 TO 2023] AND LAND: (Deutschland) AND (QUELLENANGABE: (Frankfurter Allgemeine Zeitung) OR QUELLENANGABE: (FAZ.NET)) NOT SERIENTITEL: (Politische Bücher) NOT SERIENTITEL: (Wirtschaftsbücher)

LexisNexis:

Zins-Hammer or steigende Preise or fallende Preise or Preischaos or Bezinpreis* or Teuerschock or Inflation* or niedrige Preise or hohe Preise or Disinflation or Deflation or Geldpolitik or Lohn-Preis-Spirale or Zentralbank or Stagflation or Energiepreis* or Geldsteuerung or EZB or Kaufkraftverlust or Notenb* or Leitzins or Kostenanstieg or Preissenkung* or Zinswende or Preisexplosion or Preisanstieg or Lebenshaltungskosten or Teuerung or Preisindex and not Verlag* and not Preisträger* and not Preissieger* and not *Film* and not Oscar and not Preisverleihung* and not Preisgeld* and not Billigung and not Trainer* and not Autor* and not Fußball* and not preisge* and not Trump

mediacloud:

(Zins-Hammer OR steigende Preise OR fallende Preise OR Preischaos OR Bezinpreis* OR Teuerschock OR Inflation* OR niedrige Preise OR hohe Preise OR Disinflation OR Deflation OR Geldpolitik OR Lohn-Preis-Spirale OR Zentralbank OR Stagflation OR Energiepreis* OR Geldsteuerung OR EZB OR Kaufkraftverlust OR Notenban* OR Leitzins OR Kostenanstieg OR Preissenkung* OR Zinswende OR Preisexplosion OR Preisanstieg OR Lebenshaltungskosten OR Teuerung OR Preisindex) AND NOT (Verlag OR Preisträger* OR Preissieger* OR Film OR Oscar OR Preisverleihung* OR Preisgeld* OR Billigung OR Trainer* OR Autor* OR Fußball* OR preisge* OR Trump OR Tor*)

3.9.2 ChatGPT's interpretation of Bild expressions

ChatGPT is indeed able to understand the special language of the Bild newspaper and to capture the intended tone of its expressions. I demonstrate this by asking ChatGPT to give explanations and interpretations of the tone of the in Table 3.3-1 highlighted German terms, using the following prompt: *"Explain the term [...] in the context of the newspaper headline [...] and describe the tone this expression implies."* Table 3.9-1 presents the results of this exercise. In all cases, ChatGPT accurately captures the meaning and provides a suited interpretation of the German expressions.

Term	**ChatGPT Explanation**	**Interpretation of tone**
Teuer-Schock	The term "Teuer-Schock" is a German compound term that combines "Teuer," meaning expensive or costly, with "Schock," meaning shock. As a compound term, "Teuer-Schock" is often used in the context of economic or financial discussions to describe a situation where there is a sudden, unexpected, and significant increase in prices or costs, leading to a shock-like impact on individuals, businesses, or the economy as a whole.	The tone implied by "Teuer-Schock" is generally one of concern, surprise, and possibly criticism. It suggests a negative or adverse situation in which the cost or price surge is seen as problematic, potentially disruptive, or challenging for those affected. The term is used to convey the idea that the increase in prices is not just a regular or anticipated occurrence but rather an event that has a notable and adverse impact, similar to the impact of a shock. In news headlines or economic discussions, "Teuer-Schock" is often employed to draw attention to the severity and unexpectedness of the price increase, emphasizing its potential consequences and the need for attention or action.
oberster Geld-Hüter	The term "oberster Geld-Hüter" is a German expression that can be translated to "supreme money guardian" or "chief monetary guardian" in English. This term is often used to refer to a person who holds a key position of authority or responsibility in monetary or financial matters, such as a central bank governor, finance minister, or another high-ranking official overseeing economic policy.	The tone implied by "oberster Geld-Hüter" is one of authority, responsibility, and sometimes a degree of formality. It conveys the idea that the individual being referred to is at the highest level of decision-making and influence regarding financial and monetary policies. The term is often used in news articles, economic analyses, or discussions to highlight the significance and influence of the person in question, emphasizing their role in shaping and managing the country's economic policies. The use of such a term may carry a sense of expectations regarding the individual's ability to make important decisions and steer the country's economic course, reflecting a certain level of accountability and authority associated with the position.

Term	ChatGPT Explanation	Interpretation of tone
Madame Inflation	In the context of the euro zone, ”Madame Inflation” could potentially refer to Christine Lagarde, the current President of the European Central Bank (ECB). Christine Lagarde holds a significant role in shaping and implementing monetary policies that influence economic conditions, including inflation rates, within the euro zone.	When using a term like ”Madame Inflation” to refer to Christine Lagarde, it suggests a personification of responsibility for inflationary pressures or challenges within the euro zone. The term implies a critical or at least skeptical stance toward her policies or actions, holding her accountable for the economic conditions that may lead to inflation. The use of ”Madame” adds a formal and somewhat sardonic tone, suggesting a level of authority and influence associated with Christine Lagarde’s position. Overall, it implies a critical perspective on the individual in charge of monetary policies in the euro zone, particularly in the context of concerns or shocks related to inflation.
Luxus-Lagarde	In the context of the newspaper headline ”Madame Inflation; Luxus-Lagarde macht Sparer und Rentner arm,” the term ”Luxus-Lagarde” is a metaphorical expression that combines ”Luxus” (luxury) with the surname ”Lagarde,” referring to Christine Lagarde, the President of the European Central Bank (ECB).	The term ”Luxus-Lagarde” suggests a critical tone, implying that Christine Lagarde is associated with a lifestyle of luxury or extravagance. By combining her name with ”Luxus,” it conveys the idea that Lagarde is perceived as leading a lifestyle that may be considered opulent or out of touch with the financial challenges faced by the general population, specifically savers and retirees mentioned in the headline. The tone of ”Luxus-Lagarde” in this context is likely critical, implying a disapproval of Lagarde’s perceived lifestyle in connection with the economic challenges faced by savers and retirees. It may be used as a form of criticism, suggesting a disconnect between Lagarde’s perceived luxurious lifestyle and the economic struggles of certain groups in society. This kind of expression is often employed in headlines or articles to convey a strong opinion or perspective on a public figure’s actions or policies.
Zins-Hammer	In the context of the newspaper headline ”Zins-Hammer! EZB erhöht auf 1,25 Prozent” (Interest Rate Hammer! ECB Increases to 1.25 Percent), the term ”Zins-Hammer” is a metaphorical expression. It combines ”Zins” (interest rate) with ”Hammer,” suggesting a forceful and impactful event related to interest rates.	The term ”Zins-Hammer” implies a tone of surprise, emphasis, and potentially a sense of urgency or significance. It suggests that the increase in interest rates is not just a routine adjustment but a substantial and noteworthy event that is likely to have important consequences or effects. The use of ”Hammer” adds a dramatic and attention-grabbing element to the headline, indicating that the interest rate change is being portrayed as a powerful or forceful action. It might also imply that the change in interest rates could have considerable impacts on various stakeholders, such as borrowers, savers, and the overall economy.

Table 3.9-1: Explanations and interpretations of Bild-typical expressions

3.9.3 Matching readership data with EC survey bins

The table in this section provides a description of how I match the EC survey’s household categories with those of the *Media-Analyse Tageszeitungen 2022*. The matching of the income bins is very approximate and is not based on the actual income distribution of the surveyed households.

	Media-Analyse survey bins	EC survey bins
Income Levels	net household income: less than 1,000€	1st income quartile
	1,000 - 1,250€	
	1,250 - 1,500€	2nd income quartile
	1,500 - 2,000€	
	2,000 - 2,500€	3rd income quartile
	2,500 - 3,000€	
	more than 3,000€	4th income quartile
Education Levels	Students in general education	primary education
	Lower secondary school diploma without apprenticeship	
	Lower secondary school diploma with apprenticeship	secondary education
	Secondary school without university entrance qualification	
	College entrance qualification without university degree	further (tertiary) education
	College entrance qualification with university degree	
Age Groups	16 - 39	16 - 29 (age = 1)
		30 - 49 (age = 2)
	40 - 59	50 - 64 (age = 3)
	60+	65+ (age = 4)

Table 3.9-2: Matching of Media-Analyse survey bins and EC survey bins across different income levels, education levels, and age groups

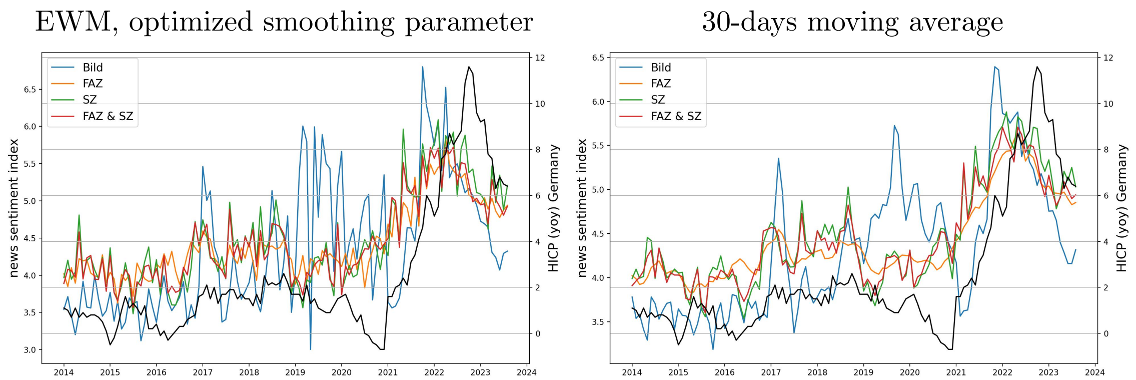

Figure 3.9-1: Alternative aggregation methods

Notes: This figure shows the resulting newspaper indices from alternative aggregation methods. The indices in left panel are constructed by computing the EWM with an optimized smoothing parameter and the panel on the right the indices constructed as the moving average of a 30-days window. All indices are in monthly frequency.

3.9.4 Alternative index construction

This section presents the results of the relative forecast errors when using alternative ways to aggregate the article-level sentiment scores to a monthly index. First, I obtain mean scores per day to then compute moving averages (MA) over a window of 10, 30 or 60 days (MA-10, MA-30 & MA-60), and then aggregate these daily indices to monthly frequency by taking the average for each month. As a second option, and analogously to Angelico et al. (2022), I compute exponentially weighted averages over the daily mean scores using three alternative smoothing parameters: 0.1, 0.3 and the one that minimizes the sum-of-squared in-sample forecast errors, which is close to 0.1 (indices are denoted by EWM, $a = 0.1$, EWM $a = 0.3$ and EWM, optimal a in Table 3.9-3). The results are of a similar magnitude as those obtained with the baseline index (i.e., the simple monthly mean over the daily sentiment scores), and also exhibit the same pattern in heterogeneity across newspapers (see Figure 3.9-1).

	Education				Income			
h	Bild	FAZ	SZ	FAZ & SZ	Bild	FAZ	SZ	FAZ & SZ
	a) EWM, $a = 0.1$							
1	1.02	1.0	1.0	1.0	1.02	1.01	1.01	1.01
2	1.03	0.96	0.97	0.96	1.03	1.02	1.02	1.02
3	1.04	0.94	0.95	0.95	1.05	1.03	1.03	1.03
4	1.05	0.95	0.96	0.96	1.07	1.05	1.05	1.05
5	1.07	0.97	0.98	0.97	1.09	1.08	1.08	1.08
6	1.08	0.98	1.0	0.99	1.09	1.09	1.1	1.09
	b) EWM, $a = 0.3$							
1	1.04	1.03	1.03	1.03	1.04	1.04	1.04	1.04
2	1.03	0.95	0.97	0.96	1.04	1.03	1.03	1.03
3	1.04	0.94	0.95	0.94	1.05	1.04	1.04	1.04
4	1.06	0.94	0.96	0.95	1.08	1.07	1.07	1.07
5	1.08	0.95	0.97	0.96	1.09	1.1	1.1	1.1
6	1.08	0.96	0.97	0.96	1.1	1.11	1.11	1.11
	c) EWM, optimal a							
1	1.05	1.04	1.04	1.04	1.04	1.04	1.04	1.04
2	1.03	0.95	0.96	0.96	1.03	1.03	1.03	1.03
3	1.04	0.93	0.95	0.94	1.05	1.05	1.05	1.05
4	1.06	0.94	0.95	0.94	1.08	1.08	1.08	1.08
5	1.07	0.94	0.96	0.95	1.09	1.1	1.11	1.11
6	1.08	0.94	0.96	0.95	1.1	1.12	1.12	1.12
	d) MA-10							
1	1.03	1.02	1.02	1.02	1.03	1.03	1.03	1.03
2	1.02	0.96	0.97	0.96	1.03	1.02	1.02	1.02
3	1.03	0.94	0.95	0.94	1.05	1.04	1.04	1.04
4	1.05	0.94	0.96	0.95	1.07	1.06	1.06	1.06
5	1.06	0.95	0.96	0.96	1.08	1.08	1.08	1.08
6	1.06	0.95	0.97	0.96	1.07	1.09	1.09	1.09
	e) MA-30							
1	1.01	0.98	0.99	0.99	1.02	1.01	1.01	1.01
2	1.02	0.96	0.97	0.96	1.03	1.01	1.01	1.01
3	1.03	0.94	0.96	0.95	1.04	1.02	1.02	1.02
4	1.04	0.96	0.98	0.97	1.06	1.03	1.03	1.03
5	1.06	0.99	1.0	1.0	1.08	1.06	1.06	1.06
6	1.08	1.02	1.03	1.03	1.1	1.08	1.08	1.08
	f) MA-60							
1	1.02	0.99	0.99	0.99	1.02	1.01	1.01	1.01
2	1.04	0.96	0.97	0.97	1.04	1.02	1.02	1.02
3	1.03	0.92	0.94	0.93	1.05	1.01	1.01	1.01
4	1.04	0.91	0.92	0.92	1.06	1.02	1.02	1.02
5	1.06	0.93	0.95	0.94	1.08	1.04	1.04	1.04
6	1.09	0.97	0.98	0.98	1.1	1.07	1.07	1.07

Table 3.9-3: RMSE: different aggregation methods

Notes: This table presents the RMSE ratios of the FAZ & SZ RMSEs over the Bild RMSEs when predicting the 12-months ahead inflation expectations weighted by each newspaper's readership composition. Panels a) to c) contain the results constructed with the exponentially weighted mean (EWM) with smoothing parameters a taking the values 0.1, 0.3, or an optimized value. Panels d) to f) contain results for indices constructed as the moving average over a window of 10, 30, or 60 days, collapsed at monthly frequency by taking the mean.

	Education				Income			
h	Bild	FAZ	SZ	FAZ & SZ	Bild	FAZ	SZ	FAZ & SZ
1	1.0	1.0	1.0	1.0	0.99	0.98	0.98	0.98
2	1.0	0.92	0.93	0.93	1.01	1.0	1.0	1.0
3	1.02	0.9	0.92	0.91	1.03	1.03	1.02	1.02
4	1.02	0.89	0.91	0.9	1.04	1.05	1.04	1.04
5	1.03	0.89	0.91	0.9	1.05	1.06	1.06	1.06
6	1.04	0.9	0.92	0.91	1.06	1.08	1.08	1.08

Table 3.9-4: RMSE: individual results for FAZ

Notes: This table presents the relative RMSE of the FAZ index compared to the Bild index. Values above 1 indicate that adding the Bild index to the benchmark model outperforms the other newspaper indices. The left panel shows results for predicting households' inflation expectations, weighted by the education characteristics of the average reader of the respective newspapers (Bild, FAZ, SZ, and FAZ & SZ combined). The right panel displays the equivalent results weighted by income characteristics. RMSEs are calculated on a sample spanning from January 2017 to August 2023.

	Education				Income			
h	Bild	FAZ	SZ	FAZ & SZ	Bild	FAZ	SZ	FAZ & SZ
1	1.02	1.01	1.01	1.01	1.03	1.02	1.02	1.02
2	1.01	0.94	0.95	0.94	1.03	1.02	1.02	1.02
3	1.03	0.94	0.95	0.95	1.05	1.05	1.05	1.05
4	1.05	0.94	0.95	0.95	1.08	1.08	1.08	1.08
5	1.07	0.96	0.97	0.96	1.1	1.1	1.1	1.1
6	1.08	0.97	0.98	0.97	1.11	1.12	1.12	1.12

Table 3.9-5: RMSE: individual results for SZ

Notes: This table presents the relative RMSE of the SZ index compared to the Bild index. Values above 1 indicate that adding the Bild index to the benchmark model outperforms the other newspaper indices. The left panel shows results for predicting households' inflation expectations, weighted by the education characteristics of the average reader of the respective newspapers (Bild, FAZ, SZ, and FAZ & SZ combined). The right panel displays the equivalent results weighted by income characteristics. RMSEs are calculated on a sample spanning from January 2017 to August 2023.

3.9.5 RMSE results for FAZ and SZ individually

This section presents the relative RMSEs between the FAZ and SZ indices, and the Bild index, in the tables 3.9-4 and 3.9-5, respectively. The results are very similar and display the same patterns in heterogeneity as the baseine.

3.9.6 Controlling for the ECB rate in the sentiment shock identification

I alternatively control for the ECB's main refinancing operations rate in the news sentiment shock identification which follows Shapiro et al. (2022) in their approach. It might be the case that changes in the key interest rate itself might be more relevant for the general media than the high-frequency financial market surprises measured by Jarociński and Karadi (2020). As visible in Figures 3.9-2 and 3.9-3, this change does not affect the results.

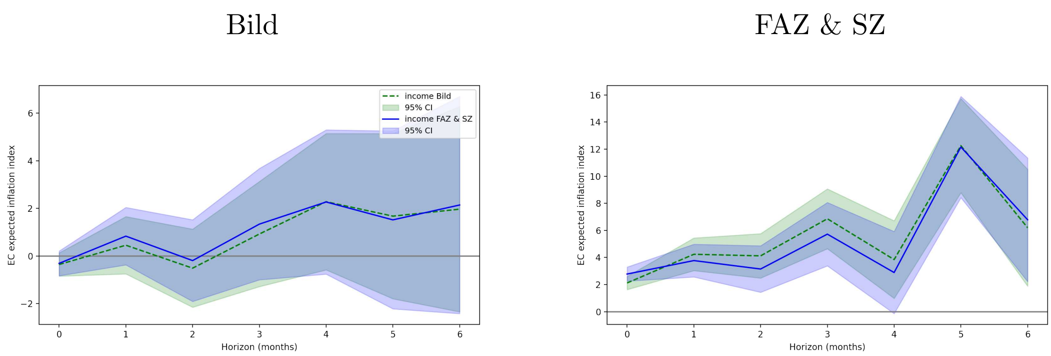

Figure 3.9-2: Controlling for ECB rate: education

Notes: Cumulative impulse responses of EC survey expectations disaggregated across newspaper-weighted readership education characteristics to a shock in newspaper sentiment. Shaded areas represent 95% confidence bands constructed with Newey-West standard errors.

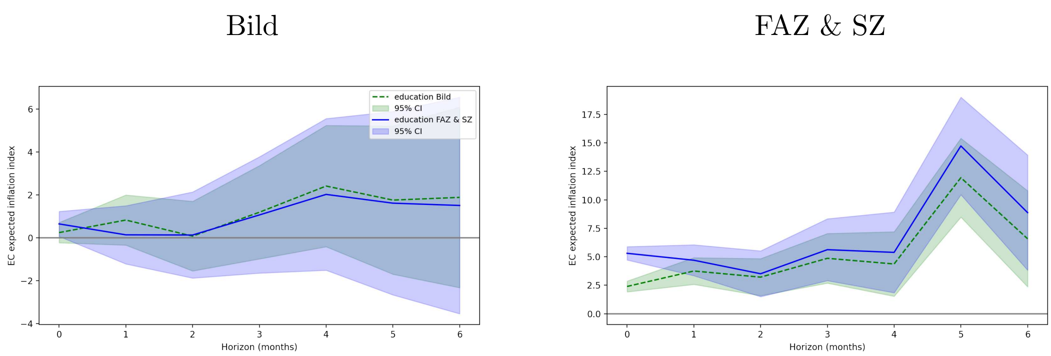

Figure 3.9-3: Controlling for ECB rate: income

Notes: Cumulative impulse responses of EC survey expectations disaggregated across newspaper-weighted readership income characteristics to a shock in newspaper sentiment. Shaded areas represent 95% confidence bands constructed with Newey-West standard errors.

4. Fiscal Narratives and Inflation

Authors

Sarah Arndt & Farah Tohme[1]

Abstract

This paper investigates how media narratives on fiscal policy shape households' inflation expectations. We collect a large corpus of newspaper articles reporting on fiscal policy from four major German newspapers spanning from 2006 to March 2025. Using a large language model (ChatGPT), we propose a strategy to automatically identify different fiscal narratives in text and construct narrative indices from these data. We then estimate the effect of these narrative indices on household inflation expectations and find that each has a positive and significant effect, varying in size. Lastly, we measure how fiscal narratives affect the transmission of a government spending shock to the economy and find that some narratives have an amplifying effect, while others dampen the impact.

Keywords

Fiscal policy, narratives, government spending shock, ChatGPT

4.1 Introduction

The recent surge in inflation in the US and Europe that began in 2021 has been accompanied by a high debt level due to deficit-financed fiscal interventions introduced to cushion the economic impact of the COVID-19 pandemic. Under these circumstances, the question of the relationship between fiscal policy and inflation gained renewed interest in the literature. Barro and Bianchi (2024), for example, show empirically that in the aggregate of 17 euro area countries in 2020–2023, inflation responded positively to a composite government spending variable. Further, Bassetto and Miller (2025) build a model in which they show how inflation and fiscal deficits can suddenly become connected through an information channel, generating sudden inflation. In a similar spirit, De Fiore et al. (2024) estimate the effect of the announcement of three major US stimulus packages (the American Rescue Plan, the CARES Act, and the Tax Cuts and Jobs Act) and find that financially literate households shift their inflation expectations upward.

In this paper, we shed more light on the aforementioned information channel that seems to play a significant role in the effects of fiscal policy shocks. Unlike experts or those active in financial markets, households typically do not directly observe fiscal policy interventions, but instead rely on intermediaries such as the news media to make sense of economic

[1] **Acknowledgements:** This work has greatly benefited from conversations with Klaus Masuch and Marco Bassetto. We further thank Zeno Enders, Alfons Weichenrieder, and seminar participants at the Deutsche Bundesbank for helpful comments.

developments. In most cases, journalists do not exclusively report on factual events, but also interpret, frame, and narrate fiscal policy in ways that can shape public beliefs about government action, debt, and inflation. The majority of households do not know about economic theories that predict the effects of a government spending shock or fiscal deficits. If they establish a link between, for example, fiscal expansion and inflation, then most likely it is because they have learned about it in the news.

So far, Coibion et al. (2021) and Andrade et al. (2025) investigate the link between information on the fiscal stance and inflation expectations in survey experiments. The former find for the US that only information on the change in the government debt level leads to an upward revision of inflation expectations, while the latter observe for German households that only those who already believe the fiscal resources are constrained expect higher inflation when learning about a growing debt-to-GDP ratio. Again, the households' perceptions of the fiscal space might be strongly influenced by the media. As an illustration, if newspapers deem the fiscal space to be limited due to, in their view, excessive debt levels, agents consuming these news might follow this view. Based on Andrade et al. (2025), this would imply that any announcement of future fiscal expansion measures would lead to an increase in inflation expectations in these households, which might translate into actual inflation.

As an alternative to these information provision experiments, we aim to directly analyze households' information sets by proxying them with news articles. In this paper, we introduce a systematic approach using Artificial Intelligence (AI) to measure fiscal narratives in Germany. We answer the question of whether media outlets establish a link between fiscal expansion and inflation in their reporting, as well as measure their narratives on government debt. Building on these fiscal narratives, we analyze how this type of information affects household inflation expectations and estimate their impact on the transmission of government spending shocks to output and inflation in Germany.

Our approach is as follows: we collect a set of newspaper articles that contain at least one of a list of terms related to fiscal policy in order to focus only on those relevant for our analysis. We focus on four major news outlets in Germany: *Die Bild*, *Frankfurter Allgemeine Zeitung*, *Die Süddeutsche Zeitung*, and *die tageszeitung*. These newspapers represent different types of media: tabloid press, conservative broadsheets, and critical left-wing journalism. By analyzing this diverse selection, we aim to assess the differences in their reporting on fiscal policy, and also to capture as wide a range as possible of public opinions.

Our findings provide empirical insights into the heterogeneity of news and the importance of news to the formation of macro expectations. In doing so, we also contribute to the development of models that incorporate heterogeneous beliefs. First, we evaluate the effect of the fiscal narratives on household inflation expectations and find that all of the four narrative indicators we construct show significant effects on 12-month ahead household inflation expectations. Articles linking fiscal expansion to inflation raise expectations by 0.105 percentage points (pp), while those framing debt as unsustainable by 0.009 pp.

Narratives on whether the government should in- or decrease its spending raise inflation expectations by 0.034 and 0.013 pp, respectively.

In a second step, we investigate the non-linear effect of the fiscal narratives on the transmission of government spending shocks. News articles conveying that the government should decrease its spending dampen the effect of a fiscal shock on output and inflation, while the opposite holds true when they report the government should increase its spending. In this way, we can show that narratives not only affect subjective expectations, but also have an effect on macro variables.

Related Literature. Our paper is part of the expanding body of research that seeks to extract structured information from textual sources, such as sentiments or expectations. See, for example, Bybee (2023), Shapiro et al. (2022), Picault et al. (2022), Ellingsen et al. (2022), and Angelico et al. (2022). Another example is Aruoba and Drechsel (2024), who use information in written reports produced by Federal Reserve staff for the Federal Open Market Committee (FOMC) meetings to identify monetary policy shocks.

Recent studies apply natural language processing techniques likewise to investigate fiscal policy-related questions. Latifi et al. (2024) use all the parliamentary speeches in the German Bundestag from 1960 to 2021 and derive fiscal sentiments (i.e., expansionary vs. contractionary fiscal policy stance) from them using machine learning techniques. Furthermore, they use these fiscal sentiments to show that changes in fiscal sentiment cause a shift in government spending. In an extension, Tillmann (2025) estimates the effect of a monetary policy shock on fiscal sentiments. He uses the same fiscal sentiments measure as in Latifi et al. (2024) and finds that monetary tightening causes a drop in fiscal sentiments. Lieb et al. (2025) use all publicly released postwar communications from US presidents to project both the direction and extent of future changes in taxes. The aforementioned studies all focus on the policy makers themselves, using their speeches to uncover their fiscal stance. In contrast, our paper puts the spotlight on those affected by fiscal policy, namely the public. By leveraging news articles, we can learn about the public perception of fiscal policy and how it affects households' expectations.

Since Shiller (2017) there has been an emerging interest in studying narratives from an economic perspective. In Shiller (2020), Shiller coins the definition of economic narratives as *"stories that offer interpretations of economic events, or morals, or hints of theories about the economy"*. In that vein, Andre et al. (2021) and Andre et al. (2024) follow a survey approach to measure economic narratives of households about inflation by analyzing open-ended survey questions in which participants explain what they think caused the latest surge in inflation. Building on this analysis, Trebbi (2024) identifies supply and demand narratives within newspaper articles on inflation. Going beyond broad economic narratives, we adopt Shiller's definition as *"stories that offer interpretations of economic events"* and apply it specifically to the context of fiscal policy. In our setting, we understand fiscal narratives as the recurrent stories in the media that interpret, explain, or give meaning to fiscal policy actions and their economic consequences.

This paper is structured as follows: Section 4.2 describes our dataset and the method of measuring fiscal narratives in news texts. Section 4.3 then presents the resulting indices that we construct, and in Section 4.4, we estimate their effect on household inflation expectations, followed by analyzing how fiscal narratives affect the transmission of government spending shocks to GDP and inflation in Section 4.5. Finally, Section 4.6 concludes.

4.2 Data and measurement

4.2.1 Newspaper data

We compile a large dataset of daily news articles from four major German daily newspapers: *Frankfurter Allgemeine Zeitung* (FAZ), *Süddeutsche Zeitung* (SZ), *Bild*, and *die tageszeitung* (TAZ). Three of them (Bild, SZ, and FAZ) are the most widespread daily newspapers in Germany, and each appeals to a different segment of the German population. Bild is the most popular tabloid in Germany and recorded 1,252,394 sold units of its Sunday edition in the fourth quarter of 2024 (Informationsgemeinschaft zur Feststellung der Verbreitung von Werbeträgern e.V. (IVW), 2025). Its style of reporting is characterized by exaggerations and polarizing language. In contrast, FAZ and SZ are the most popular German quality newspapers, which are also known to be more conservative in their views. Each of them sold 512,520 and 460,789 units, respectively, in 2024Q4 (IVW, 2025). Lastly, TAZ stands for independent, critical, and left-leaning journalism but has the comparably smallest number of sold units in 2024Q4, namely 44,765 (IVW, 2025). With this diverse range of media outlets, we aim to cover the major streams of news reporting to represent as large a share of the German news-consuming population as possible.

Our sample spans from January 2006 for FAZ, SZ, and TAZ, and from August 2011 for Bild, to March 2025. We collect articles either manually from the publishers' online archives or by web scraping with the open-source platform mediacloud.org. In both cases, we limit the number of articles we collect by applying a set of keywords associated with fiscal topics and consider only those that contain at least one of the keywords. The list of keywords comprises all words of the three fiscal dictionaries reported in Appendix 4.7.1. With this procedure, we collect 514,940 articles in total. The number of articles per news outlet is listed in Table 4.2-1.

Figure 4.2-1 depicts the total monthly distribution of articles by newspaper across time (the blue line). It is apparent that the frequency with which these newspapers report on fiscal topics varies substantially over time. For a first exploration of the dataset, we apply a dictionary of terms related to fiscal expansion and another one on fiscal consolidation (see Appendix 4.7.1) to the text corpus and filter those articles containing at least one of these terms. The results of this exercise are also shown in Figure 4.2-1. In general, newspapers tend to report more on topics related to fiscal expansion (orange line), but especially in the past year, starting in 2024, fiscal consolidation (green line) has become more prevalent in news reporting.

4.2.2 Textual analysis of fiscal narratives

We conduct the main textual analysis with the large language model ChatGPT-4o-mini by OpenAI. Our goal is to measure the prevalence of certain fiscal narratives in the media that

	Bild	FAZ	TAZ	SZ
Sample Span	2011-2025	2006-2025	2006-2025	2006-2025
Total	47,230	58,290	77,783	331,637
Fiscal Expansion	21,476	27,852	36,461	154,422
Fiscal Consolidation	3,700	9,214	10,921	36,810

Table 4.2-1: Number of articles per media outlet

Notes: The sample starts in January 2006, for FAZ, SZ and TAZ and in August 2011 for Bild. It ends in March 2025.

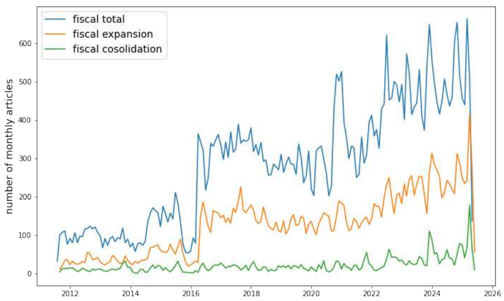

(a) Bild

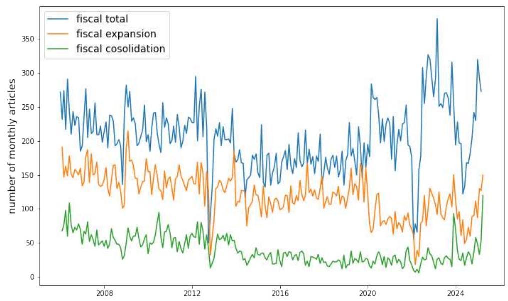

(b) FAZ

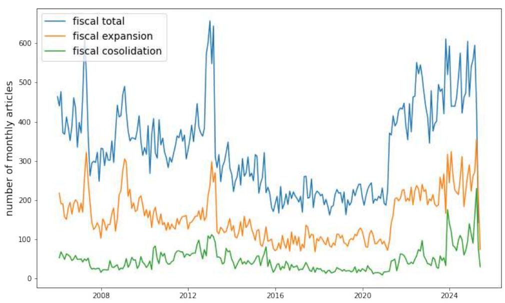

(c) TAZ

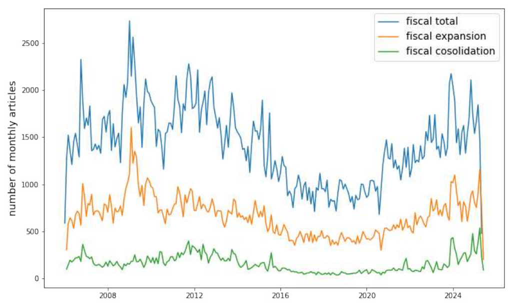

(d) SZ

Figure 4.2-1: Number of monthly articles across time by newspaper

Notes: This figure presents the number of monthly articles when filtering for keywords associated with fiscal expansion (orange), fiscal consolidation (green), or general fiscal topics (blue) for each of the media outlets Bild, FAZ, SZ, and TAZ.

allow for conclusions on the link between fiscal expansion and inflation. The narratives we want to measure are the following four:

1. Fiscal expansion will lead to inflation (referred to as *expansion leads to inflation* or *expansion* $\Rightarrow$*inflation* narrative).
2. The German government debt level is unsustainably high (referred to as *unsustainable debt* narrative).
3. The German government should spend more and conduct fiscal expansion (referred to as *spending* narrative).
4. The German government should reduce its spending and conduct fiscal consolidation (referred to as *saving* narrative).

This choice of narratives is partly motivated by Bassetto and Miller (2025), who report the numbers of Google searches for the term inflation in connection with fiscal policy-related terms. They show that these searches especially rose starting from 2022. Further, they construct a model that is based on the assumption that at some point, agents start to acquire more information on the government debt level. If more agents have this information, a transition into a regime in which expansion will generate inflation becomes more likely, in the case that the government is fiscally constrained[2]. These model assumptions and predictions can be captured by narratives that on the one hand link inflation and fiscal expansion (i.e., narrative 1), and on the other hand by narratives that convey that the government is risking becoming fiscally constrained (i.e., narratives 2 and 4). Lastly, we are also interested in the opposite case, i.e., the spending narrative.

To measure the four narratives, we design three prompts with which we automatically classify the text corpus using the OpenAI batch application programming interface (API). In the first prompt, to identify the fiscal expansion leads to inflation narrative, we ask the model to determine whether the article establishes a narrative that fiscal expansion (e.g., government spending increases, tax cuts) will lead to inflation. In the second prompt to measure the unsustainable debt narrative, we ask the model: "Does the article suggest that government debt levels in Germany are unsustainable?" And in the third prompt, we detect the spending/saving narratives by asking if the presented article conveys an opinion on whether the German government should save more or increase its spending. The exact prompts can be found in Appendix 4.7.2. For each prompt, we receive a binary yes/no answer or the answers "save" and "spend" which we use to classify the articles according to the narratives.

To then create indicators out of the raw narrative classification counts, we create a monthly weighted sum by popularity of the news outlet. This way, we account for their differences in dissemination and their likelihood of being seen by the public. The most popular outlet in our news set is the Bild newspaper with approximately 7.35 million readers per daily edition. SZ, FAZ, and TAZ each reach 1.3, 0.94, and 0.296 million readers per edition (agma, 2022)[3]. These single indicators then help assess the public's views and perceptions

[2]Fiscal constraints can also be politically driven, for example by limiting the possibility to increase taxes or to take on higher debt (Bassetto and Miller, 2025). In Germany, such a political instrument limiting government borrowing and ensuring fiscal discipline is the so-called "debt break", a constitutional rule introduced in 2009.

[3]Numbers on newspaper readerships are obtained from Arbeitsgemeinschaft Media-Analyse e.V. (agma).

on fiscal policy and its relation to inflation.

4.3 Indices of fiscal narratives

We present the results of the newspaper article classification for each media outlet and narrative in Figure 4.3-2. Each graph shows the absolute monthly numbers of articles from each newspaper that were classified into one of the four narratives. The sample of Bild articles starts in August 2011 and the others in January 2006. Three different types of variation are evident: substantial time variation within each narrative, variation across narratives, and heterogeneity between the different newspapers.

Comparing the number of articles classified into the unsustainable debt narrative across newspapers, there are similar peaks, although their relative magnitudes differ. According to our measurement, this narrative was especially prevalent in the years of the Great Financial Crisis (GFC) and later in the years of the euro crisis, with peaks in 2010 and 2012. In the mid-2010s, when the European sovereign debt was tackled, the unsustainable debt narrative lost importance, but gained popularity later on starting with the COVID-19 pandemic in 2020, with another surge in 2022 when Russia invaded Ukraine. Comparing the newspapers, for TAZ and SZ the peaks during the GFC and the invasion of Ukraine are similarly high, while the emphasis on this narrative in the FAZ was more pronounced during the GFC and euro crisis.

In the case of the spending vs. saving narratives, the differences between newspapers become even bigger. While the saving narrative seems to dominate the FAZ reporting during the GFC and the sovereign debt crisis, both are more on par in the TAZ and SZ, with larger peaks in the spending narrative. In contrast, for these newspapers the prevalence of the two narratives diverges towards the end of the sample, with the spending narrative becoming the leading narrative, especially for TAZ. For Bild, the spending narrative has its first bigger peak at the onset of the COVID-19 crisis, and after the invasion of Ukraine, both narratives surge and become equally important.

Lastly, the distribution of articles classified into the fiscal expansion leads to inflation narrative shows a similar pattern, especially across Bild, FAZ, and TAZ, with peaks in 2022. For all of these three media outlets, this type of narrative did not seem to play an important role in the years before. This picture is quite different in the case of SZ. Here, the prevalence of the fiscal expansion leads to inflation narrative was already similarly high during the GFC and euro crisis as it is at the end of the sample.

In Figure 4.3-3, we present the final fiscal narrative indices that we construct by weighting the individual newspaper results by their audience size. Overall, the majority of articles are classified into the spending narrative, directly followed by the saving narrative. Another general observation is that even though the number of articles related to fiscal policy was more or less stable across time with a small reduction in the mid-2010s (see Figure 4.2-1), this drop is much stronger in the fiscal narrative indices. We interpret this as an indication that the narratives capture something beyond the bare attention to fiscal policy.

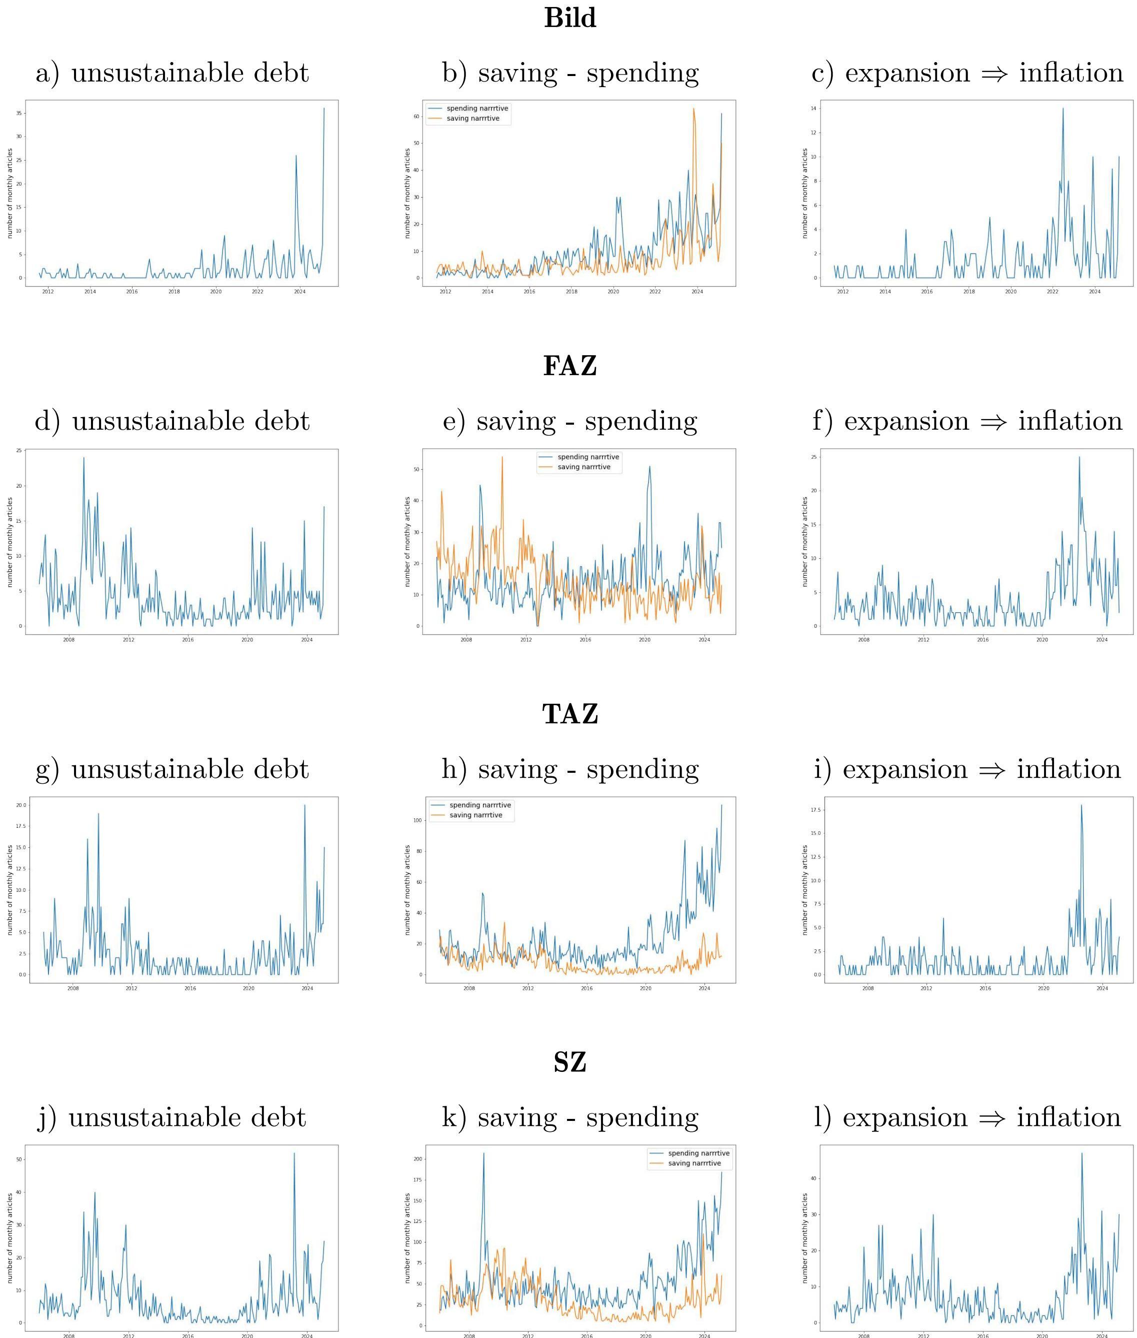

Figure 4.3-2: Narrative classification of articles by newspaper

Notes: The sample periods vary by news source.

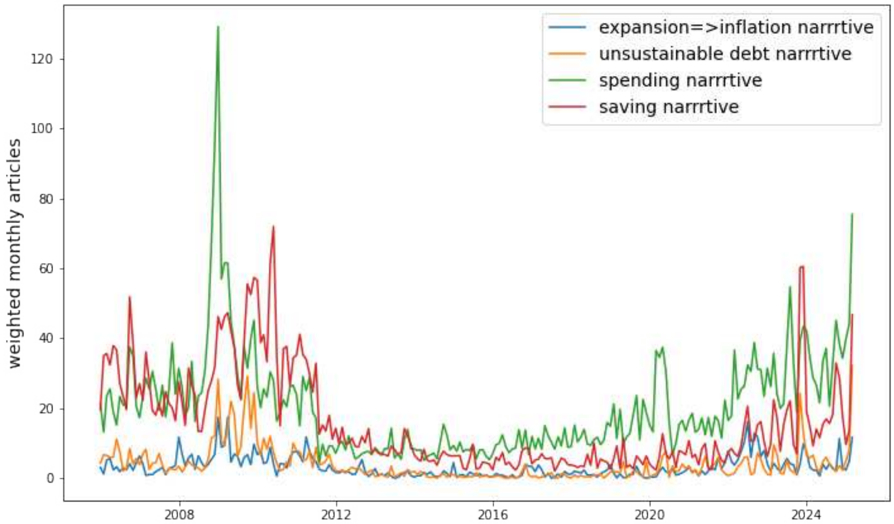

Figure 4.3-3: Final narrative indices

Notes: The final narrative indices are the by readership weighted newspaper indices.

4.4 Fiscal narratives' effect on inflation expectations

One mechanism by which fiscal narratives in the media potentially influence prices and other macro variables is by affecting households' inflation expectations. While it is quite common in theoretical models for inflation expectations to have an effect on actual inflation (see, e.g., Clarida et al., 1999; Smets, 2003; Woodford and Walsh, 2005), Moessner (2025) also provides empirical evidence for it. In the following, we test the hypothesis of whether the fiscal narratives we identify in the media increase household inflation expectations and therefore potentially affect actual inflation.

4.4.1 Data and model specification

For the estimation, we use the European Central Bank's (ECB) Consumer Expectations Survey (CES), which has been running since April 2020. Every month, the CES asks between 2,000 and 3,000 participating households in Germany for their quantitative inflation expectations: *"How much higher (lower) do you think prices in general will be 12 months from now in the country you currently live in? Please give your best guess of the change in percentage terms. You can provide a number up to one decimal place."* One can access the survey micro data online. It provides background information on participants such as gender, age, and education and income level. Our sample spans from April 2020 to December 2024.

To measure the effect of the prevalence of the four fiscal narratives in the media on households' 12-months-ahead inflation expectations, denoted by $E_{it}\pi_{t+12}$, we exploit the panel dimension of the CES and estimate the following model by OLS regression:

$$E_{it}\pi_{t+12} = c + \alpha_i + \beta Narrative^{j}_{t-1} + \gamma X_{t-1} + \epsilon_{it}. \tag{4.1}$$

Individual households are indexed by $i = 1, ..., n$ and the time dimension by $t = 1, ..., T$. We include household fixed effects (FE), denoted by α_i, to control for individual time-invariant

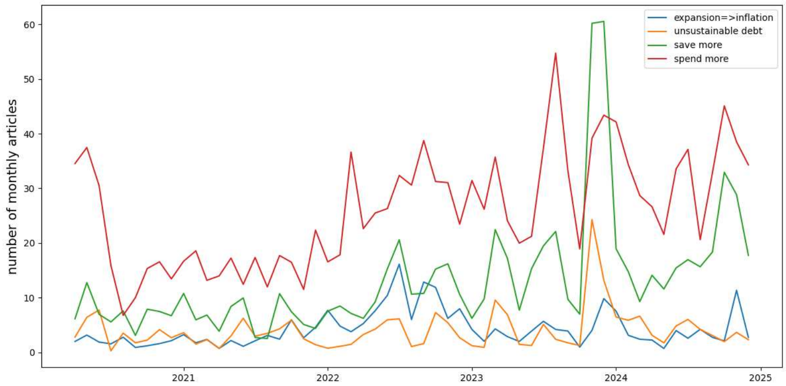

Figure 4.4-4: Aggregate fiscal narrative indices
Notes: This figure displays the four aggregate fiscal narrative indicators from April 2020 to December 2024. Each series represents the weighted number of monthly articles across FAZ, Bild, SZ, and TAZ falling into the i) fiscal expansion leads to inflation (blue), ii) public debt is unsustainably high (orange), iii) the government should increase savings (green), or iv) its spending (red) narrative.

household characteristics. In addition, following Bańkowska et al. (2021), we winsorize the sample at the 2% and 98% level per survey wave to prevent extreme responses from driving the results. To control for the macro environment, we include a set of macro variables as controls X_t. These include the German year-on-year inflation rate (π_t), industrial production (IP_t) as a monthly proxy for GDP, the ECB's policy rate (ECB_rate_t), and the total number of fiscal policy related articles, weighted by readership ($total\_fiscal_t$), to make sure effects are not driven by an overall increased attention to fiscal policy in the media. The variable $Narrative_t^j$ denotes one of the four fiscal narrative indicators j we described above and that are shown for the survey sample length in Figure 4.4-4. In Appendix 4.7.3, we also report results for the individual newspapers. Finally, ϵ_{it} is the error term.

The narrative indicators and control variables both enter the regression model with a lag, because some of the news might have been published after the household's survey participation, and also most macro data is released with a lag. The data collection for each wave of the CES typically starts on the first Thursday of a month and closes on the first Tuesday of the next month. This means, for example, in the September 2024 wave, the households were surveyed between September 5 and 30 (for further information see European Central Bank (ECB), 2024). On August 30, Eurostat published a first flash estimate of August inflation, and on September 18, 2024, the final inflation data for August was released.

All variables enter the estimation in levels—only the industrial production index is transformed to log differences. In addition, we assume that there is no reverse causality in the sense that households get influenced by the media, but the individual household has no impact on the number of articles sharing a specific fiscal policy narrative.

4.4.2 Results

Tables 4.4-2 to 4.4-5 show the regression results of Model 4.1 for the four fiscal narrative indicators, respectively. We run five separate specifications in which we add one of the

Table 4.4-2: Effect of articles implying fiscal expansion causes inflation on $E_{it}\pi_{t+12}$, baseline

	(1)	(2)	(3)	(4)	(5)
expansion_inflation$_{t-1}$	0.203***	0.107***	0.111***	0.105***	0.105***
	(0.006)	(0.005)	(0.005)	(0.005)	(0.006)
π_{t-1}		0.205***	0.203***	0.212***	0.212***
		(0.009)	(0.009)	(0.010)	(0.010)
IP_{t-1}			-3.640***	-3.652***	-3.652***
			(0.284)	(0.285)	(0.285)
ECB_rate_{t-1}				-0.078***	-0.077***
				(0.023)	(0.024)
total_fiscal$_{t-1}$					-0.000
					(0.000)
Household FE	yes	yes	yes	yes	yes
Observations	159783	159783	159783	159783	159783
R^2	0.020	0.032	0.033	0.033	0.033

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Winsorized survey data at 2% and 98%.

Table 4.4-3: Effect of articles implying unsustainable government debt on $E_{it}\pi_{t+12}$, baseline

	(1)	(2)	(3)	(4)	(5)
debt_unsustainable$_{t-1}$	0.013***	0.022***	0.024***	0.032***	0.009**
	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
π_{t-1}		0.273***	0.274***	0.283***	0.261***
		(0.009)	(0.009)	(0.010)	(0.010)
IP_{t-1}			-3.119***	-3.262***	-3.253***
			(0.283)	(0.284)	(0.284)
ECB_rate_{t-1}				-0.150***	-0.205***
				(0.023)	(0.023)
total_fiscal$_{t-1}$					0.004***
					(0.000)
Household FE	yes	yes	yes	yes	yes
Observations	159783	159783	159783	159783	159783
R^2	0.000	0.028	0.029	0.030	0.031

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Winsorized survey data at 2% and 98%.

Table 4.4-4: Effect of articles implying the government should spend more on $E_{it}\pi_{t+12}$, baseline

	(1)	(2)	(3)	(4)	(5)
spending$_{t-1}$	0.056***	0.030***	0.029***	0.035***	0.034***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
π_{t-1}		0.240***	0.241***	0.247***	0.247***
		(0.010)	(0.010)	(0.010)	(0.010)
IP_{t-1}			-2.321***	-2.311***	-2.317***
			(0.280)	(0.280)	(0.280)
ECB_rate_{t-1}				-0.194***	-0.194***
				(0.022)	(0.023)
total_fiscal$_{t-1}$					0.000
					(0.000)
Household FE	yes	yes	yes	yes	yes
Observations	159783	159783	159783	159783	159783
R^2	0.012	0.031	0.031	0.033	0.033

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Winsorized survey data at 2% and 98%.

Table 4.4-5: Effect of articles implying the government should save more on $E_{it}\pi_{t+12}$, baseline

	(1)	(2)	(3)	(4)	(5)
saving$_{t-1}$	0.019***	0.014***	0.014***	0.019***	0.013***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
π_{t-1}		0.269***	0.269***	0.279***	0.269***
		(0.009)	(0.009)	(0.010)	(0.010)
IP_{t-1}			-2.925***	-3.019***	-3.100***
			(0.282)	(0.282)	(0.283)
ECB_rate_{t-1}				-0.173***	-0.194***
				(0.023)	(0.023)
total_fiscal$_{t-1}$					0.002***
					(0.000)
Household FE	yes	yes	yes	yes	yes
Observations	159783	159783	159783	159783	159783
R^2	0.002	0.029	0.029	0.031	0.031

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Winsorized survey data at 2% and 98%.

four macro controls one at a time, plus specification (5) which includes all controls as the baseline. Clustered standard errors at the household level are reported in parentheses.

For all four narrative indicators, we find statistically significant effects on 12-months-ahead household inflation expectations. As reported in Table 4.4-2, an additional article linking fiscal expansion to inflation increases household inflation expectations by 0.105 pp. The effects of the macro variables are similarly significant, and their directions are in line with economic theory and results in the empirical literature. The effect of articles deeming the government debt level to be unsustainable (see Table 4.4-3) is positive and significant at the 5% level. An additional article on unsustainable debt increases inflation expectations by 0.009 pp. Media discussions on whether the government should in- or decrease its spending also have a significantly positive effect on household expectations, although smaller in size compared to the effects of the fiscal expansion leads to inflation narrative. Tables 4.4-5 and 4.4-4 report an increase in inflation expectations by 0.013 pp for the saving and 0.034 pp for the spending narratives, respectively.

The baseline results are obtained using the winsorized survey data, but we also estimate Model 4.1 on the non-adjusted data. Results reported in Appendix 4.7.4 show that the estimates are not particularly sensitive to this. For further robustness, we repeat the baseline estimation without weighting the newspaper data by readership in the aggregation of the indices and report the results in Appendix 4.7.5. Only the effect of the unsustainable debt narrative then becomes insignificant and zero.

4.4.3 Discussion

The relationship between household inflation expectations and the fiscal expansion leads to inflation narrative is in line with what one would expect. The more articles report on fiscal expansion leading to inflation, the more households will adopt this narrative for themselves and therefore expect higher inflation. Given that the point estimate with only household fixed effects falls substantially when adding controls (from 0.203 to 0.105, Table 4.4-2), this suggests that much of the effect is channeled through general economic sentiment and other macro-shocks. Once controlling for the standard drivers of expectations, the effect appears to be robust but economically modest. The standard deviation of the fiscal expansion⇒inflation indicator over the CES sample length is 3.31 articles. This means that a standard deviation uptick in articles following this narrative translates into an increase in inflation expectations by around 0.35 pp. As is usual in micro-level FE regressions, the overall R^2 is low—household expectations are driven by many idiosyncratic factors. Yet the narrative effect remains highly statistically significant once we condition on standard macro drivers.

The positive coefficient of the unsustainable government debt narrative indicator in column (5) of Table 4.4-3 is also in line with what Bassetto and Miller (2025) predict. An increase in coverage deeming public debt to be unsustainably high leads to more households acquiring this information, which they conclude can result in sudden inflation scares. This effect is, though statistically significant, in absolute terms much smaller than for the fiscal expansion⇒inflation narrative. The standard deviation in the survey period is 3.03, such that an increase of articles of this size translates into a rise in inflation expectations by 0.03 pp. The savings narrative is closest to the unsustainable debt narrative and its coefficient, being 0.013, is just slightly bigger.

4.5 The effect of fiscal narratives on the transmission of government spending shocks

Motivated by the result that fiscal narratives can increase household inflation expectations, we next analyze their effect on the transmission of government spending shocks to real GDP and inflation. We do so by following a two-step estimation approach. In the first step, we estimate a structural government spending shock, and in the second step, we employ the resulting shock series in local projections. We model the non-linearity in the transmission of government spending shocks by interacting the shocks with the narrative indicators, one at a time.

4.5.1 Estimation strategy

We estimate the government spending shock series in a structural VAR (SVAR) model employing a recursive identification strategy with short-run restrictions. In the selection of variables and their ordering, we follow Auerbach and Gorodnichenko (2012) by placing a measure of expected government spending first, to control for fiscal foresight ($FG_{t|t-1}$), followed by government spending (G_t), tax revenues (τ_t), and GDP (Y_t). $FG_{t|t-1}$ is ordered first because a contemporaneous shock in G_t cannot affect a forecast for period t made in $t-1$ (see Auerbach and Gorodnichenko (2012) and Leduc et al. (2007) for further discussion). The ordering of the rest of the variables follows a standard specification first introduced by Blanchard and Perotti (2002). Using the Akaike Information Criterion, we determine the optimal number of lags to be 8.

In the data selection for the vector of variables $X_t = [FG_{t|t-1}, ln(G_t), ln(\tau_t), ln(Y_t)]'$ in the SVAR, we follow Amendola et al. (2020). Government spending is the sum of the seasonally adjusted series of real general government final consumption expenditure and real gross fixed capital formation for Germany, while tax revenues are provided by the German Federal Statistical Office as total revenues across the federal and local levels. As a measure of fiscal foresight, we use the OECD government spending growth forecasts as proposed by Ilori et al. (2022). The OECD forecasts are produced twice a year and are published in the publicly available *OECD Economic Outlook*. We use the government spending forecast data provided by Ilori et al. (2022), which ends in 2019Q4, and extend it forward by interpolating the semi-annual data to quarterly frequency by applying mid-point averaging. All series except the government spending forecasts enter the model in logs, and the frequency is quarterly.

In the second step, we employ the government spending shock estimated in the SVAR in a local projections model as in Jordà (2005). To model the potentially non-linear effect of fiscal narratives on the transmission of government spending shocks, we introduce an interaction term ($shock_t \cdot narrative_t$):

$$y_{t+h} = \alpha_h + \beta_h shock_t + \gamma_h(shock_t \cdot narrative_t) + \sum_{l=1}^{n} \delta_{h,l} W_{t-l} + u_{t+h} \tag{4.2}$$

The interaction term is in the spirit of Ascari and Haber (2022), who interact a monetary policy shock variable with its absolute values to model the non-linear effect of the shock size on its transmission. Analogously to Ascari and Haber (2022), the coefficient β_h captures the baseline effect of the government spending shock ($shock_t$), irrespective of the prevalent fiscal narrative, while γ_h measures how this effect changes in the event of a unit change in the narrative indicator ($narrative_t$). As we standardize the narrative indicators before interacting them with the shock series, this unit change corresponds to one standard deviation. If γ_h has the opposite sign of β_h, the fiscal narrative dampens the effect of government spending on the variable of interest y_t, and if both have the same sign, the effect will be amplified in the respective direction.

We estimate model 4.2 for horizons $h = 1, ..., 12$, which corresponds to three years due to the quarterly data frequency. The control variables W_t comprise up to $n = 4$ lags (motivated by the quarterly frequency of the data) of real GDP, government spending and tax revenues (transformed to logs and multiplied by 100), government spending forecasts, the ECB's shadow rate constructed by Wu and Xia (2020), the GDP deflator, as well as the shock series itself. This choice follows Montiel Olea et al. (2025), who suggest including lags of the impulse and control variables as well as other variables that are strong predictors of both in the set of controls, to counteract potential dynamic misspecifications as well as imperfections in the shock measure. The reported error bands are computed with Newey-West standard errors to control for autocorrelation.

4.5.2 Results

4.5.2.1 The effect of individual narrative indicators

Figures 4.5-5 and 4.5-6 present the results of the local projections described above, measuring the effect of a government spending shock on GDP and inflation, respectively. The blue responses are the linear effects of the fiscal shock, while the orange ones depict the interaction term between the shock and one of the four narrative indicators. Each narrative acts differently on the transmission of a fiscal shock.

The narrative that the government is supposed to decrease its spending (the saving narrative) has a dampening effect on impact and towards the end of the projection horizon on the transmission of a fiscal shock to GDP. Similarly, it has such a dampening effect on inflation in the first four quarters after impact, as the linear and the interaction term coefficients have opposite signs. The opposite is the case for the narrative that the government should increase its spending, for which both have the same sign, and hence this narrative amplifies the transmission of a fiscal shock to both GDP and inflation.

The effect of the unsustainable debt narrative on the fiscal shock transmission to GDP seems to be negligible, as the interaction term coefficient rarely differs from zero. In those

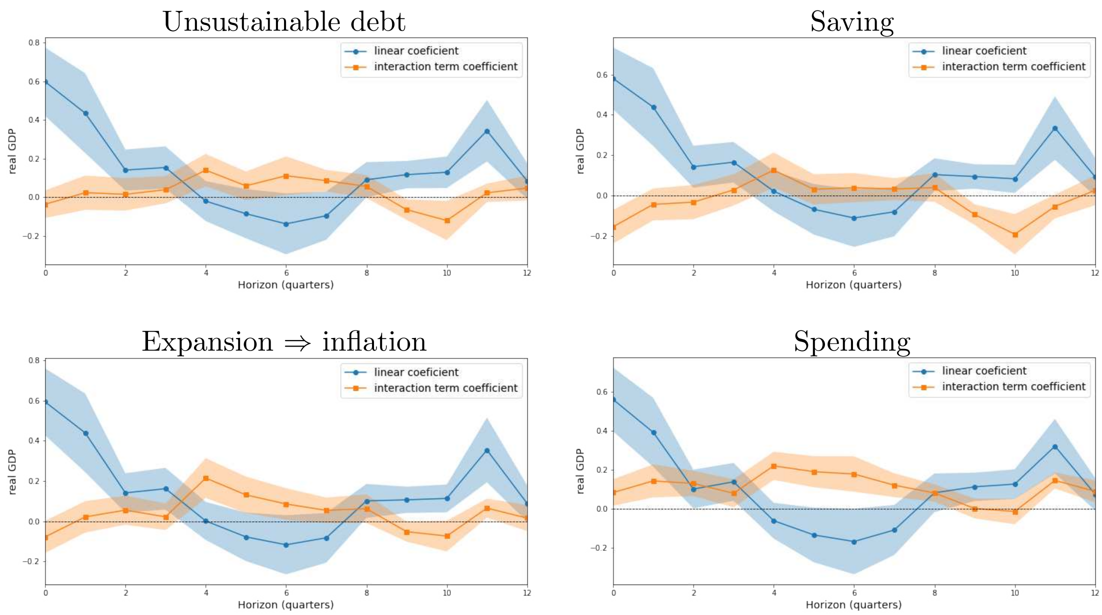

Figure 4.5-5: Impulse responses real GDP

Notes: This figure shows the impulse responses of real GDP after and government spending shock (blue) and its interaction (orange) with one of the four narrative indicators. Responses are estimated across a 12 quarters projection horizon and are depicted with 68% confidence bands computed with Newey-West standard errors.

cases where it does, it has the same sign as the linear coefficient, which means that if there is a non-linear effect of this narrative, then it seems to be amplifying the effect of a fiscal shock. This effect is even more pronounced for inflation. The narrative that fiscal expansion leads to inflation has very limited effects on the transmission of a fiscal shock, as its interaction term coefficient rarely significantly differs from zero.

4.5.2.2 The effect of a composite narrative indicator

Based on the results of the effect of the individual narrative indicators on the transmission of government spending shocks, we create a composite index that signals if government spending shocks will be amplified by the fiscal narratives news outlets convey. To this end, we create the sum of the individual narrative indicators with their signs depending on whether the narrative is likely to amplify or dampen the effect of a government spending shock. The spending narrative has a clearly amplifying effect both on GDP and prices, while the saving narrative has the opposite effect. Hence, we multiply the latter with -1. The unsustainable debt narrative has a significantly amplifying effect at least for inflation, while for the fiscal expansion leads to inflation narrative, the case is not as clear. Nonetheless, we weight both of them with a positive sign. As a robustness check, we also report a composite indicator consisting only of the saving and spending narrative as well as their effects on GDP and inflation in Appendix 4.7.6. These results look very similar.

We present the resulting index in Figure 4.5-7. It is evident that the index had very low values in the aftermath of the GFC and at the beginning of the European sovereign debt crisis in 2010. Then, since 2012, the index has constantly increased with a local peak during COVID-19 and heightened volatility after Russia invaded Ukraine, which sparked

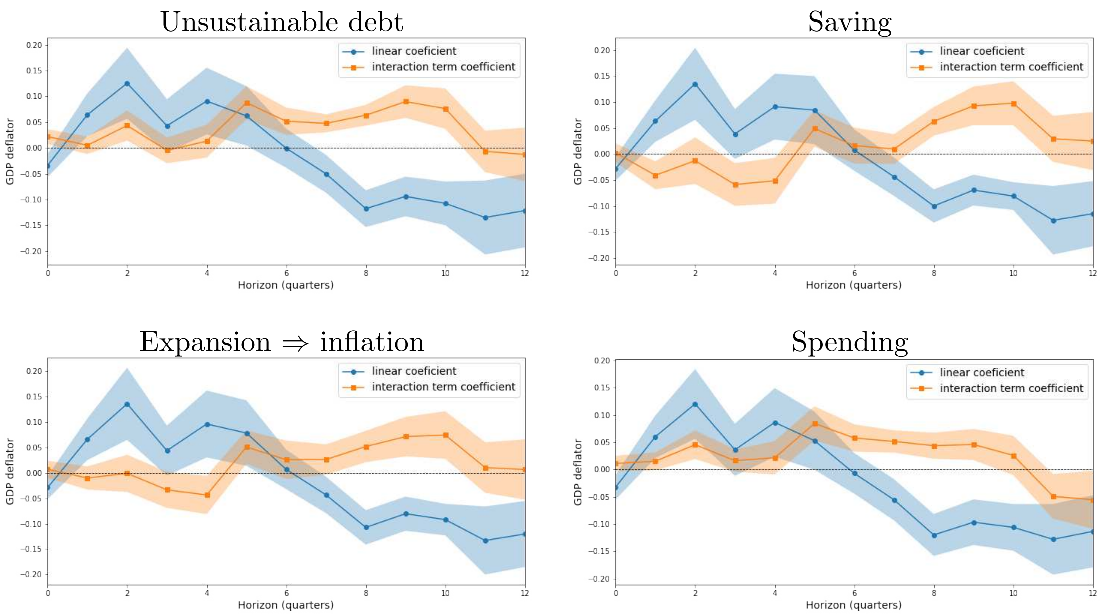

Figure 4.5-6: Impulse responses GDP deflator

Notes: This figure shows the impulse responses of the GDP deflator after and government spending shock (blue) and its interaction (orange) with one of the four narrative indicators. Responses are estimated across a 12 quarters projection horizon and are depicted with 68% confidence bands computed with Newey-West standard errors.

a large debate on increasing military spending in Germany.

Figure 4.5-8 shows that the resulting composite indicator indeed amplifies the effect of a government spending shock on both GDP and inflation. This means that if the index takes higher values, then the effects of government spending will be larger. Hence, in such periods, a government spending shock is likely to be more inflationary and at the same time more stimulating for output.

4.6 Conclusion

In this paper, we analyze a large set of articles from German newspapers using AI, and evaluate the fiscal narratives conveyed in these articles. Our results show that these fiscal

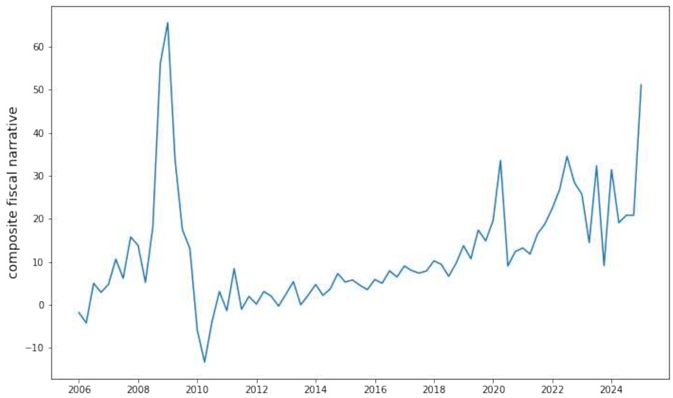

Figure 4.5-7: Composite fiscal narrative index

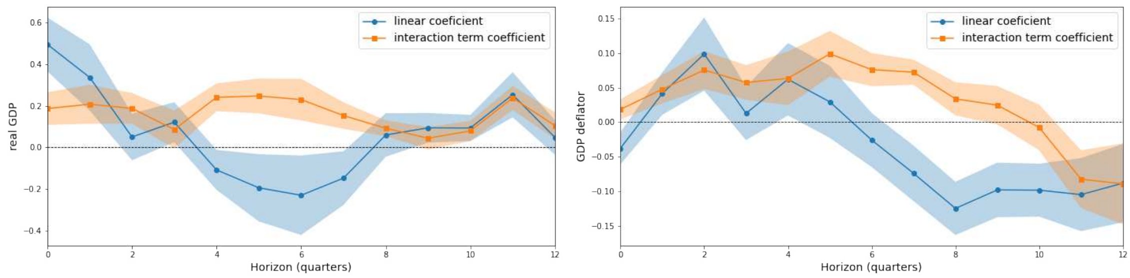

Figure 4.5-8: Impulse responses interaction with composite index
Notes: This figure presents the impulse responses of a government spending shock (blue) on GDP (left panel) and the GDP deflator (right panel) and its interaction with the composite fiscal narrative indicator (orange) together with 68% confidence intervals.

narratives are not merely descriptive features of public discourse, but instead play an active role in influencing household inflation expectations. In particular, we find that articles linking fiscal expansion to inflation have the strongest effect, raising household inflation expectations by 0.105 percentage points. Narratives emphasizing unsustainable debt, government spending, or fiscal restraint also have statistically significant, though smaller, impacts. These findings highlight the importance of the information environment and the expectations channel in the transmission of fiscal policy to the real economy.

Going further, we demonstrate that the prevalence of specific fiscal narratives in the media can either amplify or dampen the macroeconomic impact of government spending shocks. Both the spending and unsustainable debt narratives are found to amplify the effects of fiscal shocks on GDP and inflation, while narratives advocating saving or directly linking fiscal expansion to inflation tend to moderate these effects. By combining the individual indicators into a composite fiscal narrative index, we construct a real-time measure that can signal periods when fiscal shocks are likely to have stronger effects on output and prices.

Overall, our findings underline the critical role of public narratives in shaping economic expectations and policy effectiveness. The composite index we develop offers policymakers a new tool for monitoring the fiscal information climate and anticipating the potential impact of fiscal interventions. More broadly, our approach demonstrates how textual analysis and AI can be harnessed to better understand the evolving interplay between public discourse, expectations, and the macroeconomy.

4.7 Appendix

4.7.1 Fiscal dictionaries

General fiscal terms:

Fiskal*, Haushalt*, Finanz*, Finanzminist*, Staatsfinanzie-rung, Staatsbudget, Haushaltsplan

Fiscal expansion:

Staatsausgabe, Staatsverschuldung, Kreditaufnahme, Schuldenpolitik, Staatsdefizit, Ausgabenerhöhung, Investition, Subvention, Konjunkturprogramm, Konjunkturpolitik, Wachstumsimpuls, Konjunkturförderung, Finanzspritze, Sondervermögen, Steuersenkung, Prämie

Fiscal consolidation:

Haushaltskonsolidierung, Konsolidier*, Schuldenabbau, Sparpolitik, Sparmaßnahme, Defizitreduzierung, Haushaltsdisziplin, Budgetkürzung, Haushaltsüberschuss, Steuererhöhung, Ausgabenkürzung, Ausgabensenkung, Austerität*, Schuldenbremse, Budgetkonsolidierung, Kostensenkung, Sparhaushalt, Rentenkürzung

4.7.2 ChatGPT prompts

```
You will be provided with a news article. Your task is to determine whether the
article establishes a narrative that fiscal expansion (e.g., government spending
increases, tax cuts) will lead to inflation. Classify an article as Yes if it
clearly presents a narrative—whether stated as a fact or as a speculation—that
fiscal expansion will cause or eventually result in inflation. Do not flag
articles where fiscal expansion and inflation are mentioned independently
without linking them in a cause-and-effect manner. Do not hypothesize a
relationship based on economic theory or indirect reasoning.

Steps:
1. Read the article carefully.
2. Look for any statement or narrative—explicit, combined, or speculative—where
fiscal expansion is connected to inflation as a cause or expected outcome.
3. If the article only discusses these topics in isolation or does not establish
a cause-and-effect link, classify it as No.
4. Respond with: Response: Yes or No, Explanation (brief, 1-2 sentences): Justify
your answer by quoting or summarizing the relevant text that establishes (or
does not establish) the connection.

Examples: 1. Article Excerpt: The government's recent increase in spending is
expected to drive up prices; experts warn that this fiscal expansion will likely
result in higher inflation rates.
Response: Yes,
Explanation: The article clearly links increased government spending to higher
inflation through an explicit narrative.
2. Article Excerpt: While government spending has increased, some economists
argue that other factors, like global supply chain issues, are keeping inflation
in check.
Response: No,
Explanation: The article mentions fiscal expansion and inflation separately but
does not establish a cause-and-effect link.
3. Article Excerpt: If the government continues to boost spending, it might
eventually set off a chain reaction that results in inflation.
Response: Yes,
Explanation: The article speculates that continued fiscal expansion could
eventually lead to inflation, thus establishing the desired narrative.
Article:
```

Figure 4.7-1: Prompt to assess if an article establishes a causal link between fiscal expansion and inflation

```
Analyze the following article to determine whether it suggests that government
debt levels in Germany are unsustainable. Consider both explicit statements and
implicit impressions, while also accounting for speculative or contradictory
arguments. Classify the article into one of the following three categories:

'Yes': If the article states or strongly implies that German government debt
levels are unsustainable now or in the near future.
'No': If the article mentions government debt but does not suggest that it is
unsustainable, OR if it only raises speculative concerns without claiming current
unsustainability.
'NA': If the article does not mention government debt at all, OR if it is not
directly mentioning government debt in Germany. If the article discusses economic
risks (e.g., inflation, recession) without directly linking them to government
debt, classify as 'NA' unless the connection is explicitly made.

Respond with only one of the three categories: 'Yes', 'No', or 'NA'. Provide no
additional text or explanation.
Article:
```

Figure 4.7-2: Prompt to assess if an article conveys that the debt level is unsustainable

```
Analyze the following article to determine whether it explicitly conveys an
opinion on whether the German government should save more or increase spending.
Classify the article into one of the following three categories:

‘Save’: If the article explicitly states or strongly argues that the German
government should reduce spending, save more, or focus on fiscal restraint.
‘Spend’: If the article explicitly states or strongly argues that the German
government should spend more, invest in public programs, or prioritize economic
stimulus.
‘NA’: If the article does not express an explicit opinion on German government
saving or spending. The opinion must be explicitly stated - general discussions
or reporting on government spending without a clear stance should be classified
as ‘NA’. If the article presents both perspectives, classify based on the
dominant opinion expressed.

Respond with only one of the three categories: ‘Save’, ‘Spend’, or ‘NA’. Provide
no additional text or explanation.
Article:
```

Figure 4.7-3: Prompt to analyze whether an article expresses the opinion that the government should increase its savings or spending

4.7.3 Effects of individual newspapers

Table 4.7-1: Effect of articles implying fiscal expansion causes inflation on $E_{it}\pi_{t+12}$, individual newspapers

	SZ	FAZ	TAZ	Bild
expansion_inflation$_{t-1}$	0.023***	-0.005	0.099***	0.085***
	(0.002)	(0.004)	(0.005)	(0.006)
π_{t-1}	0.223***	0.260***	0.206***	0.239***
	(0.010)	(0.010)	(0.009)	(0.010)
IP_{t-1}	-3.494***	-3.152***	-3.522***	-3.361***
	(0.284)	(0.286)	(0.284)	(0.283)
ECB_rate_{t-1}	-0.151***	-0.214***	-0.158***	-0.119***
	(0.023)	(0.023)	(0.023)	(0.023)
total_fiscal$_{t-1}$	0.002***	0.005***	0.004***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Household FE	yes	yes	yes	yes
Observations	159783	159783	159783	159783
R^2	0.032	0.031	0.035	0.032

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Winsorized survey data at 2% and 98%.

Table 4.7-2: Effect of articles implying unsustainable government debt on $E_{it}\pi_{t+12}$, individual newspapers

	SZ	FAZ	TAZ	Bild
debt_unsustainable$_{t-1}$	-0.019***	0.020***	0.051***	0.023***
	(0.002)	(0.004)	(0.004)	(0.003)
π_{t-1}	0.252***	0.263***	0.269***	0.266***
	(0.010)	(0.010)	(0.010)	(0.010)
IP_{t-1}	-3.698***	-3.454***	-3.028***	-3.422***
	(0.287)	(0.289)	(0.283)	(0.286)
ECB_rate_{t-1}	-0.201***	-0.206***	-0.183***	-0.197***
	(0.023)	(0.023)	(0.023)	(0.023)
total_fiscal$_{t-1}$	0.006***	0.004***	0.003***	0.003***
	(0.000)	(0.000)	(0.000)	(0.000)
Household FE	yes	yes	yes	yes
Observations	159783	159783	159783	159783
R^2	0.032	0.031	0.032	0.031

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Winsorized survey data at 2% and 98%.

Table 4.7-3: Effect of articles implying the government should save more on $E_{it}\pi_{t+12}$, individual newspapers

	SZ	FAZ	TAZ	Bild
saving_{t-1}	0.005***	0.009***	0.049***	0.010***
	(0.001)	(0.003)	(0.003)	(0.002)
π_{t-1}	0.262***	0.261***	0.275***	0.268***
	(0.010)	(0.010)	(0.010)	(0.010)
IP_{t-1}	-3.084***	-3.196***	-2.765***	-3.147***
	(0.286)	(0.282)	(0.283)	(0.282)
ECB_rate_{t-1}	-0.206***	-0.212***	-0.171***	-0.194***
	(0.023)	(0.023)	(0.023)	(0.023)
$\text{total\_fiscal}_{t-1}$	0.003***	0.004***	0.000	0.002***
	(0.000)	(0.000)	(0.000)	(0.000)
Household FE	yes	yes	yes	yes
Observations	159783	159783	159783	159783
R^2	0.031	0.031	0.032	0.031

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Winsorized survey data at 2% and 98%.

Table 4.7-4: Effect of articles implying the government should spend more on $E_{it}\pi_{t+12}$, individual newspapers

	SZ	FAZ	TAZ	Bild
spending_{t-1}	0.010***	0.009***	0.018***	0.028***
	(0.001)	(0.002)	(0.001)	(0.002)
π_{t-1}	0.257***	0.263***	0.244***	0.245***
	(0.010)	(0.010)	(0.010)	(0.010)
IP_{t-1}	-2.606***	-3.102***	-2.937***	-2.597***
	(0.284)	(0.280)	(0.282)	(0.279)
ECB_rate_{t-1}	-0.207***	-0.202***	-0.220***	-0.195***
	(0.023)	(0.023)	(0.023)	(0.023)
$\text{total\_fiscal}_{t-1}$	0.001***	0.004***	0.002***	0.002***
	(0.000)	(0.000)	(0.000)	(0.000)
Household FE	yes	yes	yes	yes
Observations	159783	159783	159783	159783
R^2	0.032	0.031	0.034	0.032

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Winsorized survey data at 2% and 98%.

4.7.4 Non-winsorized survey data

Table 4.7-5: Effect of articles implying fiscal expansion causes inflation on $E_{it}\pi_{t+12}$, non-winsorized data

	(1)	(2)	(3)	(4)	(5)
expansion_inflation$_{t-1}$	0.205***	0.107***	0.112***	0.104***	0.106***
	(0.007)	(0.006)	(0.006)	(0.006)	(0.008)
π_{t-1}		0.209***	0.206***	0.217***	0.217***
		(0.011)	(0.011)	(0.011)	(0.011)
IP_{t-1}			-3.848***	-3.861***	-3.864***
			(0.350)	(0.350)	(0.350)
ECB_rate_{t-1}				-0.094***	-0.089***
				(0.027)	(0.028)
total_fiscal$_{t-1}$					-0.000
					(0.000)
Household FE	yes	yes	yes	yes	yes
Observations	159783	159783	159783	159783	159783
R^2	0.014	0.022	0.023	0.023	0.023

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4.7-6: Effect of articles implying unsustainable government debt on $E_{it}\pi_{t+12}$, non-winsorized data

	(1)	(2)	(3)	(4)	(5)
debt_unsustainable$_{t-1}$	0.012***	0.021***	0.023***	0.032***	0.010**
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
π_{t-1}		0.277***	0.277***	0.288***	0.267***
		(0.011)	(0.011)	(0.011)	(0.011)
IP_{t-1}			-3.318***	-3.476***	-3.467***
			(0.350)	(0.349)	(0.349)
ECB_rate_{t-1}				-0.165***	-0.217***
				(0.027)	(0.027)
total_fiscal$_{t-1}$					0.004***
					(0.000)
Household FE	yes	yes	yes	yes	yes
Observations	159783	159783	159783	159783	159783
R^2	0.000	0.019	0.020	0.021	0.021

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4.7-7: Effect of articles implying the government should save more on $E_{it}\pi_{t+12}$, non-winsorized data

	(1)	(2)	(3)	(4)	(5)
saving_{t-1}	0.019***	0.013***	0.013***	0.018***	0.013***
	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)
π_{t-1}		0.273***	0.273***	0.284***	0.275***
		(0.011)	(0.011)	(0.011)	(0.011)
IP_{t-1}			-3.132***	-3.234***	-3.308***
			(0.348)	(0.348)	(0.348)
ECB_rate_{t-1}				-0.187***	-0.206***
				(0.027)	(0.028)
$\text{total\_fiscal}_{t-1}$					0.002***
					(0.001)
Household FE	yes	yes	yes	yes	yes
Observations	159783	159783	159783	159783	159783
R^2	0.001	0.020	0.020	0.022	0.022

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4.7-8: Effect of articles implying the government should spend more on $E_{it}\pi_{t+12}$, non-winsorized data

	(1)	(2)	(3)	(4)	(5)
spending_{t-1}	0.056***	0.029***	0.028***	0.034***	0.035***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
π_{t-1}		0.244***	0.246***	0.252***	0.252***
		(0.011)	(0.011)	(0.011)	(0.011)
IP_{t-1}			-2.543***	-2.533***	-2.523***
			(0.346)	(0.346)	(0.345)
ECB_rate_{t-1}				-0.208***	-0.207***
				(0.027)	(0.028)
$\text{total\_fiscal}_{t-1}$					-0.000
					(0.000)
Household FE	yes	yes	yes	yes	yes
Observations	159783	159783	159783	159783	159783
R^2	0.008	0.021	0.022	0.023	0.023

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

4.7.5 Results non-weighted fiscal narrative indicators

Table 4.7-9: Effect of articles implying fiscal expansion causes inflation on $E_{it}\pi_{t+12}$, non-weighted data

	(1)	(2)	(3)	(4)	(5)
expansion_inflation$_{t-1}$	0.046***	0.024***	0.026***	0.024***	0.027***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
π_{t-1}		0.187***	0.182***	0.191***	0.191***
		(0.010)	(0.010)	(0.010)	(0.010)
IP_{t-1}			-3.932***	-3.924***	-4.020***
			(0.287)	(0.287)	(0.289)
ECB_rate_{t-1}				-0.073***	-0.046**
				(0.023)	(0.023)
total_fiscal$_{t-1}$					-0.000***
					(0.000)
Household FE	yes	yes	yes	yes	yes
Observations	159783	159783	159783	159783	159783
R^2	0.024	0.032	0.033	0.033	0.033

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Winsorized survey data at 2% and 98%. Non-weighted newspaper data.

Table 4.7-10: Effect of articles implying unsustainable government debt on $E_{it}\pi_{t+12}$, non-weighted data

	(1)	(2)	(3)	(4)	(5)
debt_unsustainable$_{t-1}$	0.001	0.003***	0.003***	0.005***	0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
π_{t-1}		0.272***	0.273***	0.282***	0.261***
		(0.009)	(0.009)	(0.010)	(0.010)
IP_{t-1}			-2.968***	-3.053***	-3.059***
			(0.282)	(0.282)	(0.282)
ECB_rate_{t-1}				-0.144***	-0.173***
				(0.023)	(0.023)
total_fiscal$_{t-1}$					0.001***
					(0.000)
Household FE	yes	yes	yes	yes	yes
Observations	159783	159783	159783	159783	159783
R^2	0.000	0.028	0.028	0.030	0.030

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Winsorized survey data at 2% and 98%. Non-weighted newspaper data.

Table 4.7-11: Effect of articles implying the government should save more on $E_{it}\pi_{t+12}$, non-weighted data

	(1)	(2)	(3)	(4)	(5)
saving$_{t-1}$	0.007***	0.005***	0.005***	0.007***	0.008***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)
π_{t-1}		0.267***	0.267***	0.276***	0.287***
		(0.009)	(0.009)	(0.010)	(0.010)
IP_{t-1}			-2.843***	-2.906***	-2.867***
			(0.283)	(0.283)	(0.283)
ECB_rate_{t-1}				-0.190***	-0.181***
				(0.023)	(0.023)
total_fiscal$_{t-1}$					-0.000***
					(0.000)
Household FE	yes	yes	yes	yes	yes
Observations	159783	159783	159783	159783	159783
R^2	0.002	0.029	0.029	0.031	0.031

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Winsorized survey data at 2% and 98%. Non-weighted newspaper data.

Table 4.7-12: Effect of articles implying the government should spend more on $E_{it}\pi_{t+12}$, non-weighted data

	(1)	(2)	(3)	(4)	(5)
spending$_{t-1}$	0.010***	0.006***	0.006***	0.007***	0.009***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
π_{t-1}		0.244***	0.245***	0.252***	0.262***
		(0.010)	(0.010)	(0.010)	(0.010)
IP_{t-1}			-2.406***	-2.411***	-2.303***
			(0.281)	(0.281)	(0.281)
ECB_rate_{t-1}				-0.207***	-0.188***
				(0.022)	(0.022)
total_fiscal$_{t-1}$					-0.001***
					(0.000)
Household FE	yes	yes	yes	yes	yes
Observations	159783	159783	159783	159783	159783
R^2	0.011	0.031	0.031	0.034	0.034

Notes: Standard errors clustered at the household level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Winsorized survey data at 2% and 98%. Non-weighted newspaper data.

4.7.6 Composite indicator consisting of saving and spending narratives

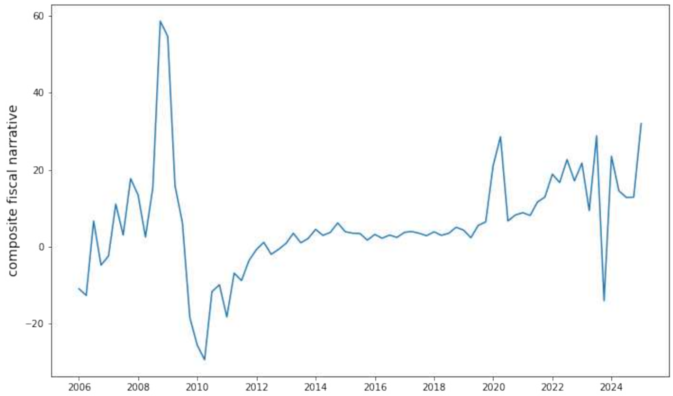

Figure 4.7-4: Standardized composite fiscal narrative index
Notes: This index consists of the spending narrative (weighted positively) and the saving narrative (weighted negatively) only.

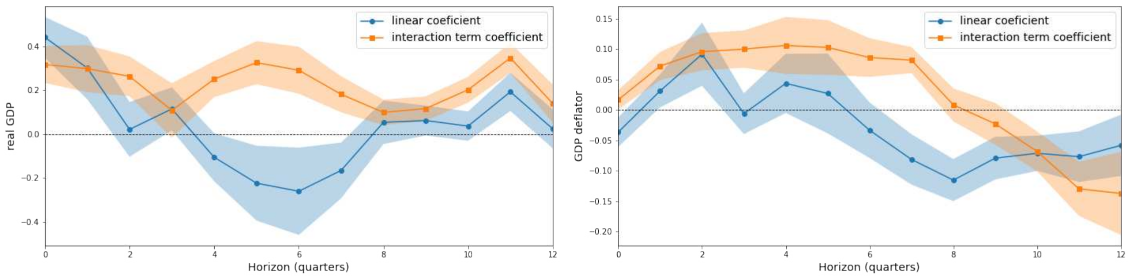

Figure 4.7-5: IRFs robustness composite index
Notes: This figure presents the impulse responses of a government spending shock (blue) on GDP (left panel) and the GDP deflator (right panel) and its interaction with the composite fiscal narrative indicator (orange) together with 68% confidence intervals.

References

agma, A. M.-A. e. (2022). *Media-Analyse Tageszeitungen 2022*. Arbeitsgemeinschaft Media-Analyse e.V.

Altavilla, C., Brugnolini, L., Gürkaynak, R. S., Motto, R., and Ragusa, G. (2019). Measuring euro area monetary policy. *Journal of Monetary Economics*, 108:162–179.

Álvarez, F., Beraja, M., Gonzalez-Rozada, M., and Neumeyer, P. A. (2019). From hyperinflation to stable prices: Argentina's evidence on menu cost models. *The Quarterly Journal of Economics*, 134(1):451–505.

Álvarez, F., Lippi, F., and Paciello, L. (2018). Monetary shocks in models with observation and menu costs. *Journal of the European Economic Association*, 16(2):353–382.

Álvarez, F., Lippi, F., and Passadore, J. (2017). Are state and time dependent models really different? In Eichenbaum, M., Hurst, E., and Parker, J., editors, *NBER Macroeconomics Annual*, pages 379–457. University of Chicago Press.

Amendola, A., Di Serio, M., Fragetta, M., and Melina, G. (2020). The euro-area government spending multiplier at the effective lower bound. *European Economic Review*, 127:103480.

Amiti, M., Itskhoki, O., and Konings, J. (2019). International shocks, variable markups, and domestic prices. *The Review of Economic Studies*, 86(6):2356–2402.

Andrade, P., Gautier, E., Mengus, E., Moench, E., and Schmidt, T. (2025). Household beliefs about fiscal dominance. *HEC Paris Research Paper forthcoming*.

Andre, P., Haaland, I., Roth, C., Wiederholt, M., and Wohlfart, J. (2024). Narratives about the macroeconomy. SAFE Working Paper.

Andre, P., Haaland, I., Roth, C., and Wohlfart, J. (2023). Narratives about the macroeconomy. *CESifo Working Paper*.

Andre, P., Roth, C., Haaland, I., and Wohlfart, J. (2021). Inflation narratives. *CEPR Discussion Paper No. DP16758*.

Angelico, C., Marcucci, J., Miccoli, M., and Quarta, F. (2022). Can we measure inflation expectations using Twitter? *Journal of Econometrics*, 228(2):259–277.

Aruoba, S. B. and Drechsel, T. (2024). Identifying monetary policy shocks: A natural language approach. *NBER Working Paper*.

Ascari, G. and Haber, T. (2022). Non-linearities, state-dependent prices and the transmission mechanism of monetary policy. *The Economic Journal*, 132:37–57.

Auer, R. A. and Schoenle, R. S. (2016). Market structure and exchange rate pass-through. *Journal of International Economics*, 98:60–77.

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

Bank for International Settlements (2022). Annual economic report. Chapter II.

Bańkowska, K., Borlescu, A., Charalambakis, E., Dias da Silva, A., Di Laurea, D., Dossche, M., Georgarakos, D., Honkkila, J., Kennedy, N., Kenny, G., et al. (2021). ECB consumer expectations survey: An overview and first evaluation. *ECB Occasional Paper*, (2021/287).

Barro, R. J. and Bianchi, F. (2024). Fiscal influences on inflation in OECD countries, 2020-2023. *NBER Working Paper*.

Bassetto, M. and Miller, D. S. (2025). A monetary-fiscal theory of sudden inflations. *The Quarterly Journal of Economics*, page 1959–2000.

Baumeister, C. and Hamilton, J. D. (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review*, 109:1873–1910.

Belal, M., She, J., and Wong, S. (2023). Leveraging ChatGPT as text annotation tool for sentiment analysis. *arXiv e-prints*, pages arXiv–2306.

Berger, D. and Vavra, J. (2019). Shocks versus responsiveness: What drives time-varying dispersion? *Journal of Political Economy*, 127(5):2104–2142.

Blanchard, O. and Perotti, R. (2002). An empirical characterization of the dynamic effects of changes in government spending and taxes on output. *The Quarterly Journal of Economics*, 117(4):1329–1368.

Blanco, A., Ottonello, P., and Ranošová, T. (2025). The dynamics of large inflation surges. *The Review of Economics and Statistics*, pages 1–31.

Blinder, A. S., Canetti, E. R. D., Lebow, D. E., and Rudd, J. B. (1998). *Asking About Prices: A New Approach to Understanding Price Stickiness*. Russell Sage Foundation.

Bobeica, E., Ciccarelli, M., and Vansteenkiste, I. (2020). The link between labor cost inflation and price inflation in the euro area. *Serie Banca Central, análisis y políticas económicas, No. 27.*

Bobeica, E., Ciccarelli, M., and Vansteenkiste, I. (2021). The changing link between labor cost and price inflation in the United States. ECB Working Paper.

Bonadio, B., Fischer, A., and Sauré, P. (2019). The speed of exchange rate pass-through. *Journal of the European Economic Association*, 18:506–538.

Borio, C., Disyatat, P., Xia, D., and Zakrajšek, E. (2021). Monetary policy, relative prices and inflation control: Flexibility born out of success. In *BIS Quarterly Review*. Bank for International Settlements. September.

Borio, C., Lombardi, M., Yetman, J., and Zakrajšek, E. (2023). The two-regime view of inflation. BIS Papers No 133.

Born, B., Dalal, H., Lamersdorf, N., and Steffen, S. (2023). Monetary policy in the age of social media: A twitter-based inflation analysis. *Available at SSRN*.

Bybee, J. L. (2023). The ghost in the machine: Generating beliefs with large language models. *arXiv preprint arXiv:2305.02823*.

Campa, J. M. and Goldberg, L. S. (2005). Exchange rate pass-through into import prices. *The Review of Economics and Statistics*, 87(4):679–690.

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

Chauvet, M. and Hamilton, J. D. (2006). Dating business cycle turning points. *Contributions to Economic Analysis*, 276:1–54.

Clarida, R., Gali, J., and Gertler, M. (1999). The science of monetary policy: A new keynesian perspective. *Journal of Economic Literature*, 37(4):1661–1707.

Cloyne, J., Ferreira, C., Froemel, M., and Surico, P. (2023). Monetary policy, corporate finance, and investment. *Journal of the European Economic Association*, 21(6):2586–2634.

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

Conrad, C., Enders, Z., and Glas, A. (2022). The role of information and experience for households' inflation expectations. *European Economic Review*, 143:104015.

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

De Fiore, F., Lombardi, M., and Pierres-Tejada, A. (2024). Fiscal stimulus plans and households' expectations. BIS Working Papers No 1238.

De Santis, R. A. and Tornese, T. (2023). Energy supply shocks' nonlinearities on output and prices. ECB Working Paper No. 2834.

Dedola, L., Gautier, E., Nakov, A., Santoro, S., Veirman, E. D., Henkel, L., and Fagandini, B. (2023). Some implications of micro price setting evidence for inflation dynamics and monetary transmission. ECB Occasional Paper No. 321.

Dedola, L., Kristoffersen, M. S., and Züllig, G. (2022). The extensive and intensive margin of price adjustment to cost shocks: Evidence from Danish multiproduct firms. *American Economic Journal: Macroeconomics, forthcoming*.

Devereux, M. B. (2006). Exchange rate policy and endogenous price flexibility. *Journal of the European Economic Association*, 4(4):735–769.

Devereux, M. B. and Yetman, J. (2002). Menu costs and the long-run output-inflation trade-off. *Economic Letters*, 76:95–100.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805.*

Die Ärzte (2003). Deine Schuld. Song on the album *Geräusch*. Hot Action Records.

Ellingsen, J., Larsen, V. H., and Thorsrud, L. A. (2022). News media versus FRED-MD for macroeconomic forecasting. *Journal of Applied Econometrics*, 37(1):63–81.

Enders, A., Enders, Z., and Hoffmann, M. (2018). International financial market integration, asset compositions, and the falling exchange rate pass-through. *Journal of International Economics*, 110:151–175.

European Central Bank (ECB) (2024). ECB consumer expectations survey: Methodological information on the survey. Technical report. Available at: `https://www.ecb.europa.eu/`.

Fabiani, S., Druant, M., Hernando, I., Kwapil, C., Landau, B., Loupias, C., Martins, F., Mathä, T. Y., Sabbatini, R., Stahl, H., and Stokman, A. C. (2005). The pricing behaviour of firms in the euro area: New survey evidence. ECB Working Paper No. 535.

Galeone, P. and Gros, D. (2023). The ECB in the face of an unprecedented energy price shock. IEP@BU Policy Brief October.

Garzon, A. J. and Hierro, L. A. (2021). Asymmetries in the transmission of oil price shocks to inflation in the eurozone. *Economic Modelling*, 105:105665.

Gautier, E., Conflitti, C., Faber, R. P., Fabo, B., Fadejeva, L., Jouvanceau, V., Menz, J.-O., Messner, T., Petroulas, P., Roldan-Blanco, P., Rumler, F., Santoro, S., Wieland, E., and Zimmer, H. (2024). New facts on consumer price rigidity in the euro area. *American Economic Journal: Macroeconomics*, 16(4):386–431.

Giacomini, R., Kitagawa, T., and Read, M. (2022). Narrative restrictions and proxies. *Journal of Business and Economic Statistics*, 40(4):1415–1425.

Gödl-Hanisch, I. and Menkhoff, M. (2023). *Frims' pass-through dynamics: A survey approach.* Centre for Economic Policy Research.

Golosov, M. and Lucas, R. E. (2007). Menu costs and Phillips curves. *Journal of Political Economy*, 115(2):171–199.

Gonçalves, S., Herrera, A. M., Kilian, L., and Pesavento, E. (2024). State-dependent local projections. *Journal of Econometrics*, 244(2):105702.

Hall, R. E. (2023). A major shock makes prices more flexible and may result in a burst of inflation or deflation. NBER Working Paper No. 31025.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57:357–384.

Hamilton, J. D. (1990). Analysis of time series subject to changes in regime. *Journal of Econometrics*, 45:39–70.

Ilori, A. E., Paez-Farrell, J., and Thoenissen, C. (2022). Fiscal policy shocks and international spillovers. *European Economic Review*, 141:103969.

Informationsgemeinschaft zur Feststellung der Verbreitung von Werbeträgern e.V. (IVW) (2023). Ivw-auflagenliste q3 2023. Corporate report.

Informationsgemeinschaft zur Feststellung der Verbreitung von Werbeträgern e.V. (IVW) (2025). Ivw-auflagenliste q4 2024. Corporate report.

International Monetary Fund (2006). How has globalization affected inflation? In *World Economic Outlook 04/06*, chapter III.

Jarociński, M. and Karadi, P. (2020). Deconstructing monetary policy surprises - the role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.

Kapetanios, G. and Tzavalis, E. (2010). Modeling structural breaks in economic relationships using large shocks. *Journal of Economic Dynamics and Control*, 34(3):417–436.

Khalil, M. and Lewis, V. (2024). Product turnover and endogenous price flexibility in uncertain times. Deutsche Bundesbank Discussion Paper No. 14/2024.

Kimura, T. and Kurozumi, T. (2010). Endogenous nominal rigidities and monetary policy. *Journal of Monetary Economics*, 57(8):1038–1048.

Kling, M.-U. (2019). *Qualityland*. Ullstein Buchverlage, Berlin. Novel.

Krippner, L. (2013). Imeasuring the stance of monetary policy in zero lower bound environments. *Economics Letters*, 118:135–138.

Lane, P. (2022). Inflation diagnostics. *The ECB Blog, November*, 25.

Latifi, A., Naboka-Krell, V., Tillmann, P., and Winker, P. (2024). Fiscal policy in the Bundestag: Textual analysis and macroeconomic effects. *European Economic Review*, 168:104827.

Leduc, S., Sill, K., and Stark, T. (2007). Self-fulfilling expectations and the inflation of the 1970s: Evidence from the Livingston survey. *Journal of Monetary Economics*, 54(2):433–459.

Lewis, D. J. and Mertens, K. (2022). A robust test for weak instruments with multiple endogenous regressors. FRB of New York Staff Report No. 1020.

Li, D., Hong, Y., Wang, L., Xu, P., and Pan, Z. (2022). Extreme risk transmission among bitcoin and crude oil markets. *Resources Policy*, 77:102761.

Lieb, L., Jassem, A., Almeida, R. J., Baştürk, N., and Smeekes, S. (2025). Min (d) ing the president: A text analytic approach to measuring tax news. *American Economic Journal: Macroeconomics*, 17(2):285–314.

Liu, F. T., Ting, K. M., and Zhou, Z.-H. (2012). Isolation-based anomaly detection. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 6(1):1–39.

Lloyd, S. and Manuel, E. (2024). *Controls, Not Shocks: Estimating Dynamic Causal Effects in Macroeconomics*. London School of Economics and Political Science.

Mack, C. E. (1946). Escalator clauses in purchase contracts. *The Management Review*, 35(8):350–352.

Mackowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803.

Meeks, R. and Monti, F. (2023). Heterogeneous beliefs and the Phillips curve. *Journal of Monetary Economics*, 139:41–54.

Menz, J.-O. and Poppitz, P. (2013). Households' disagreement on inflation expectations and socioeconomic media exposure in Germany. Bundesbank Discussion Paper.

Moessner, R. (2025). Effects of inflation expectations on inflation. *National Institute Economic Review*, page 1–9.

Montiel Olea, J., Qian, E., Wolf, C. K., et al. (2025). Local projections or vars? A primer for macroeconomists. *arXiv e-prints*, pages arXiv–2503.

Montiel Olea, J. L. and Plagborg-Møller, M. (2021). Local projection inference is simpler and more robust than you think. *Econometrica*, 89(4):1789–1823.

Muehlegger, E. and Sweeney, R. L. (2022). Pass-through of own and rival cost shocks: Evidence from the U.S. fracking boom. *The Review of Economics and Statistics*, 104(6):1361–1369.

Oxford University Press (n.d.). Inflation, n. Retrieved July 18, 2025.

Paciello, L. and Wiederholt, M. (2014). Exogenous information, endogenous information, and optimal monetary policy. *The Review of Economic Studies*, 81(1 (286)):356–388.

Panetta, F. (2022). Normalising monetary policy in non-normal times. Speech by Fabio Panetta at a policy lecture hosted by the SAFE Policy Center at Goethe University and the Centre for Economic Policy Research (CEPR).

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Picault, M., Pinter, J., and Renault, T. (2022). Media sentiment on monetary policy: Determinants and relevance for inflation expectations. *Journal of International Money and Finance*, 124:102626.

Ramey, V. A. and Zubairy, S. (2018). Government spending multipliers in good times and in bad: Evidence from US historical data. *Journal of Political Economy*, 126(2):850–901.

Romer, C. D. and Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 94:1055–1084.

Shapiro, A. H., Sudhof, M., and Wilson, D. J. (2022). Measuring news sentiment. *Journal of Econometrics*, 228(2):221–243.

Shiller, R. J. (2017). Narrative economics. *American Economic Review*, 107(4):967–1004.

Shiller, R. J. (2020). Popular economic narratives advancing the longest us expansion 2009–2019. *Journal of Policy Modeling*, 42(4):791–798.

Sinn, H.-W. (2021). Wir haben Inflation, wie sie in einem Menschenleben einmal vorkommt. *Finanzen100*. December 15.

Smets, F. (2003). Maintaining price stability: How long is the medium term? *Journal of Monetary Economics*, 50(6):1293–1309.

Stock, J. H. and Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal*, 128(610):917–948.

Taylor, J. B. (2000). Low inflation, pass-through and the pricing power of firms. *European Economic Review*, 44:1389–1408.

Tillmann, P. (2025). Monetary policy shocks and the fiscal talk in the Bundestag. *Available at SSRN 5260481*.

Trebbi, G. (2024). Inflation narratives and expectations. *Available at SSRN 4881324*.

Vavra, J. (2014). Inflation dynamics and time-varying volatility: New evidence and an Ss interpretation. *The Quarterly Journal of Economics*, 129(1):215–258.

Walsh, C. E. (2003). *Monetary Theory and Policy*. Cambridge, MA: MIT Press. Second Edition.

Wang, Z., Xie, Q., Ding, Z., Feng, Y., and Xia, R. (2023). Is ChatGPT a good sentiment analyzer? A preliminary study. *arXiv preprint arXiv:2304.04339*.

Weber, M., Candia, B., Afrouzi, H., Ropele, T., Lluberas, R., Frache, S., Meyer, B. H., Kumar, S., Gorodnichenko, Y., Georgarakos, D., Coibion, O., Kenny, G., and Ponce, J. (2023). Tell me something I don't already know: Learning in low and high-inflation settings. *NBER Working Paper*, (31485).

Weinhagen, J. (2002). An empirical analysis of price transmission by stage of processing. *Monthly Labor Review*, 125:3.

Weinhagen, J. C. (2016). Price transmission within the producer price index final demand-intermediate demand aggregation system. *Monthly Labor Review*, 139:1.

Woodford, M. and Walsh, C. E. (2005). Interest and prices: Foundations of a theory of monetary policy. *Macroeconomic Dynamics*, 9(3):462–468.

Wu, J. C. and Xia, F. D. (2020). Negative interest rate policy and the yield curve. *Journal of Applied Econometrics*, 35(6):653–672.