

Discussion Paper Series – CRC TR 224

Discussion Paper No. 724  
Project C 02

# From Tweets to Transactions: High-Frequency Inflation Expectations, Consumption, and Stock Returns

Benjamin Born<sup>1</sup>  
Nora Lamersdorf<sup>2</sup>  
Jana-Lynn Schuster<sup>3</sup>  
Sascha Steffen<sup>4</sup>

January 2026

<sup>1</sup> University of Bonn, CEPR, CESifo, and ifo Institute, Email: b.born@uni-bonn.de

<sup>2</sup> BI Norwegian Business School, Email: nora.lamersdorf@bi.no

<sup>3</sup> : Frankfurt School of Finance & Management, Email: j.schuster@fs.de

<sup>4</sup> Frankfurt School of Finance & Management, Email: s.steffen@fs.de

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)  
through CRC TR 224 is gratefully acknowledged.

# From Tweets to Transactions: High-Frequency Inflation Expectations, Consumption, and Stock Returns\*

Benjamin Born, Nora Lamersdorf, Jana-Lynn Schuster, Sascha Steffen

January 2026

## Abstract

Using modern natural language processing, we construct a high-frequency inflation expectations index from German-language tweets. This index closely tracks realized inflation and aligns even more closely with household survey expectations. It also improves short-run forecasts relative to standard benchmarks. In response to monetary policy tightening, the index declines within about a week, with the effects concentrated in tweets by private individuals and during the recent period of elevated inflation. Using 117 million online transactions from German retailers, we show that higher inflation expectations are followed by lower household spending on discretionary goods. By linking these shifts in demand to stock returns, we find that, during periods of elevated inflation, firms operating in discretionary sectors experience significantly lower stock returns when inflation expectations rise. Thus, our Twitter-based index provides market participants and policymakers with a timely tool to monitor inflation sentiment and its economic consequences.

*Keywords:* Inflation expectations, social media (Twitter/X), large language models (LLMs), NLP, household consumption, stock returns, monetary policy

*JEL-Codes:* E31, D84, E58, C45, C81

---

\*Born: University of Bonn, CEPR, CESifo, and ifo Institute, b.born@uni-bonn.de, Lamersdorf: BI Norwegian Business School, nora.lamersdorf@bi.no, Schuster: Frankfurt School of Finance & Management, j.schuster@fs.de, Steffen: Frankfurt School of Finance & Management, s.steffen@fs.de. We thank Hrishbh Dalal for excellent research assistance on earlier drafts of this paper. We also thank Sarah Arndt (discussant), Felix Chopra (discussant), Francesco D'Acunto (discussant), Sanjiv Das, Philip Schnorpfeil, Michael Weber, as well as numerous seminar and conference audiences for helpful comments and suggestions. We thank the BMWK for funding within the safeFBDC-project. Survey data on inflation expectations are provided by Deutsche Bundesbank under the project number 2022\0105 and by the Gesellschaft für Konsumforschung (GfK). Support by the German Research Foundation (DFG) through CRC TR 224 (Project C02) is gratefully acknowledged.

## 1 Introduction

The post-COVID inflation surge highlighted the central role of household inflation expectations in shaping both macro outcomes and the transmission of monetary policy. Expectations formed through lived price experiences and salient supply shocks can spill over into price setting and inflation dynamics beyond the initially affected sectors (Acharya et al., 2025). They also constitute a key channel through which central banks communicate and potentially anchor policy (see, e.g., Baumann et al., 2021; Coibion et al., 2020). Yet we still lack timely, high-frequency measures of household inflation expectations that support clean event-time analysis.

Surveys remain the benchmark for household beliefs (Coibion et al., 2022), but they are costly, infrequent, and arrive with delays. Survey responses can be sensitive to question wording, priming, and limited sample sizes, and their monthly frequency typically precludes within-month event studies. Market-based measures (e.g., inflation swaps) are available intraday, but they reflect institutional investors' beliefs and hedging demand rather than households', and they embed time-varying risk and liquidity premia (Coibion et al., 2022; D'Acunto et al., 2024; Binder et al., 2024).

We address these limitations by constructing a high-frequency proxy for household inflation expectations from German-language tweets. We identify inflation-related tweets and use modern natural language processing (NLP) to map tweet-level language into a daily expectations series. We validate the measure by comparing it to survey- and market-based expectations and realized inflation, and by benchmarking its short-horizon forecasting performance against standard alternatives. We then use the daily series in an event-study design around high-frequency ECB monetary policy shocks to trace how household inflation expectations adjust in the days following policy news.

Finally, we combine the index with 117 million online transactions from German retailers and firm-level stock returns to study how expectation shifts during the post-COVID period translate into household spending adjustments—particularly a contraction in discretionary goods—and the pricing of firms exposed to these demand shifts. Our contribution is threefold: (i) we introduce and validate a new high-frequency measure of household inflation expectations; (ii) we quantify how expectation shifts predict spending on discretionary goods using granular transaction data; and (iii) we link expectation-driven demand shifts to equity pricing, with effects strongest during periods of elevated inflation.

In greater detail, we use advanced natural language processing (NLP) on a comprehensive dataset of German-language tweets spanning January 2011 to June 2024 to construct a high-frequency inflation expectations index for Germany. Twitter is a leading social media platform that offers a rich, high-frequency snapshot of diverse individual perspectives.<sup>1</sup> To isolate relevant content from noise and bot interference, we employ unsupervised machine learning within the BERTopic framework (Grootendorst, 2022), filtering the corpus down to over 1.3 million tweets centered on inflation-related topics such as general price increases, energy costs, and monetary policy.<sup>2</sup>

To classify whether inflation expectations are increasing or decreasing, we adopt a hybrid machine learning approach that combines the reasoning capabilities of large language models (LLMs) with the efficiency of specialized classifiers. Specifically, we use the ChatGPT API to generate a high-quality labeled training set. We then use this set to fine-tune a pre-trained BERT model optimized for social media text. This process enables us to categorize the sentiment of hundreds of thousands of tweets with high accuracy while minimizing the biases inherent in manual labeling. Finally, we compile the tweets that refer to rising or falling inflation into a daily index.

When benchmarked against Germany’s CPI inflation and the Bundesbank Online Panel – Households (BOP-HH) inflation expectations, the Twitter-based inflation expectations index reveals a close alignment for most periods. Notably, the Twitter index mirrors both the surge and decline in realized inflation—such as in early 2015 and 2017. However, during the spring 2020 outbreak of the novel coronavirus, the index diverged from actual inflation, which fell sharply, while our measure remained relatively stable. By contrast, household inflation expectations stayed elevated during this period, aligning more closely with our index. This suggests that the index may capture short-term inflation expectations rather than actual inflation dynamics. Comparative evidence from the Gesellschaft für Konsumforschung (GfK) survey reinforces this interpretation. Our index closely tracks survey-based expectations in earlier periods and continues to do so during salient events, such as the ECB’s PEPP announcement in March 2020. Correlation analyses with CPI data and both BOP-HH and GfK inflation expectations reveal strong associations, with values reaching up to

---

<sup>1</sup>Twitter was rebranded as “X” in July 2023; we use “Twitter” for clarity and because it was the official name for most of our sample. We find no evidence of a structural break in our index around the rebranding or subsequent migration debates: its correlation with official household expectations remains high and stable through June 2024.

<sup>2</sup>We take rigorous steps to filter out bot activity and extraneous content. As shown in Appendix OA.2, our index remained robust during Twitter’s intensified bot removal efforts in 2018, indicating that our filtering removed artificial amplification without discarding economically relevant information.

0.95 for quantitative household expectations. When benchmarked against market-based measures of inflation expectations, the correlations are consistently lower but remain high. Taken together, these findings suggest that our Twitter-based index reliably proxies for households' short-term inflation expectations at a high frequency.

We then assess the predictive power of our Twitter-based inflation expectations index, showing that it adds additional explanatory value to survey-based inflation expectations and actual CPI, over and above traditional predictors such as lagged expectations and past CPI. Using three autoregressive AR(1) models as benchmarks, we compare the performance of models that incorporate a lag of our Twitter-based index against those that do not. Our results reveal that, in most cases, models incorporating our Twitter-based index outperform the benchmarks, with the most significant improvement observed in forecasting the BOP-HH's quantitative measure of inflation expectations. These results further validate the index as a useful high-frequency measure of short-term household inflation expectations.

Our Twitter-based inflation expectation index has a clear advantage over traditional sources of inflation expectation due to its high-frequency, real-time nature. This makes it particularly valuable for tracking shifts in inflation sentiment following relevant events, such as monetary policy announcements. Using intraday changes in two-year OIS rates around the ECB's press release window, as described in Altavilla et al. (2019), and a local projection framework, our analysis reveals that the index notably responds to shifts in the ECB's monetary policy. For example, when monetary policy tightens unexpectedly, our index declines after just over a week. This suggests that the ECB can influence public inflation expectations within days.

While recent studies find that households typically respond little to monetary policy communication (see, *e.g.* Coibion et al., 2020; D'Acunto et al., 2021), we find a measurable reaction in our data. When we restrict the index to tweets posted by private individuals, the response to ECB announcements remains surprisingly strong. Consistent with Afrouzi et al. (2025), this effect is particularly pronounced during periods of elevated inflation, when households seem to pay closer attention to inflation news and monetary policy decisions. In essence, during periods of elevated inflation, households appear to be more responsive to monetary policy announcements, aligning their inflation expectations accordingly. Therefore, our index presents a real-time tool for market participants and policymakers to assess their impact on public sentiment.

Having validated our index as a reliable real-time proxy for inflation expectations, we next examine how changes in these expectations affect the real economy through household consumption. Leveraging the high-frequency nature of our data, we combine our Twitter-based inflation expectation index with detailed information on online retail transactions of German households from January 2021 to March 2024. This information is obtained from the FinTech company *Grips Intelligence*. The dataset comprises more than 1,000 firms and over 600 million transactions across a broad set of product categories.<sup>3</sup>

The ideal experiment to identify the causal effect of inflation expectations would exogenously shift households' beliefs about future inflation and trace the resulting consumption response. Since targeted shocks are not feasible, we use the predetermined release dates of the consumer price index (CPI) as instruments for our Twitter-based inflation expectations index within a two-stage least squares (2SLS) framework. This provides a quasi-experimental setting that approximates random shocks to inflation beliefs. CPI release dates are fixed in advance and unrelated to households' contemporaneous consumption decisions. They raise public attention to inflation and thereby shift expectations. Our key identifying assumption is that CPI announcements affect consumption primarily by shifting inflation expectations. This is plausible because we explicitly control for contemporaneous financial-market reactions (DAX returns and bond yields), which helps absorb alternative channels such as wealth effects and monetary-policy news. Moreover, the high-frequency data allow us to estimate responses within a tight 12-day window around each CPI release.

Our results show that household consumption patterns systematically respond to changes in inflation expectations. When inflation expectations rise, particularly around CPI release days, households reduce spending on discretionary goods such as clothing, electronics, and lifestyle products. The decline is most pronounced six to nine days after the CPI announcement. Observed through the combination of Twitter-based expectations and real consumption data, these patterns reflect behavior in everyday economic conditions rather than in a controlled setting. This strengthens the external validity of our findings. Taken together, these results suggest that shifts in inflation

---

<sup>3</sup>*Grips Intelligence* is a FinTech company that specializes in e-commerce analytics. The firm collects detailed information on product sales, prices, and competitive activity across online platforms. We obtain daily transaction data for customers in Germany from them, including the website on which the purchase was made and the associated consumption category.

expectations can quickly affect the real economy through household spending decisions. This underscores the importance of real-time monitoring for policymakers and firms.

We then investigate how shifts in household consumption, as documented above, affect firms, with a focus on stock market performance. Using German stock prices and restricting our analysis to CPI release days—when investors receive a significant update on inflation-related information—we document two key results. First, during the recent period of elevated inflation, increases in our Twitter-based inflation expectations index are associated with a significant decline in stock returns on CPI announcement days. This finding is consistent with the evidence presented in Knox and Timmer (2025) showing that positive inflation surprises lead to lower equity valuations. This effect is economically meaningful: a one-standard-deviation increase in the index on CPI release days reduces the CDAX by 14 basis points.

Second, these effects are concentrated among firms in the consumer discretionary sector. During periods of elevated inflation, the daily returns of discretionary firms fall by approximately 11 basis points more than those of non-discretionary firms following a one-standard-deviation increase in the index on CPI announcement days. Furthermore, discretionary firms' return sensitivity becomes increasingly negative at higher inflation levels, indicating that financial markets discount these firms more heavily when inflation is high.

Taken together, the results show that our Twitter-based index captures inflation-related information that becomes especially important in periods of elevated inflation. In such environments, increased inflation expectations weaken consumer spending and lower valuations of firms exposed to discretionary demand. This provides a coherent transmission mechanism from household beliefs to real economic activity and financial markets.

**Related literature.** Our paper contributes to four strands of the literature. First, we add to the work on eliciting inflation expectations. Coibion et al. (2022) and D'Acunto et al. (2024) provide comprehensive overviews of how expectations are measured, emphasizing the reliance on surveys. We contribute by providing a high-frequency measure derived from unprompted statements at a large scale, overcoming the frequency and framing limitations of traditional surveys. Closely related is Angelico et al. (2022), who use Italian tweets to measure expectations. In contrast to their work, we employ advanced Large Language Models (LLMs) to minimize manual error, extend the analysis

to the recent high-inflation regime, and specifically link the index to real consumption and firm-level equity returns. While Müller et al. (2022) measure media coverage of inflation, our index captures the direction of expectations—whether households perceive prices as rising or falling.

Second, we contribute to the growing application of textual analysis and AI in finance. Bartov et al. (2018) and Adams et al. (2023) demonstrate that social media tone predicts earnings surprises and financial conditions. At the firm level, tweet sentiment has been shown to correlate with returns and volatility consistent with gradual information incorporation (Sul et al., 2017; Gu and Kurov, 2020). We extend this line of research by isolating a specific macroeconomic signal—household inflation expectations—from social media noise and tracing its transmission to firm-level valuations.

Third, we connect to the literature on the real effects of inflation expectations on consumption. Theoretical and empirical evidence on this link is mixed. On one hand, the intertemporal substitution channel suggests that higher expected inflation should accelerate spending, especially for durables (D'Acunto et al., 2021). On the other hand, higher inflation can reduce real disposable income or trigger precautionary savings, leading to reduced spending (Bachmann et al., 2015; Coibion et al., 2023; Christelis et al., 2020). Our high-frequency setting allows us to test these channels around information shocks. Consistent with the income and precautionary motive dominating during high-inflation regimes, we find that shocks to our Twitter-based index predict an immediate decline in discretionary spending.

Finally, we contribute to the literature linking inflation to asset prices. Flannery and Protopapadakis (2002) and Ang et al. (2012) identify inflation announcements as key macro-factors affecting equity returns, though recent evidence on the direction of this relationship is mixed. Chaudhary and Marrow (2024) find that higher market-based long-term inflation expectations were associated with higher stock returns during the 2000–2020 period, interpreting inflation as a signal of economic growth. In contrast, Knox and Timmer (2025) show that over a longer horizon (1977–2022), investors hold a “stagflationary view,” where positive inflation surprises lead to falling stock returns. Kroner (2025) further emphasizes that investor attention to CPI amplifies these price responses during high-inflation episodes. We bridge the gap between household sentiment and asset pricing by providing micro-evidence for this stagflationary channel: we show that the discretionary spending cut we document in our consumption data directly translates into lower equity returns for firms in discretionary sectors.

The rest of the paper is structured as follows: After presenting our data sources in section 2, we describe the methods to build our Twitter-based inflation expectation index in 3.1 before validating the index and showing the effects of monetary policy in sections 3.2 and 3.3, respectively. Sections 4 and 5 then look at the link between inflation expectations and private consumption and stock performance, respectively. Section 6 concludes.

## 2 Data sources

**Twitter data.** We download potentially relevant tweets using the Twitter Application Programming Interface (API) for academic research. After excluding retweets, we select all German-language tweets posted between the start of Twitter on March 21, 2006 and June 09, 2024, that contain at least one of the following keywords: *price, cost of living, high bill, inflation, expensive, gasoline price, high rent, low rent, energy costs, deflation, disinflation, sale, sell-off, low bill, low cost, cheap*.<sup>4</sup> This yields a data set of 12,014,955 tweets from over one million users, along with their metadata, such as the Twitter user’s biography, number of likes, retweets, etc.

Although this set of tweets contains keywords that are important in the context of inflation, they could also be used in other contexts. For example, the word “price” could be related to any form of advertisement, which could be considered noise for our purposes. Therefore, as we describe below, we need to reduce the amount of noise in our data set.

**Other data sources.** To compare our Twitter-based index with survey measures of households’ inflation expectations, we use two monthly data sources. The first one is the Bundesbank Online Panel Households (BOP-HH) from the Deutsche Bundesbank. This is an online survey conducted regularly since April 2020. Each month, a sample of between 2,500 and 5,000 individuals is surveyed.<sup>5</sup> We take the weighted mean of the answers to the question: “What do you think the rate of inflation/deflation will roughly be over the next twelve months as a measure of inflation expectations. We trim the data below -12% and above 12%, as is usually done by the Deutsche Bundesbank.<sup>6</sup> In addition, we use the BOP-HH survey measure of inflation perceptions, which

---

<sup>4</sup>The original list of German keywords can be found in Appendix OA.1.1.

<sup>5</sup>Data DOI: 10.12757/Bbk.BOPHH.202204.01, see Schmidt et al. (2022)

<sup>6</sup>We also run robustness checks for all our analyses, trimming the data below -24% and above 24%. The results are very similar (available on request).

was available quarterly until the end of 2020 and monthly starting in January 2021. The exact question is, “What do you think the rate of inflation or deflation in Germany was over the past twelve months?” We trim the data below -12% and above 12% as well. We access the data through the Scientific Use Files provided by the Bundesbank’s Research Data and Service Centre (RDSC).

Since the BOP-HH data only begins in April 2020, we use the micro data underlying the GfK Consumer Climate Indicator for Germany as a second source of inflation expectations. The GfK conducts a monthly survey in which they ask 2,000 individuals questions about inflation expectations, buying propensities, and personal economic situations. We are interested in the answer to the following question: “How will consumer prices evolve during the next twelve months compared to the previous twelve months?” Unlike the BOP-HH, respondents do not provide an exact number, but rather a qualitative response.<sup>7</sup> We take the share of respondents who answered “Prices will increase more” as a measure of inflation expectations.

To compare our index with market-based measures of inflation expectations, we use Euro Area 1-year and 5-year inflation swap rates and German 10-year breakeven inflation expectations, obtained from Bloomberg. Destatis, Germany’s Federal Statistical Office, provides monthly data on realized consumer price inflation (CPI). The exact CPI release dates are obtained from Investing.com. Data on ECB monetary policy surprises are taken from the dataset by Altavilla et al. (2019), which is regularly updated and covers our sample period up to June 2024.

We use anonymized, transaction-level data from *Grips Intelligence*, a FinTech company, covering online retail purchases in Germany from January 2021 to March 2024. The dataset includes time-stamped transactions by retailer, with customer locations inferred from IP addresses. To ensure a consistent sample and mitigate bias from firm entry or exit, we restrict our baseline analysis to firms reporting sales in at least 33 of the 39 months. This filtering yields a balanced panel of 212 firms and over 117 million transactions, mapped into 20 primary categories and 60 subcategories. Given that the dataset does not cover major grocery chains and is inherently dominated by non-essential spending, we focus our main analysis on firms that primarily sell discretionary goods. This restriction results in a final sample of 153 discretionary firms and approximately 106 million transactions. To account for broader financial conditions in the consumption analysis, we include daily control

---

<sup>7</sup>Specifically, respondents could answer, “Prices will increase more,” “Prices will increase by the same,” “Prices will increase less,” “Prices will stay the same,” or “Prices will decrease.”

variables: the DAX index from the London Stock Exchange Group (LSEG), the 10-year German government bond yield from the OECD, and the one-year Euro Area yield-curve spot rate from the ECB Statistical Data Warehouse.

To investigate the relationship between our index and the stock market, we employ two sets of equity data. First, to capture the aggregate market response, we use daily returns on the German CDAX (Composite DAX) index obtained from Bloomberg. Second, to analyze cross-sectional heterogeneity, we compile daily stock prices for 210 German firms from LSEG, along with their industry classification according to the Global Industry Classification Standard (GICS).<sup>8</sup> For the firm-level analysis, we match these stock prices with quarterly fundamentals from Compustat to control for firm size, profitability (ROA, ROE), and capital structure (leverage, short-term debt).

### 3 The Twitter-based inflation expectations index

#### 3.1 Building the index

Our Twitter-based inflation expectations index is built in three steps, once all the tweets that are potentially related to inflation have been downloaded. First, we use topic modeling to clean the remaining tweets of noise and pin down the ones that are actually about inflation. Next, we classify the tweets belonging to an inflation topic as referring to increasing or decreasing inflation. Finally, we aggregate these two classes of tweets into a daily index.

##### 3.1.1 Selecting the relevant tweets

To filter out irrelevant tweets, we rely on topic modeling to detect tweets that actually deal with inflation. Cleaning the tweets of topics not directly related to inflation also eliminates a significant amount of Twitter bot activity, which is often associated with advertising.<sup>9</sup> However, before splitting the tweets into different topics, we exploit a specific feature of Twitter bots: they often repeat tweets verbatim or post extremely similar ones. For example, we noticed that many tweets contain the

---

<sup>8</sup>We limit the sample to the constituents of the DAX, MDAX, SDAX, and Deutsche Börse Prime Standard as of September 2025 to exclude delisted stocks. Delistings may depress returns for reasons unrelated to our mechanism and in ways that differ systematically across firms, thereby potentially biasing our estimates.

<sup>9</sup>Twitter bots are software applications that control Twitter accounts via the Twitter API and can tweet autonomously. Since they can perform many Twitter activities, such as tweeting, retweeting, liking, or direct messaging, on a large scale, they can influence many users. They can be used for (mis)information campaigns or advertising purposes, for example.

same sentence(s) with a link that varies from tweet to tweet. Therefore, first, we clean the tweets by removing hashtags, user mentions, unnecessary white spaces, and importantly, links. Then, we drop duplicates based on the cleaned text. This simple approach already removes more than two million tweets, around 18% of the total. This approach ensures that we retain relevant content from these bot activities that could influence people's inflation expectations, while preventing our index from becoming biased by their volume.<sup>10</sup>

We apply the topic-modeling technique BERTopic (Grootendorst, 2022) to classify the remaining tweets into different topics. BERTopic builds on pre-trained transformer-based language models.<sup>11</sup> BERTopic outperforms traditional topic models, such as Latent Dirichlet Allocation (LDA), in several ways: it generates more coherent and interpretable topics, can handle large text volumes, and is computationally more efficient. We use the TwHIN-BERT model by Zhang et al. (2023) for text embeddings because it is specifically trained on multilingual Twitter data, making it well-suited for our German-language corpus. The model reduces the high-dimensional embeddings, clusters them, and extracts topic representations using a class-based TF-IDF weighting scheme. This yields easy-to-interpret topics. Since the brevity of tweets can pose a challenge for topic models, we combine five consecutive tweets per user into one document during the training process. Furthermore, we clean the text by removing web links, user mentions, and symbols. Setting the number of topics to 150 yields granular, well-separated themes that capture various aspects of consumer price inflation rather than unrelated topics, such as cryptocurrencies or gold prices. Based on the most important words within each topic, we manually identify 19 topics that clearly relate to inflation.

Next, we use our trained topic model to infer the most likely topic for each tweet in our text corpus. Then, we select the tweets assigned to one of the 19 inflation topics. These 1,357,609 tweets comprise around 11% of all the tweets we downloaded. For visualization, we group these 19 topics into four sub-topics. The word clouds of these sub-topics are shown in Figure 1 based on the frequency with which the words appear in the tweets belonging to these sub-topics. The larger a word appears in a word cloud, the more important it is for the specific topic.

Based on the most significant keywords, we can distinguish between four main topics: general inflation, energy prices, inflation in the context of (monetary) policy, and inflation in the context of

---

<sup>10</sup>See Online Appendix OA.2 on details on our bot-detection procedure, bot features, and the impact of bot activity on the index.

<sup>11</sup>A detailed description of the BERTopic pipeline and implementation choices is provided in Online Appendix OA.1.

housing. Figure 2 shows the number of tweets per year addressing these four categories of inflation topics over time. Note that the data for 2024 covers the first five months of the year. The general inflation topic is the largest, accounting for about 64% of all inflation-related tweets and increasing sharply in 2022. Discussions on Twitter about energy prices and (monetary) policy in the context of inflation have been present throughout the years, but naturally increased in 2021 and especially in 2022 compared to previous years.

### 3.1.2 Classifying the tweets

We now need to classify the selected tweets based on whether they describe increasing or decreasing inflation or take a neutral stance on it.<sup>12</sup> Earlier approaches, such as Angelico et al. (2022), rely on manually constructed dictionaries of bi- and trigrams labeled as “up” or “down.” While this method is intuitive, it becomes infeasible for our corpus, which contains over ten million unique n-grams. Furthermore, it remains subjective and context-insensitive because it focuses on individual word sequences rather than the full meaning of a tweet. We adopt a supervised machine learning approach based on TwHIN-BERT (Zhang et al., 2023), a transformer model pre-trained on multilingual Twitter data. We fine-tune TwHIN-BERT for our specific classification task. To generate the labeled data necessary for this process, we combine a small, manually annotated sample with additional tweets labeled using the gpt-3.5-turbo model of ChatGPT. The tweets are categorized into three classes: increasing, decreasing, and other. This approach enables the model to capture the semantic context of entire tweets, resulting in more accurate and scalable classifications than those achieved with dictionary-based methods.

The resulting model performs well. It achieves an accuracy of about 0.76 on the validation set and 0.64 when evaluated against manually labeled data. When applied to our full corpus, the final model yields 405,184 tweets indicating increasing inflation or prices, 86,666 tweets indicating decreasing inflation or prices, and 865,759 tweets assigned to the *other* category. This classification step serves as an additional filter to retain only tweets that are clearly related to inflation, while assigning ambiguous cases to the residual *other* category.

---

<sup>12</sup>For a more detailed description of this step, see Online Appendix OA.1.3.

### 3.1.3 Aggregating the index

To create an index from the labeled tweets, we first calculate the daily sum of tweets in the increasing and decreasing inflation classes, respectively. This gives us an up index and a down index for each day. To obtain a daily index, we subtract the down index from the up index. That is, for each day  $t$ , the index is given by  $Index_t = UpIndex_t - DownIndex_t$ .

Figure 3 shows the daily index from April 28, 2007, when the first inflation-related tweet was posted, to June 09, 2024. Since the volume of inflation-related tweets was very low at the beginning of the sample, we start our analyses in the remainder of this paper on January 1, 2011, when the volume of tweets began to increase. In the earlier years of the sample, the index falls below 0 on several days, particularly at the end of 2014 and the beginning of 2015. During this time, inflation in Germany was close to 0 and, in January 2015, fell below 0. However, in the summer of 2022, the index increases significantly as inflation rises, before decreasing again toward the end of the year.

Figure 4 zooms in on the recent period of elevated inflation from January 2021 to June 2024. It marks selected events that are likely to affect inflation expectations. The index spikes on days when there is inflation news (e.g., CPI flash releases) and reacts to global shocks, as well as to monetary and fiscal policy announcements. The largest spike occurred in September 2022. Russia's halt of gas flows through Nord Stream 1 raised fears of shortages and soaring energy bills. Later that month, undersea explosions damaged the Nord Stream pipelines, further increasing supply uncertainty. In response, the German government introduced a €65 billion relief package, followed by a €200 billion "energy shield." Meanwhile, the ECB raised policy rates by 75 basis points and voiced strong concerns about the inflation outlook. Together, the energy supply shock and policy responses signaled severe and persistent price pressures, pushing inflation to the forefront of public debate. Our index spiked accordingly.

## 3.2 Validating the index

In this section, we validate our Twitter-based index by comparing it with realized CPI data and survey-based measures of inflation perceptions and expectations. We also assess its forecasting performance.

### 3.2.1 Benchmarking against CPI and established measures of inflation expectations

Figure 5a shows the time series of the monthly averages of the Twitter-based index (left y-axis) alongside CPI inflation in Germany and inflation expectations for the next 12 months from the Bundesbank Online Panel – Households (BOP-HH) (right y-axis) from January 2011 to June 2024 and from April 2019 to June 2024, respectively. Since the latter are reported monthly, we calculated monthly averages of our index. Furthermore, to facilitate visual comparison, we standardize the daily series by dividing it by three times the standard deviation before taking averages. We use a rolling window standard deviation with a window length of ten years because the level and variance of our index differ greatly between the beginning and end of the sample period.<sup>13</sup> For most of our sample period, our index closely captures the evolution of inflation. Notably, the Twitter-based index mimics both increases, such as at the beginning of 2017, and decreases, such as at the beginning of 2015, in inflation. The large increase in inflation from 2021 onward and the subsequent decline from summer 2022 onward are also present in our index. However, in spring 2020, during the outbreak of the pandemic, the Twitter-based index and CPI data diverge significantly. While our index remains within the previous years' range and even increases slightly, actual inflation in Germany falls sharply to its lowest value in the entire sample period.

However, as shown in Figure 5a, while actual inflation fell below 0% in spring 2020, quantitative inflation expectations remained around 3%. Accordingly, our index tracks this expectations measure closer than the CPI. Moreover, the challenges in measuring CPI inflation during the pandemic were amplified by pronounced shifts in consumption patterns (Cavallo, 2024; Kouvas et al., 2020).

Since the BOP-HH data on quantitative inflation expectations only became available regularly starting in April 2020 (with three pilot waves in 2019), we complement this analysis using the qualitative GfK survey, which provides a longer time series from January 2011 to May 2023. Figure 5b plots the share of respondents expecting prices to increase over the next 12 months as an alternative expectations measure. Throughout the entire period—and importantly, also in earlier years—the figure shows that our index closely mirrors the evolution of GfK inflation expectations. As with the BOP-HH data, the share of respondents in the GfK survey expecting higher prices rose sharply in April and May 2020 before declining again in subsequent months. In contrast,

---

<sup>13</sup>Using a rolling window standard deviation has the additional advantage of making our index less dependent on the exact sample period used to compute it and easier to update.

our index does not display such a pronounced spike and does not fall with actual inflation. This divergence appears to be driven, at least in part, by elevated Twitter discussions surrounding the ECB's announcement of the Pandemic Emergency Purchase Programme (PEPP) on March 18, 2020. Table 3 reports in the upper panel correlations across several subsamples, comparing our Twitter-based index with actual CPI data and survey-based measures of inflation expectations. The last column also includes correlations with BOP-HH inflation perceptions. For the pre-COVID period (2011–2019), correlations with CPI and GfK inflation expectations are already relatively high (0.63 and 0.66, respectively). In the more recent period from 2020 to mid-2024, the index continues to correlate strongly with benchmark measures, particularly with BOP-HH expectations (0.94) and actual CPI (0.83). Focusing on 2020–2022, which includes the inflation surge, correlations are especially pronounced for CPI (0.88), BOP-HH expectations (0.95), and perceptions (0.85). Even in the most recent period (2023–2024), correlations remain strong, especially with quantitative BOP-HH measures (0.88 for expectations and 0.84 for perceptions).

Overall, the correlation with qualitative GfK expectations weakens in the later sample, but our Twitter-based index continues to closely track quantitative BOP-HH inflation expectations and CPI across all periods. This indicates that the index captures short-run movements in households' inflation expectations rather than merely reflecting the inflation spike of 2021–2022.

The lower panel of Table 3 compares our index with market-based measures of inflation expectations.<sup>14</sup> While correlations with inflation swap rates and breakeven inflation expectations are positive throughout, they are systematically lower than those with survey-based household expectations. Correlations strengthen during the 2020–2022 inflation surge, reaching 0.92 for the 1-year Euro Area swap rate, and remain sizeable in 2023–2024, particularly for short-horizon measures (e.g., 0.78 for the EA 1-year swap).

Taken together, these patterns show that our Twitter-based index co-moves with both survey-based and market-based expectations, but most strongly with household short-term expectations. This is intuitive, as the average Twitter user more closely resembles the population represented in household surveys than financial market participants. We therefore view our index as a reliable high-frequency proxy for households' short-term inflation expectations.

---

<sup>14</sup>All correlations are based on monthly averages of the respective series. Correlations computed from daily data are somewhat lower, potentially reflecting the higher noise and idiosyncratic day-to-day variation in both market-based measures and our Twitter-based index.

### 3.2.2 Can the index explain actual CPI data and inflation expectations?

In this section, we examine whether our Twitter-based index helps explain household inflation expectations and actual CPI data, and whether it can therefore serve as an alternative, and especially an earlier, indicator.<sup>15</sup> We assess whether the index contains information beyond what is already captured by lagged values of expectations and inflation. Specifically, we regress BOP-HH and GfK inflation expectations, as well as actual CPI inflation, on their own lagged values, one lag of CPI inflation, and our Twitter-based index.

To align the timing of the index with the survey fieldwork, we do not compute monthly averages over the entire month. Instead, we average only the first 16 days of each month, since approximately 90% of BOP-HH respondents complete the survey after the 16th. This ensures that the Twitter-based measure reflects the information set available to households when forming their expectations.

We estimate the following baseline specification:

$$y_t = \alpha + \beta_1 \text{Index}_t + \beta_2 y_{t-1} + \beta_3 \text{CPI}_{t-1} + u_t , \quad (1)$$

where  $y_t$  denotes either survey-based inflation expectations or actual CPI inflation,  $\text{Index}_t$  is the average of our Twitter-based index over the first half of month  $t$ , and  $\text{CPI}_{t-1}$  is the consumer price index in the previous month.

We run four regressions with the following dependent variables: BOP-HH inflation expectations ( $E_t^{BOP} \pi_{t,t+12}$ ), GfK inflation expectations ( $E_t^{GfK} \pi_{t,t+12}$ ), and actual CPI inflation ( $\text{CPI}_t$ ). For the CPI regressions, we estimate two variants: one that includes only lagged CPI as a control, and a second specification that additionally includes lagged BOP-HH and GfK inflation expectations to account for survey-based information that may help predict CPI movements.

Table 4 reports the regression results. All variables have been standardized such that, e.g., a one-standard-deviation increase in our index is associated with a 0.29-standard-deviation increase in BOP-HH inflation expectations. The table shows positive and significant coefficients for our Twitter-based index for all regressions, indicating that it has additional explanatory power for both inflation expectations and actual CPI beyond their lags and, in the case of the former, actual inflation data. Even when regressing actual CPI data on their lagged values and both measures of

---

<sup>15</sup>Since correlations are highest for survey-based inflation expectations, we focus on these in the following analysis.

inflation expectations, our index provides additional explanatory power (column four). Including our inflation expectations index in these four regressions increases the Adjusted  $R^2$  by around 0.03. Thus, our Twitter-based index provides some additional information about survey-based measures of inflation beyond existing covariates, such as lagged survey variables or actual CPI data. It can also explain actual CPI data beyond past CPI data and traditional measures of inflation expectations.

### 3.2.3 Forecasting exercise

We now explore whether our Twitter-based index can be used to forecast inflation expectations and actual inflation. For both measures of inflation expectations and actual CPI data, we use three AR(1) autoregressive models as benchmark models, for which we choose a lag length of 1 according to the BIC criterion. Since the BOP-HH inflation expectations data only covers a relatively short time period, we first use as in-sample 12 monthly observations starting in January 2011 for the CPI and GfK inflation expectations, and in April 2020 for the BOP-HH inflation expectations. Then, we incrementally add one additional month to the in-sample. In addition to these benchmark models, we estimate three corresponding competing models. For these models, we add one lag of the Twitter-based index to the aforementioned AR(1) models, respectively. We produce one- and three-month-ahead forecasts for models including and excluding our index.

Table 5 shows the root mean squared errors (RMSE) for the three benchmark AR(1) models without our Twitter-based index in the first three rows. Rows four through six report the RMSE of the competing models relative to the benchmark models. Therefore, a value below one indicates that the competing model outperforms the benchmark. The table shows that including our Twitter-based index improves the forecasting performance for almost all three variables and forecasting horizons, as the relative RMSEs are below one. However, the index does not improve the forecast accuracy of the one-month-ahead ( $h = 1$ ) BOP-HH and GfK expectations forecasts, likely because the simple AR(1) model already performs very well at such short horizons, leaving little room for additional gains. The improvement is most pronounced for BOP-HH's quantitative measure of inflation expectations. Therefore, we conclude that our Twitter-based index is well-suited to capture consumer expectations about near-term inflation.<sup>16</sup>

---

<sup>16</sup>It could also improve the forecast of actual CPI data. However, we acknowledge that a proper forecasting of actual inflation would require the inclusion of many other variables.

### 3.3 Can monetary policy influence inflation expectations?

Figure 4 illustrates how our Twitter-based inflation expectations index responds to major events that are likely to shape inflation expectations. In this section, we study its response to monetary policy announcements—events with direct implications for inflation and expectations—to assess these effects systematically.

#### 3.3.1 Inflation expectations and monetary policy surprises

One key advantage of our inflation expectations index is its high-frequency nature, which sets it apart from traditional expectations measures, such as surveys. The index provides real-time insights into how inflation expectations shift in response to specific events. Monetary policy announcements are of particular interest because guiding expectations is a central objective of monetary policy. Therefore, assessing the effectiveness of such announcements at a relatively high frequency is highly relevant for central banks.

To analyze the effects of ECB monetary policy announcements on inflation expectations, we estimate local projections à la Jordà (2005) at a daily frequency for each horizon  $h$ :

$$y_{t+h} - y_{t-1} = \alpha_h + \beta_h \epsilon_t + \gamma_h X_t + u_{t+h} , \quad (2)$$

where  $y_t$  denotes the daily Twitter-based inflation expectations index,  $\alpha_h$  is a horizon-specific constant,  $\epsilon_t$  are intraday changes in 2-year OIS rates around the ECB press release window taken from Altavilla et al. (2019) (updated through June 2024), and  $X_t$  is a vector of controls including up to five lags of the index and last month's CPI. The coefficients  $\beta_h$  trace the response of the index to monetary policy surprises after  $h$  days.

Figure 6a shows that following a contractionary monetary policy surprise, our inflation expectations index declines within just over a week. A surprise of ten basis points leads to a drop of about 20 in the index, corresponding to an increase of 20 tweets anticipating lower rather than higher inflation. Therefore, the results suggest that the ECB can influence public inflation perceptions within days. This also confirms that our index captures economically relevant inflation events.

### 3.3.2 How do households respond to monetary policy during a surge in inflation?

Given recent findings suggesting that households generally react only weakly to central bank communication (see, for instance, Coibion et al., 2020; D’Acunto et al., 2021),<sup>17</sup> the responsiveness of our inflation expectations index to monetary policy announcements might appear surprising. This raises the question of whether our index reflects household reactions or instead captures media coverage or other institutional content related to monetary policy.

To explore this issue, we identify private individuals among Twitter users and construct sub-indices based on tweets from different groups of users. We classify accounts tweeting about inflation into categories such as private individuals, media organizations, firms, and others using the pre-trained TwHIN-BERT model, which we fine-tune for this specific task. We create most of the training data for this fine-tuning via the OpenAI API and supplement it with manually curated entries for media companies. In total, we classify 194,422 accounts as private users.<sup>18</sup>

Figure 6b shows the impulse responses of the inflation expectations index for private individuals. These responses closely resemble those of the overall index shown in Figure 6a: Following a contractionary monetary policy shock, the index declines after approximately one week.

The apparent responsiveness of private individuals seems to contradict earlier evidence.<sup>19</sup> However, recent work by Afrouzi et al. (2025) suggests that households become more informed about inflation dynamics when inflation is high. Similarly, Cavallo et al. (2017) show that individuals are less certain about inflation in a low-inflation environment than in periods of elevated inflation. Since inflation is a greater concern during these periods, individuals may pay closer attention to related news, including communication about monetary policy.

Therefore, we split the sample into two phases: 2011–2020, a period of low inflation, and 2021–2024, a period characterized by elevated inflation. Figures 6c and 6d show the corresponding regression results. During the low-inflation period (lower left panel), the index shows minimal

---

<sup>17</sup>Binder et al. (2024) find that households respond to some, but not many, FOMC announcements in their inflation expectations. The largest effects of these announcements on expectations occurred toward the end of their sample period in 2020 and 2021.

<sup>18</sup>The classification task distinguishes six categories: *private individuals*, *individual journalists*, *influencers*, *media organizations*, *business organizations*, and *other organizations*. The private individuals category—our main focus—appears to be classified most reliably, so we concentrate on this group. A full description of the classification procedure is provided in Online Appendix OA.3.

<sup>19</sup>An exception is Ehrmann and Wabitsch (2022), who show that non-experts also react to ECB communication. However, their sample is more selective, focusing on Twitter users who post specifically about the ECB. Their analysis does not examine whether communication shapes inflation expectations.

reaction to monetary policy surprises, with no significant responses at the 90% level. In contrast, during the elevated-inflation period (lower right panel), the index responds strongly. Within a week of a tightening shock, it declines significantly. This suggests that monetary policy announcements influence household inflation expectations.

In sum, our results indicate that our index not only captures important economic events, but also that households have become more responsive to monetary policy announcements in recent years. During periods of elevated inflation that affect daily life, for example, monetary policy appears to be able to shape inflation expectations within days. Thus, our index serves as a real-time barometer for policymakers seeking to gauge their influence on households.

## 4 Inflation expectations and consumer behavior

We have documented the ability to measure high-frequency inflation expectations using social media data, as well as how these expectations respond to significant events, such as shifts in monetary policy—particularly during periods of elevated inflation. Our next inquiry focuses on the tangible impact of these expectations on the real economy, with a specific focus on the interplay between inflation expectations and household consumption dynamics.

In theory, there are several channels that can explain the link between inflation expectations and household consumption. One such channel is the substitution channel, which operates through the consumption Euler equation. Higher expected inflation lowers the perceived real interest rate, incentivizing households to consume more and save less. Conversely, the income channel reflects concerns about the erosion of real purchasing power. If households expect nominal income growth to fall short of inflation, their expected real income will decline, prompting them to cut back on current spending and increase precautionary savings. In this section, we examine how these opposing forces shape the consumption response of discretionary goods to changes in inflation expectations, as measured by our inflation expectations index. Focusing on discretionary spending is particularly informative because Andreolli et al. (2025) show that discretionary sectors react most strongly to macroeconomic shocks such as monetary policy, highlighting their sensitivity to shifts in households' economic outlook.

For our analysis, we leverage data on online transactions spanning a vast spectrum of products and retailers in Germany from January 2021 to March 2024. We obtained the dataset from *Grips Intelligence*, a FinTech startup that provides e-commerce market intelligence by tracking product sales, pricing, and other competitive dynamics across online marketplaces.

To avoid sample bias arising from firm entry or exit, we restrict our analysis to firms that consistently report their sales. Specifically, we include transactions from German customers (based on their IP addresses) for firms that report in at least 33 of the 39 months covered by our dataset. This yields 212 firms and over 117 million transactions. We winsorize transactions at the 1st and 99th percentiles. Purchases are classified into 20 primary categories, which have a total of 60 subcategories. Figure 7 depicts these primary categories and their respective shares of total transactions. “Lifestyle”—primarily clothing—is the largest category, followed by “Computers, Electronics and Technology.”

As Figure 7 shows, the vast majority of transactions fall within discretionary consumption. This is plausible, as firms in discretionary industries may be more inclined to use transaction-level analytics services such as *Grips Intelligence*. Nevertheless, to ensure that we do not inadvertently mix discretionary and non-discretionary spending, which could yield heterogeneous effects, we restrict our analysis to firms that primarily sell discretionary products.<sup>20</sup> For this classification, we rely on the discretionary and necessity consumption categories definitions provided by Andreolli et al. (2025). Specifically, we supplied their category list to ChatGPT and used it to classify firms in our sample by letting ChatGPT examine their websites in batches.<sup>21</sup> Reducing the sample to firms that can be unambiguously classified as discretionary decreases the number of firms to 168 and the total number of transactions to approximately 107 million.<sup>22</sup>

Our objective is to estimate the causal effect of inflation expectations on discretionary consumption using our Twitter-based high-frequency index. While the index captures short-term fluctuations

---

<sup>20</sup>Differentiating between durable and non-durable goods, as in Coibion et al. (2022), would also be insightful. However, given that our data focus on online purchases, few items pertain to immediate consumption, such as grocery shopping. Most of the products are intended for intermediate or long-term use, such as clothing, cameras, or tableware. Even within the “Food and Drink” category, the items are not daily essentials, but rather more specialized goods, such as liquors, wines, and teas.

<sup>21</sup>We subsequently cross-checked firm names and classifications manually and reviewed cases where the automated classification appeared uncertain or counterintuitive. Overall, the procedure performed very well, with only a few firms requiring reclassification. For data protection reasons, firm names cannot be disclosed.

<sup>22</sup>Table OA.3 in the Online Appendix reports results using the full dataset, which are very similar to those obtained for our baseline sample.

in expectations, concerns about endogeneity and reverse causality remain, as expectations and consumption may influence each other or be jointly driven by unobserved factors. To address these concerns, we employ a two-stage least squares (2SLS) approach. We use calendar-determined CPI release days to instrument our inflation expectations index. These days are fixed in advance and are therefore exogenous to contemporaneous consumption dynamics. CPI release days should not directly affect consumption—conditional on our controls for financial market reactions and associated wealth effects—and therefore satisfy the exclusion restriction. However, they are salient events that increase attention to inflation and thereby shift inflation expectations, ensuring the relevance of the instruments.<sup>23</sup> We therefore use both preliminary and official CPI release days as instruments for our expectations index and estimate the following system of equations:

$$\begin{aligned} \Delta_h \log(C_{t+h}^i) &= \beta_1^h \widehat{\log(Index_t)} + \beta_2^h r_t^{DAX} + \beta_3^h \Delta Bund10Y_t \\ &\quad + \beta_4^h \Delta Yields1YEA_t + \gamma^h \mathbf{X}_{i,t} + \mu_i + \lambda_t + u_{i,t}^h, \end{aligned} \quad (3)$$

$$\begin{aligned} \log(Index_t) &= \alpha_1 \mathbf{1}_{\text{CPI release}} + \alpha_2 r_t^{DAX} + \alpha_3 \Delta Bund10Y_t \\ &\quad + \alpha_4 \Delta Yields1YEA_t + \gamma \mathbf{X}_{i,t} + \mu_i + \lambda_t + v_{i,t}, \end{aligned} \quad (4)$$

where  $\Delta_h \log(C_{t+h}^i)$  denotes the log-change in transactions between  $t + h$  and  $t - 1$ , that is, the percentage growth in transactions  $h$  days after the CPI release relative to the day before the release.  $Index_t$  is our daily Twitter-based inflation expectations index,  $r_t^{DAX}$  are daily DAX returns, and  $\Delta Bund10Y_t$  and  $\Delta Yields1YEA_t$  are the daily changes in the 10-year Bund and one-year Euro Area yield curve spot rate, respectively. These financial variables serve as controls to account for any direct effects that CPI announcements may have on consumption through changes in financial conditions—for instance, via stock market valuations or shifts in interest rates. The control vector  $\mathbf{X}_{i,t}$  includes one to seven lags of log transactions and the log index. We explicitly include firm fixed effects, denoted by  $\mu_i$ , as well as a set of time-related fixed effects,  $\lambda_t$ , which absorb weekend and month-year specific variation.<sup>24</sup> Standard errors are two-way clustered by date and firm.

<sup>23</sup>Binder et al. (2025) show that CPI announcements play a central role in shaping inflation expectations. Furthermore, in Appendix OA.4.1 we show that during the period from 2021 to 2024, our Twitter-based inflation expectations index rises significantly on CPI announcement days and remains elevated for about one week.

<sup>24</sup>Since  $\log(Index_t)$  varies only across days, the first-stage regression in Equation (4) is effectively identified at the daily level. The firm fixed effects  $\mu_i$  and firm-specific controls are included here merely for symmetry with the second-stage specification.

Table 6 reports the results. Column (1) presents the first-stage estimates corresponding to equation (4), confirming that our index is significantly elevated on CPI release days. The 0.21 coefficient implies that, conditional on the control variables, our inflation expectations index increases by 23 percent ( $= e^{0.21} - 1$ ) on a release day relative to non-release days. Furthermore, the first-stage  $F$ -statistics in our 2SLS specifications exceed 20, corroborating the relevance of our instrument.

Columns (2) through (6) report the 2SLS estimates for different horizons relative to the CPI release.<sup>25</sup> Exploiting the exogenous variation in inflation expectations induced by CPI releases, our IV estimates show that about one week after the announcement, higher inflation expectations lead to lower consumption. Specifically, a CPI-release-induced 10 percent increase in the index is associated with a 1.0 to 1.5 percent decline in discretionary consumption six to nine days after the release.<sup>26</sup>

Table 7 further breaks down this aggregate effect, showing that households reduce their spending across a broad range of categories in response to higher inflation expectations. Particularly pronounced declines are observed in typical discretionary sectors such as 'Lifestyle', 'Games', and 'Computers, Electronics & Techn.'. The significant negative effect in the 'Health' category may appear surprising; however, this can be plausibly explained by the fact that the online retailers in this segment primarily sell wellness, fitness, and beauty products rather than traditional medical or healthcare services. Taken together, the category-level results confirm that consumers specifically cut back on non-essential goods when their inflation expectations rise.

Overall, using our high-frequency inflation expectations index, we show that rising inflation expectations cause a reduction in discretionary consumption. This finding suggests that the real income channel—where households fear the erosion of purchasing power—dominates the intertemporal substitution channel in shaping the response of household spending to inflation expectations. Crucially, observing these dynamics through Twitter and real consumption data allows us to measure expectations in an uncontrolled, natural environment, thereby enhancing the external validity of our results. Given that such shifts in expectations can quickly transmit to the real economy, real-time monitoring becomes essential for policymakers, businesses, and investors.

---

<sup>25</sup>For completeness, we also report the OLS results in Table OA.2 in the Online Appendix, which use all days in the sample. The OLS estimates show a positive contemporaneous correlation between our Twitter-based measure of inflation expectations and discretionary consumption, but a negative correlation between current inflation expectations and consumption nine to twelve days ahead.

<sup>26</sup>These magnitudes are obtained by multiplying the reported coefficients in Table 6 by  $\ln(1.10)$ .

## 5 Inflation expectations and the stock market

Building on the previous section, which documented that households reduce spending on discretionary goods when inflation expectations rise, we now ask whether these shifts are reflected in firms' valuations, focusing on stock market performance.

Figure 8 plots the seven-day moving average of our Twitter-based inflation expectations index from January 2021 to June 2024, together with the average stock prices of 31 firms in the “Consumer Discretionary” industry and 179 firms in all other sectors. These prices are normalized to 100 on January 1, 2021.<sup>27</sup> Both stock price series exhibit pronounced negative co-movement with our index. However, as the index begins to rise toward the end of 2021 and thereafter, the stock prices of discretionary firms decline significantly more than those of non-discretionary firms. This suggests that discretionary firms may react more strongly to increases in households' inflation expectations. Moving beyond mere correlations, this section tests whether the market systematically discounts firms active in the discretionary goods market when inflation expectations rise.

The empirical evidence regarding the relationship between stock performance and inflation expectations is mixed. Chaudhary and Marrow (2024) show that an increase in market-based medium- and long-term inflation expectations is associated with higher stock returns in the US from 2000 to 2020. This is because higher long-term inflation expectations are linked to expectations of stronger future economic growth. In contrast, Knox and Timmer (2025) study the period from 1977 to 2022. They use inflation surprises around CPI announcements to show that investors view inflation as bad news for the economy: stock returns tend to fall in response to a surprise increase in inflation, even though such surprises raise short-term inflation expectations.

To study the relationship between inflation expectations and stock returns for Germany, we use daily returns on the German CDAX (Composite DAX) index, which includes all stocks listed on the Frankfurt Stock Exchange, and estimate the following regression:

$$r_t^{CDAX} = \alpha + \beta \Delta \mathbb{E}_t[\pi_t^s] + u_t, \quad (5)$$

where  $\mathbb{E}_t[\pi_t^s]$  denotes inflation expectations for different horizons and from different sources. Specifi-

<sup>27</sup>These are the daily stock prices of all DAX, MDAX, and SDAX constituents, as well as Deutsche Börse Prime Standard stocks. Firms are classified as belonging to the “Consumer Discretionary” sector according to the Global Industry Classification Standard (GICS), which is provided by LSEG.

cally, we use daily changes in our Twitter-based inflation expectations index and in three market-based measures: the Euro Area 1-year and 5-year inflation swap rates, and German 10-year breakeven inflation expectations.<sup>28</sup> We scale all series by their standard deviations, where the standard deviation is computed using daily changes in the respective series on CPI announcement days.

To move beyond reduced-form correlations, we focus on CPI release days. Since release dates are scheduled in advance, these days offer a quasi-experimental environment in which market participants receive a concentrated and salient inflation signal. While daily data are less precise than intraday event-study methods, German CPI releases are usually the only major scheduled domestic macro announcement on those days and commonly dominate media coverage. Accordingly, CPI days provide a natural window in which inflation-related information is the main systematic driver of belief updating.

For identification, we additionally assume that CPI releases are the primary source of systematic revisions in inflation expectations on those days. We recognize that CPI news may also shift other macro expectations (e.g., the policy-rate path). Still, to the extent that our Twitter-based index reflects households' concerns about real income erosion distinct from mechanical discount-rate channels, we interpret our estimates as capturing the equity-market response to inflation-related belief revisions, consistent with markets pricing in weaker real activity through a demand-expectations channel.

Table 8 shows that, for the full sample, an increase in *longer-run* inflation expectations is associated with an increase in stock returns on CPI announcement days, in line with Chaudhary and Marrow (2024). For the 1-year horizon and for our Twitter-based index, however, we do not find a significant relationship. When we focus on the more recent period characterized by the sharp rise in inflation in columns (5)–(8), the previously significant relationship between longer-run inflation expectations and stock returns disappears. Instead, an increase in our Twitter-based index leads to a significant decline in stock returns on CPI release days. This pattern is consistent with Knox and Timmer (2025), who document a negative relationship between increases in inflation perceptions and stock returns for the U.S. Since our Twitter-based index primarily captures short-term inflation expectations, which are generally shaped by inflation perceptions, the negative association we find

---

<sup>28</sup>Due to limited market liquidity, we do not use German breakeven inflation rates at the 5-year horizon or shorter maturities and instead rely on the more liquid Euro Area inflation swap rates.

on CPI days fits well with their study.<sup>29</sup> The effect is also economically meaningful: restricting the sample to CPI announcement days, a one-standard deviation increase in our index leads to a decline in the CDAX of 14 basis points.

This pattern is consistent with the evidence presented above: only during the recent period of elevated inflation do household inflation expectations measured from Twitter respond to major events such as monetary policy announcements or CPI releases. Moreover, we showed that discretionary consumption decreases following an increase in Twitter-based inflation expectations during this period, indicating that households interpret inflation developments through a supply-side lens, viewing higher inflation as bad news for future economic conditions. In line with this and with the findings of Knox and Timmer (2025) for the U.S., investors appear to associate higher inflation with an adverse economic outlook in the recent period. As a result, stock returns decline in response to an increase in Twitter-based inflation expectations, suggesting that, in periods of pronounced inflation, the Twitter-based index captures inflation-related news that is also relevant for financial markets and thus becomes a meaningful proxy for the signals investors pay attention to.

To illustrate this regime shift and the specific vulnerability of discretionary firms, Figure 9 plots the cumulative average returns around CPI release days for a major constituent of the DAX Consumer Discretionary sector—a leading European online fashion retailer. The figure includes only those CPI releases for which our index increases on the announcement day, capturing instances in which inflation expectations rise. It contrasts the low-inflation period (2011–2020) with the recent inflationary episode (2021–2024). Strikingly, in the earlier period, the firm’s stock price tended to increase following CPI releases, consistent with inflation signaling robust demand. In the recent period, however, the pattern reverses completely: the stock price drops sharply after CPI announcements. This anecdotal evidence suggests that the recent negative market reaction is particularly pronounced for firms relying on discretionary spending, motivating our formal test of this sector in the next step.

Motivated by this observation, we formally test whether the returns of firms in the discretionary consumption sector systematically react negatively to an increase in inflation expectations—measured with our Twitter-based index—during the period of elevated inflation. To this end, we match daily

---

<sup>29</sup>For evidence on how inflation perceptions translate into short-term inflation expectations, see, *e.g.*, D’Acunto et al. (2023).

stock returns of all DAX, MDAX, and SDAX constituents, as well as Deutsche Börse Prime Standard stocks, with quarterly firm fundamentals from Compustat, resulting in a final sample of 175 unique firms. We control for a broad set of firm characteristics, including size, profitability, and capital structure. Specifically, we measure profitability using return on assets (ROA) and return on equity (ROE), and capital structure using leverage and the short-term debt share.<sup>30</sup> Restricting the analysis again to CPI release days, we employ two approaches. First, we focus on episodes of above-average inflation by limiting the sample to periods in which inflation exceeded the sample mean of 2.21% and estimate the following specification:

$$ret_t^i = \beta_1 \Delta \mathbb{E}_t[\pi_t^{Twitter}] \times \mathbf{1}_{Discretion}^i + \beta_2 \Delta \mathbb{E}_t[\pi_t^{Twitter}] + \gamma \mathbf{X}_{i,t} + \mu_i + \lambda_t + u_{i,t}, \quad (6)$$

where  $\Delta \mathbb{E}_t[\pi_t^{Twitter}]$  denotes standardized daily changes in our Twitter-based inflation expectations index on CPI release days, and  $\mathbf{1}_{Discretion}$  is a dummy variable that equals 1 if firm  $i$  belongs to the “Consumer Discretionary” industry. The vector  $\mathbf{X}_{i,t}$  includes firm size, return on assets (RoA), return on equity (RoE), leverage, and the short-term-debt share in the previous quarter. We always include firm fixed effects,  $\mu_i$ . In our strictest specifications, we add date fixed effects,  $\lambda_t$ , to absorb all aggregate daily variation, which subsumes the direct effect  $\beta_2$ . Standard errors are two-way clustered at the firm and date level.

In a second specification, we use the entire sample period, again restricted to CPI release days only, and interact our index with the discretionary dummy and the level of CPI to directly test whether the effect of Twitter-based inflation expectations on discretionary firms’ stock returns depends on the prevailing inflation level:

$$\begin{aligned} ret_t^i = & \beta_1 \Delta \mathbb{E}_t[\pi_t^{Twitter}] \times \mathbf{1}_{Discretion}^i \times CPI_{m(t)} + \beta_2 \Delta \mathbb{E}_t[\pi_t^{Twitter}] \times \mathbf{1}_{Discretion}^i + \beta_3 CPI_{m(t)} + \\ & \beta_4 \mathbf{1}_{Discretion}^i \times CPI_{m(t)} + \beta_5 \Delta \mathbb{E}_t[\pi_t^{Twitter}] + \beta_6 \Delta \mathbb{E}_t[\pi_t^{Twitter}] \times CPI_{m(t)} + \\ & \gamma \mathbf{X}_{i,t} + \mu_i + \lambda_t + u_{i,t}, \end{aligned} \quad (7)$$

where  $CPI_{m(t)}$  is the German CPI in month  $m$  corresponding to day  $t$  and the other variables are defined as in equation (6).

---

<sup>30</sup>Size is measured as the logarithm of total assets. ROA is defined as quarterly income before extraordinary items scaled by average total assets, and ROE as income divided by average common equity. Leverage is the ratio of total debt to total assets and the short-term debt share is the ratio of short-term debt to total debt. All variables are winsorized at the 1st and 99th percentiles.

Columns (1) and (2) in Table 9 show that, indeed, stock returns of firms in the consumer discretionary sector decline when inflation expectations increase during periods of elevated inflation. Specifically, a one-standard-deviation increase in our Twitter-based index reduces daily returns of discretionary firms by about 11 basis points more than those of non-discretionary firms when inflation is above its average level of 2.21%.

The coefficients on the triple interaction in columns (3) and (4) confirm the additional role of the inflation level: discretionary firms' stock returns react particularly negatively to an increase in inflation expectations when inflation is elevated. More precisely, over the entire sample period, for firms in the discretionary sector, the effect of a one-standard-deviation increase in our Twitter-based inflation expectations index on daily returns becomes four basis points more negative for each additional percentage point of inflation, relative to non-discretionary firms.

Overall, these findings again demonstrate that our Twitter-based index captures inflation-related information that becomes especially relevant in periods of elevated inflation. During such times, increases in the index are associated with both reductions in discretionary consumption and lower realized returns for discretionary firms, providing a coherent picture in which higher expected inflation weakens consumer demand and, in turn, contributes to lower stock performance in the most exposed sectors.

## 6 Conclusion

Inflation expectations shape consumption and investment decisions and, by extension, aggregate outcomes. Conventional measures are costly, infrequent, and released with delay. We construct a high-frequency inflation expectations index from German-language tweets using modern NLP. We first denoise inflation-related tweets and classify them as signaling rising versus falling inflation, then aggregate these signals into a daily index. The index tracks realized CPI and, more closely, survey measures of household inflation expectations (and perceived inflation). Forecasting exercises show that it adds predictive power beyond standard benchmarks—especially for quantitative measures of household inflation expectations—suggesting it is a timely indicator of short-run beliefs.

We next examine whether the index responds to policy news in economically intuitive ways. In monetary policy announcement windows, unexpected tightening is followed—within roughly one week—by a decline in the index. This response is primarily driven by private individuals and is more pronounced during the recent high-inflation period, consistent with prior evidence that households are more attentive to inflation when inflation is elevated.

Beyond providing a new way of measuring inflation expectations, we also use our index to document real effects. Leveraging the exogenous timing of CPI release days as instruments in a 2SLS framework, we identify a causal impact of expectations on spending: when inflation expectations rise, households cut purchases within days in discretionary categories that are easier to postpone or substitute. We then link these demand-side adjustments to firm valuations. Using German stock prices and again exploiting CPI release days—when investors receive a sharp update of inflation-related information—we show that increases in inflation expectations, as captured by our Twitter-based index, lead to lower stock returns during the recent high-inflation period, with effects concentrated among firms in the consumer discretionary sector.

Overall, expectation dynamics seem to be reflected rapidly in both spending and equity prices. A high-frequency, real-time measure such as our Twitter-based index provides timely information for policymakers and market participants: it improves short-run forecasting of survey expectations, responds to monetary policy surprises within days, and helps trace how expectation updates propagate from consumption behavior to equity valuations.

## References

Acharya, V. V., M. Crosignani, T. Eisert, and C. Eufinger (2025). How Do Supply Shocks to Inflation Generalize? Evidence from the Pandemic Era in Europe. *Federal Reserve Bank of New York Staff Reports* 1164.

Adams, T., A. Ajello, D. Silva, and F. Vázquez-Grande (2023). More than Words: Twitter Chatter and Financial Market Sentiment. *Finance and Economics Discussion Series* 034.

Afrouzi, H., B. Candia, O. Coibion, S. Frache, D. Georgarakos, Y. Gorodnichenko, G. Kenny, S. Kumar, R. Lluberas, B. Meyer, J. Ponce, T. Ropele, and M. Weber (2025). Tell Me Something I Don't Already Know: Learning in Low-and High-Inflation Settings. *Econometrica* 93(1), 229–264.

Altavilla, C., L. Brugnolini, R. S. Gürkaynak, R. Motto, and G. Ragusa (2019). Measuring Euro Area Monetary Policy. *Journal of Monetary Economics* 108, 162–179.

Andreolli, M., N. Rickard, P. Surico, and C. Vergeat (2025). Discretionary Spending is the Cycle, and Why it Matters for Monetary Policy. *ECB Forum on Central Banking*.

Ang, A., M. Brière, and O. Signori (2012). Inflation and Individual Equities. *Financial Analysts Journal* 68(4).

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

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

Bartov, E., L. Faurel, and P. S. Mohanram (2018). Can Twitter Help Predict Firm-Level Earnings and Stock Returns? *The Accounting Review* 93(3), 25–57.

Baumann, U., M. Darracq Paries, T. Westermann, M. Riggi, E. Bobeica, A. Meyler, B. Böninghausen, F. Fritzer, R. Trezzi, J. Jonckheere, D. Kulikov, D. Popova, S. Pert, M. Paloviita, F. Brázdkik, H. Pönkä, M. Bess, P.-A. Robert, A. Al-Haschimi, P. Gmehling, M. Banbura, E. Charalampakis, J.-O. Menz, B. Hartwig, J. Paredes, U. Volz, L. Reiche, Z. Bragoudakis, M. Tirpak, E. Kasimati, T. Łyziak, A. Tagliabruni, E. Stanisławska, A. Bessonovs, N. I. Iskrev, O. Krasnopjorovs,

M. Gavura, T. Reichenbachas, M. Damjanović, R. Colavecchio, M. Maletic, G. Galati, D. Leiva, I. Kearney, and P. Stockhammar (2021). Inflation Expectations and their Role in Eurosystem Forecasting. *ECB Occasional Paper Series 264*.

Binder, C., P. Frank, and J. M. Ryngaert (2025). The Causal Effect of News on Inflation Expectations. *NBER Working Paper 34088*.

Binder, C. C., J. R. Campbell, and J. M. Ryngaert (2024). Consumer inflation expectations: Daily dynamics. *Journal of Monetary Economics 145*, 103613.

Brown, T. B., B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei (2020). Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems 33*, 1877–1901.

Cavallo, A. (2024). Inflation with Covid Consumption Baskets. *IMF Economic Review 72*(2), 902–917.

Cavallo, A., G. Cruces, and R. Perez-Truglia (2017). Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments. *American Economic Journal: Macroeconomics 9*(3), 1–35.

Chaudhary, M. and B. Marrow (2024). Inflation Expectations and Stock Returns. *mimeo*.

Christelis, D., D. Georgarakos, T. Jappelli, and M. van Rooij (2020, 03). Consumption Uncertainty and Precautionary Saving. *The Review of Economics and Statistics 102*(1), 148–161.

Coibion, O., F. D'Acunto, Y. Gorodnichenko, and M. Weber (2022). The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications. *Journal of Economic Perspectives 36*(3), 157–84.

Coibion, O., D. Georgarakos, Y. Gorodnichenko, and M. van Rooij (2023, July). How Does Consumption Respond to News about Inflation? Field Evidence from a Randomized Control Trial. *American Economic Journal: Macroeconomics 15*(3), 109–52.

Coibion, O., Y. Gorodnichenko, S. Kumar, and M. Pedemonte (2020). Inflation Expectations as a Policy Tool? *Journal of International Economics* 124, 103297.

Coibion, O., Y. Gorodnichenko, and M. Weber (2022). Monetary Policy Communications and Their Effects on Household Inflation Expectations. *Journal of Political Economy* 130(6), 1537–1584.

D'Acunto, F., U. Malmendier, and M. Weber (2023). Chapter 5 - What do the Data tell us about Inflation Expectations? In R. Bachmann, G. Topa, and W. van der Klaauw (Eds.), *Handbook of Economic Expectations*, pp. 133–161. Academic Press.

D'Acunto, F., E. Charalambakis, D. Georgarakos, G. Kenny, J. Meyer, and M. Weber (2024). Household Inflation Expectations: An Overview of Recent Insights for Monetary Policy. *NBER Working Paper 32488*.

D'Acunto, F., D. Hoang, and M. Weber (2021). Managing Households' Expectations with Unconventional Policies. *The Review of Financial Studies* 35(4), 1597–1642.

Ehrmann, M. and A. Wabitsch (2022). Central Bank Communication with Non-Experts – A Road to Nowhere? *Journal of Monetary Economics* 127, 69–85.

Flannery, M. J. and A. A. Protopapadakis (2002). Macroeconomic Factors Do Influence Aggregate Stock Returns. *Review of Financial Studies* 15(3), 751–782.

Grootendorst, M. (2022). BERTopic: Neural Topic Modeling with a Class-Based TF-IDF Procedure. *arXiv Preprint 2203.05794*.

Gu, C. and A. Kurov (2020). Informational Role of Social Media: Evidence from Twitter Sentiment. *Journal of Banking & Finance* 121, 105969.

He, P., J. Gao, and W. Chen (2021). DeBERTaV3: Improving DeBERTa using ELECTRA-Style Pre-Training with Gradient-Disentangled Embedding Sharing. *arXiv Preprint 2111.09543*.

Jordà, Ò. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review* 95(1), 161–182.

Knox, B. and Y. Timmer (2025). Stagflationary Stock Returns. *Finance and Economics Discussion Series 2025-056. Washington: Board of Governors of the Federal Reserve System*.

Kouvavas, O., R. Trezzi, M. Eiglsperger, B. Goldhammer, and E. Gonçalves (2020). Consumption Patterns and Inflation Measurement Issues during the COVID-19 Pandemic. *ECB Economic Bulletin* (7). Box in Economic Bulletin Issue 7, 2020.

Kroner, N. (2025). How Markets Process Macro News: The Importance of Investor Attention. *Finance and Economics Discussion Series 022*.

Laurer, M., W. van Atteveldt, A. S. Casas, and K. Welbers (2024). Less Annotating, More Classifying: Addressing the Data Scarcity Issue of Supervised Machine Learning with Deep Transfer Learning and BERT - NLI. *Political Analysis* 32(1), 84–100.

McInnes, L., J. Healy, and S. Astels (2017). hdbscan: Hierarchical Density Based Clustering. *Journal of Open Source Software* 2(11), 205.

McInnes, L., J. Healy, and J. Melville (2018). UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv preprint 1802.03426*.

Müller, H., T. Schmidt, J. Rieger, L. M. Hufnagel, and N. Hornig (2022). A German Inflation Narrative - How the Media frame Price Dynamics: Results from a RollingLDA Analysis. *DoCMA Working Paper 9*.

Schmidt, T., D. Boddin, S. Helmschrott, O. Goldfayn-Frank, M. Cato, and et. al. (2022). Bundesbank-Online-Panel-Households (BOP-HH). Version: 01. *Deutsche Bundesbank Dataset*.

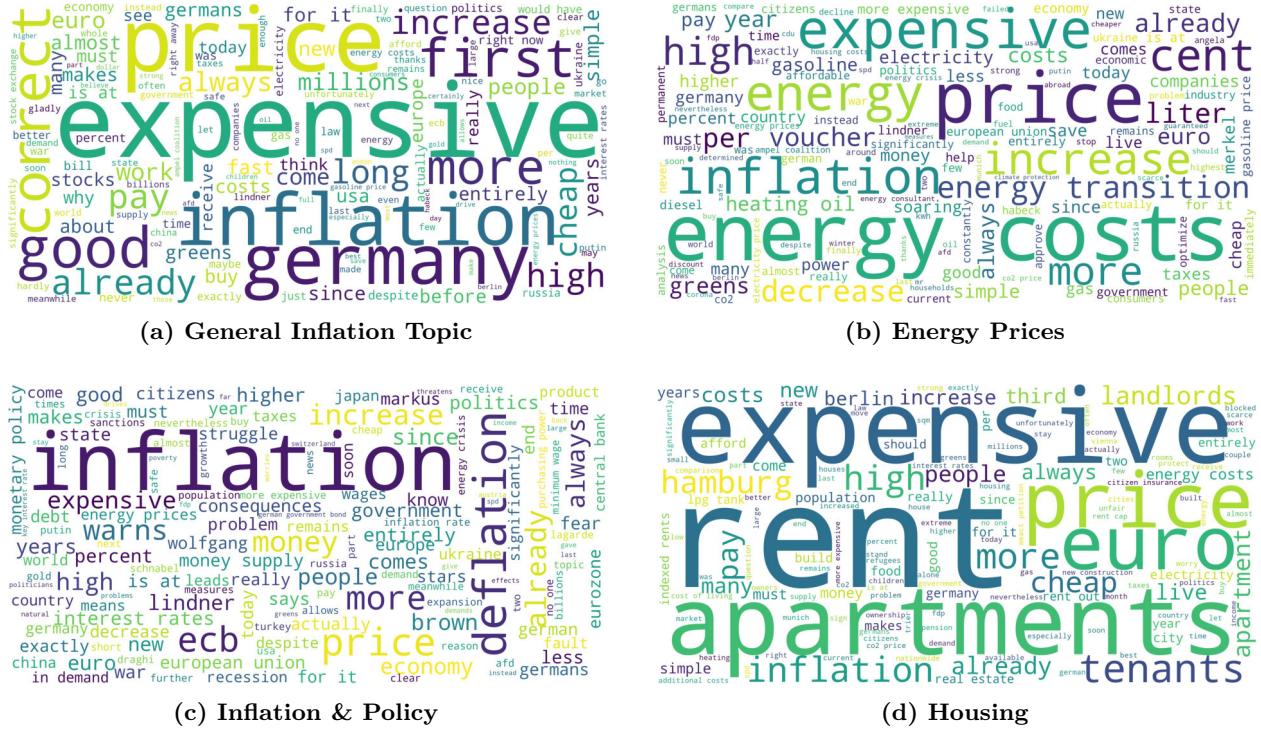
Sul, H. K., A. R. Dennis, and L. I. Yuan (2017). Trading on Twitter: Using Social Media Sentiment to Predict Stock Returns. *Decision Sciences* 48(3), 454–488.

Zhang, X., Y. Malkov, O. Florez, S. Park, B. McWilliams, J. Han, and A. El-Kishky (2023). TwHIN-BERT: A Socially-Enriched Pre-trained Language Model for Multilingual Tweet Representations at Twitter. *arXiv preprint 2209.07562*.

## A Figures

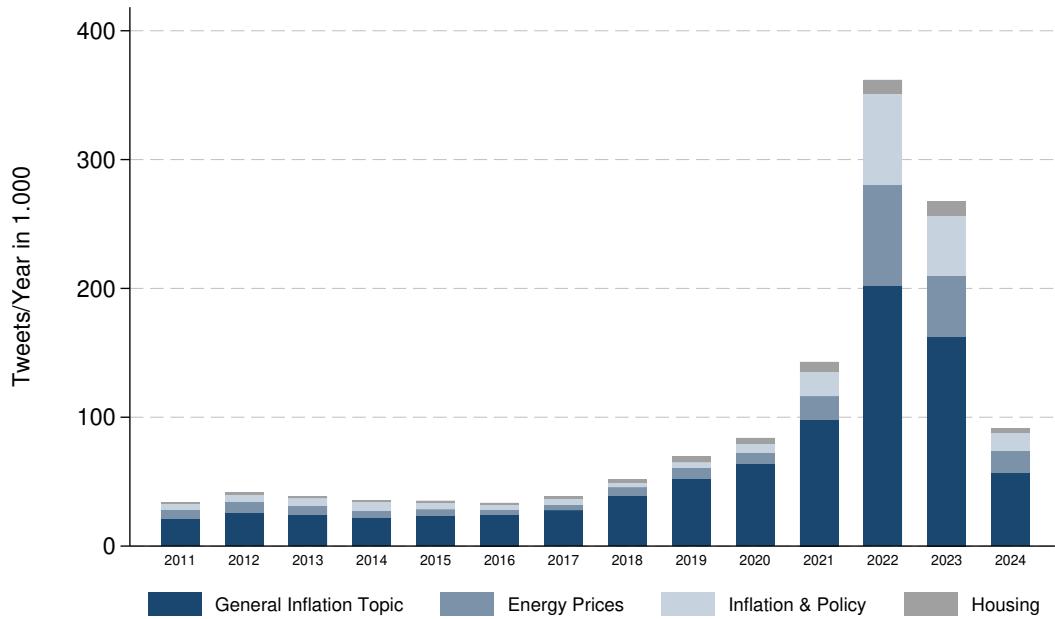
Figure 1

Grouped word clouds translated to English based on tweets within the 19 inflation topics



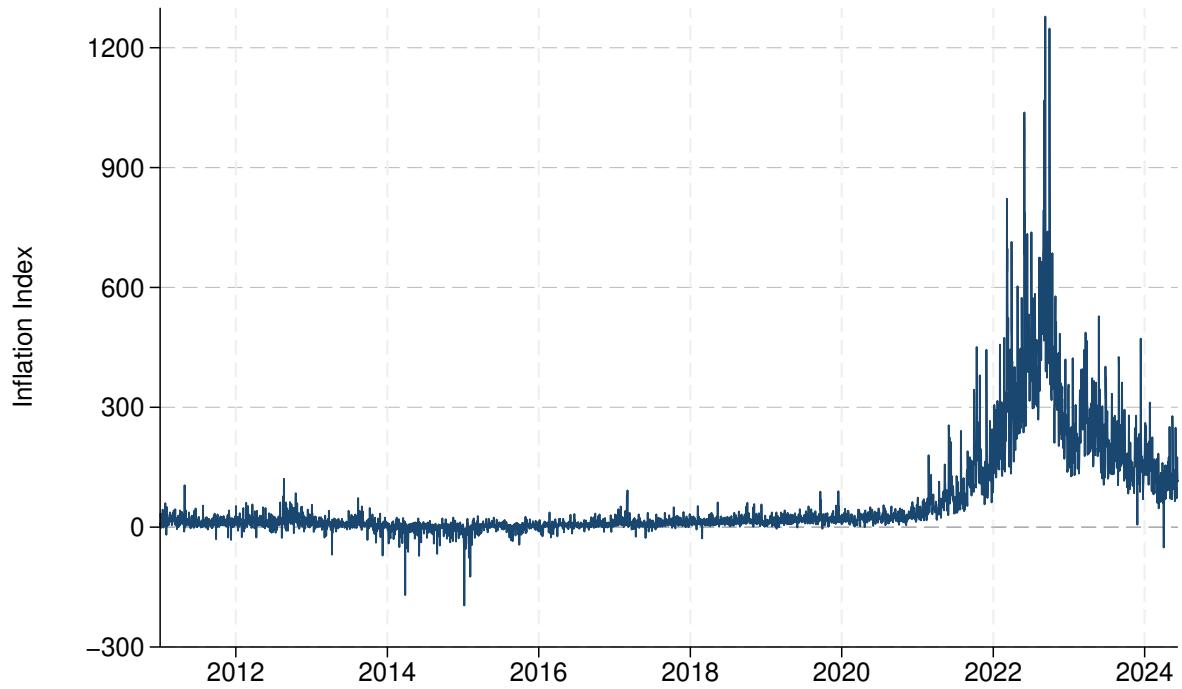
These four word clouds summarize tweets assigned by our topic model to 19 inflation-related topics, grouped into four clusters for visualization. Word size indicates the relative frequency of a term within each cluster, with larger words representing more important or frequently used terms. The figure is based on the original German tweets, with all words translated into English.

**Figure 2**  
**Number of tweets per year within the selected inflation topics over time**



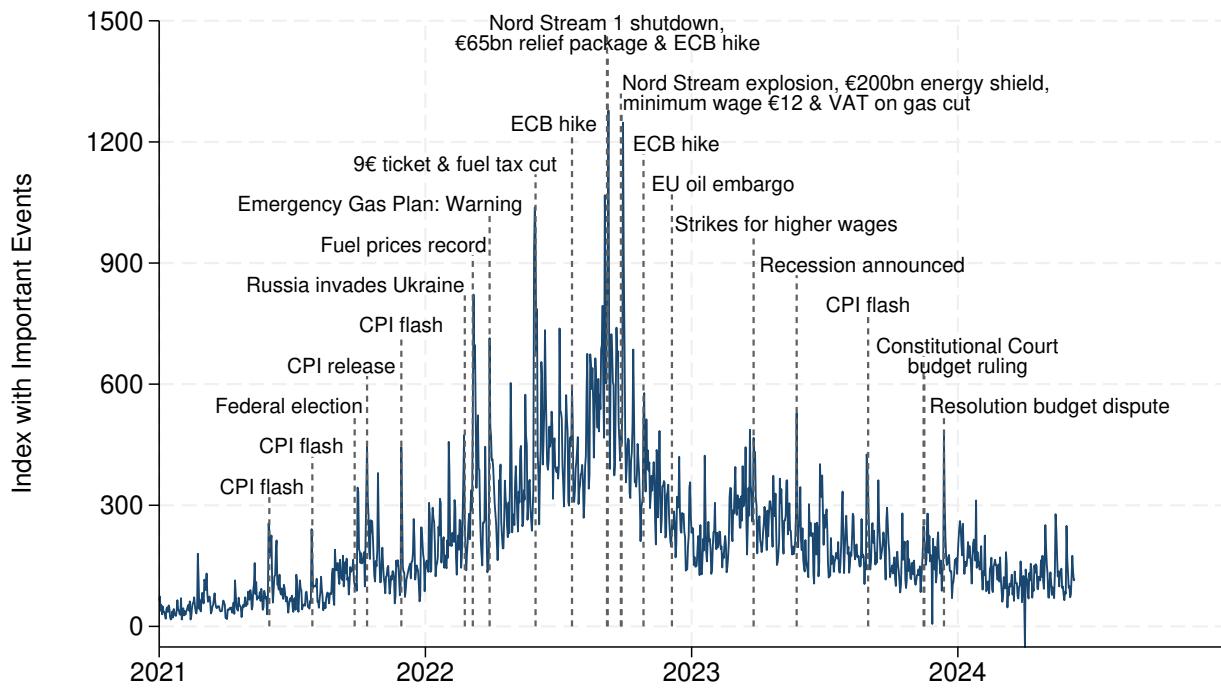
The stacked bars show the annual number of tweets (in thousands) in our inflation-related sample, grouped into four subtopics: *General*, *Energy Prices*, *Inflation & Policy*, and *Housing*. Tweets are filtered and assigned using the topic model described in the text, with embeddings based on the TwHIN-BERT model (Zhang et al., 2023).

**Figure 3**  
Twitter-based inflation expectations index between April 2007 and June 2024



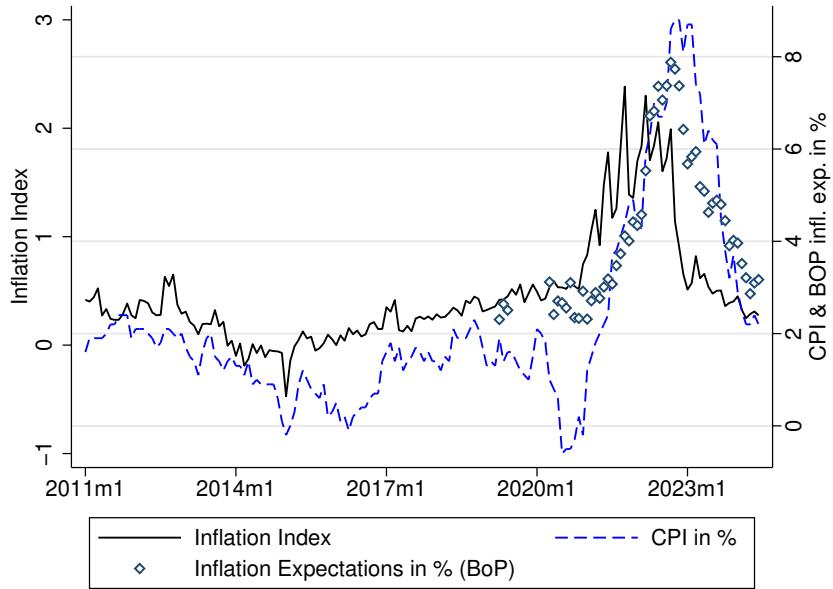
The figure plots our daily Twitter-based inflation expectations index for Germany from 28 April 2007 to 9 June 2024. Tweets are classified as "up" (rising prices/inflation), "down" (falling prices/inflation), or "other". The final index is the difference between the daily counts of up and down tweets ( $\text{Index}_t = \text{Up Index}_t - \text{Down Index}_t$ ).

**Figure 4**  
Twitter-based inflation expectations index with major inflation-relevant events (2021–2024)

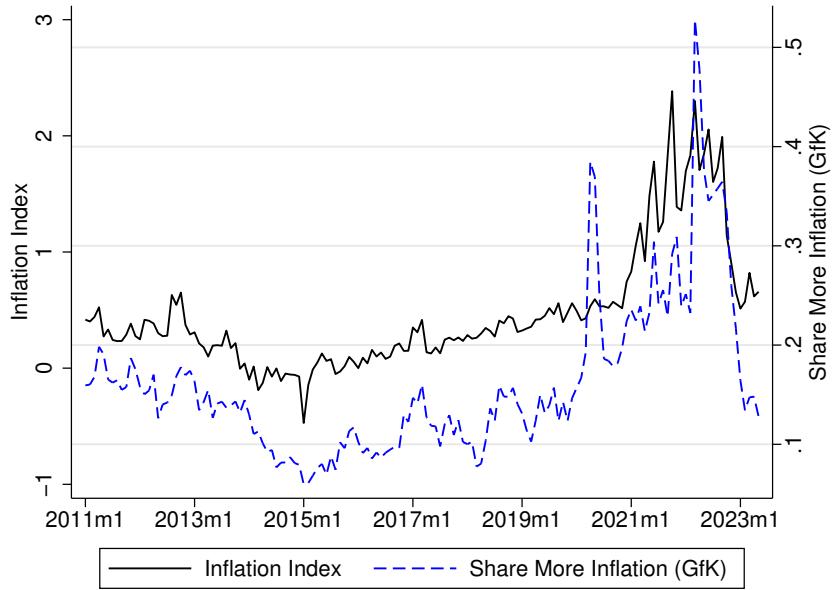


The blue line represents our daily Twitter-based inflation expectations index, zooming in on the 1 January 2021–9 June 2024 period to assess whether major inflation-relevant events are reflected in our index. Vertical dashed lines mark selected events that plausibly move inflation expectations (e.g., CPI releases, monetary policy actions, energy-price shocks). The displayed events are illustrative, not exhaustive.

**Figure 5**  
Twitter-based index, CPI, and survey expectations



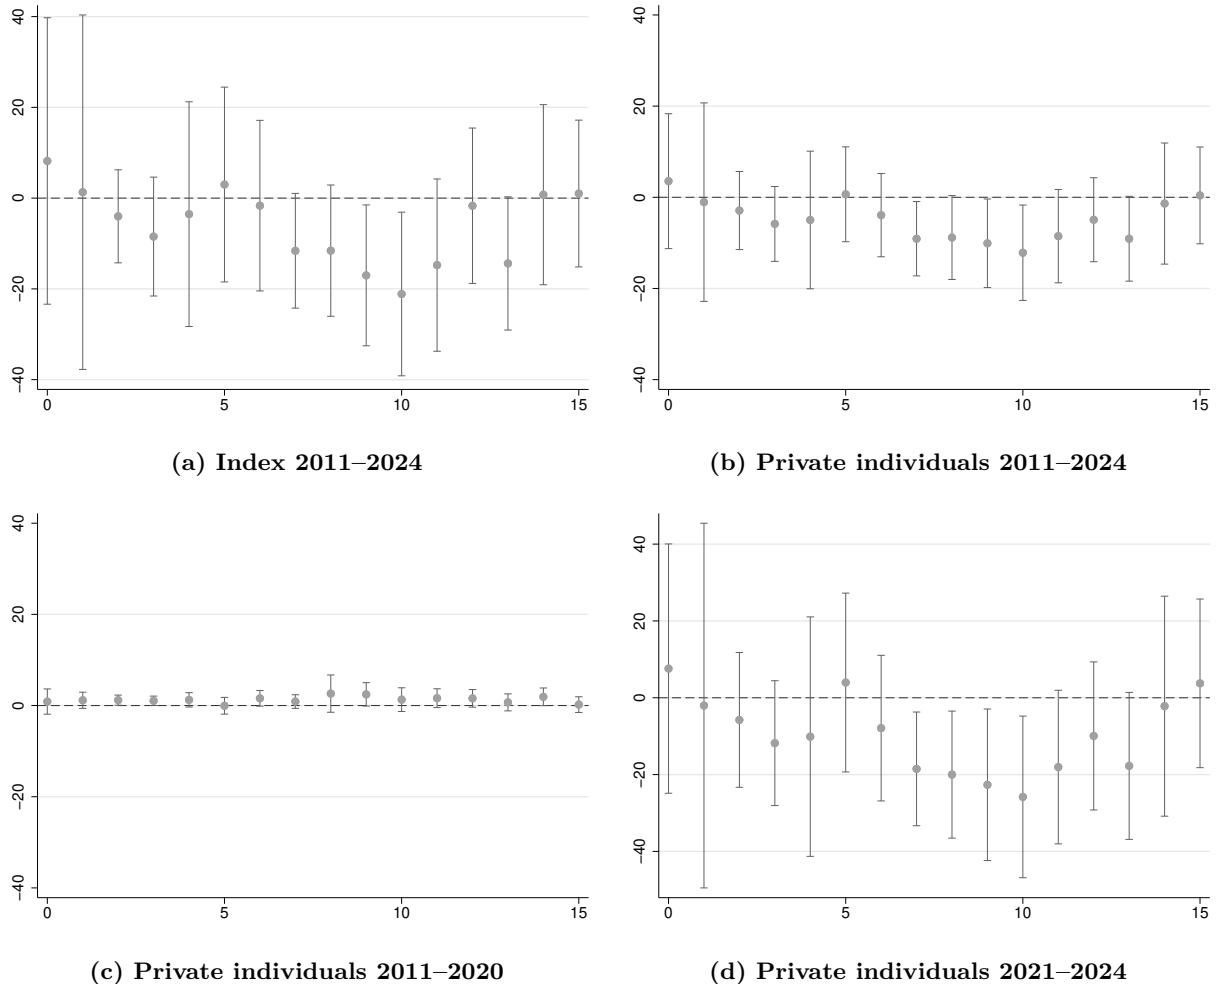
(a) Twitter-based index, CPI and BOP-HH inflation expectations



(b) Twitter-based index and GfK inflation expectations

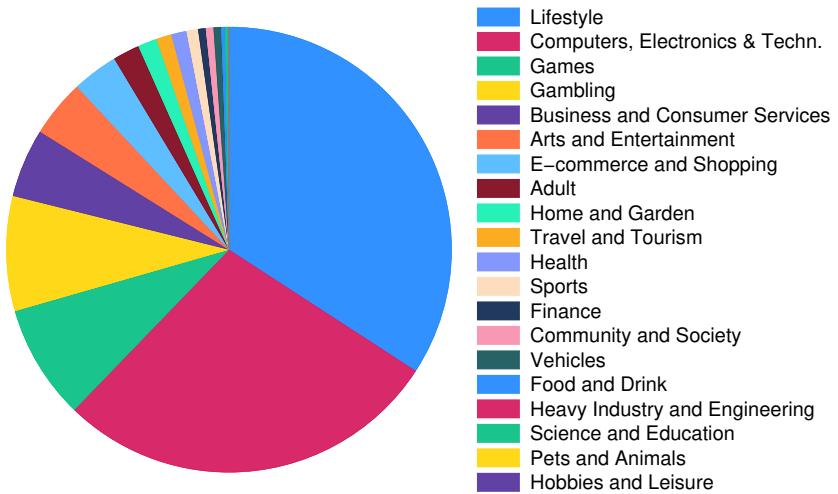
Panel (a) reports inflation expectations from the Bundesbank Online Panel – Households (BOP-HH), obtained from the Research Data and Service Centre (RDSC) of Deutsche Bundesbank, April 2019 to June 2024 (own calculations), together with the German CPI. Panel (b) reports inflation expectations from the GfK survey, measured as the share of respondents expecting prices to increase over the next 12 months, January 2011 to May 2023. In both panels, we also plot our Twitter-based inflation expectations index, constructed by standardizing the daily series (dividing by three times its 10-year rolling standard deviation) and then averaging to monthly frequency.

**Figure 6**  
**Sensitivity of private-individual inflation tweets to ECB surprises**



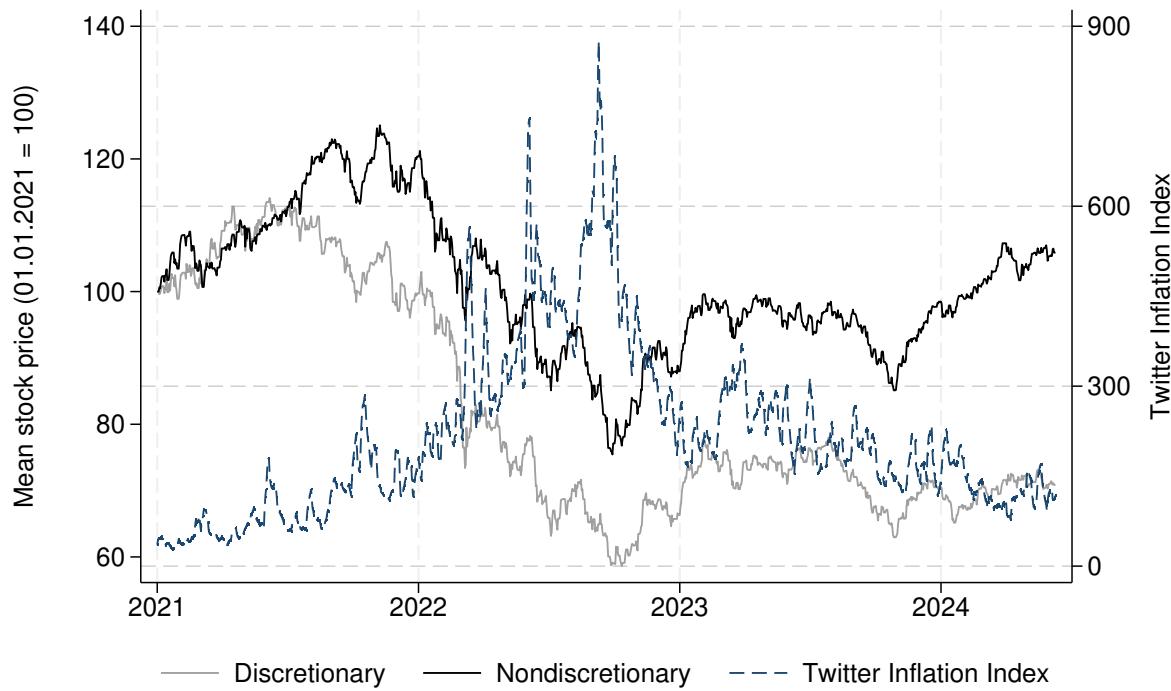
The figures plot the horizon-specific coefficients  $\beta^h$  from Equation (2), where we regress the cumulative change in our Twitter-based index between one day before the announcement and  $h$  days after on a 10-basis-point monetary policy surprise, measured as the intraday change in the 2-year OIS rate around the ECB press-release window (Altavilla et al., 2019). Panel (a) shows results for the general index; panels (b)-(d) show results for the private-individual index. We use Newey-West standard errors and show confidence bands at the 90% level.

**Figure 7**  
Shares of the number of transactions across different product categories



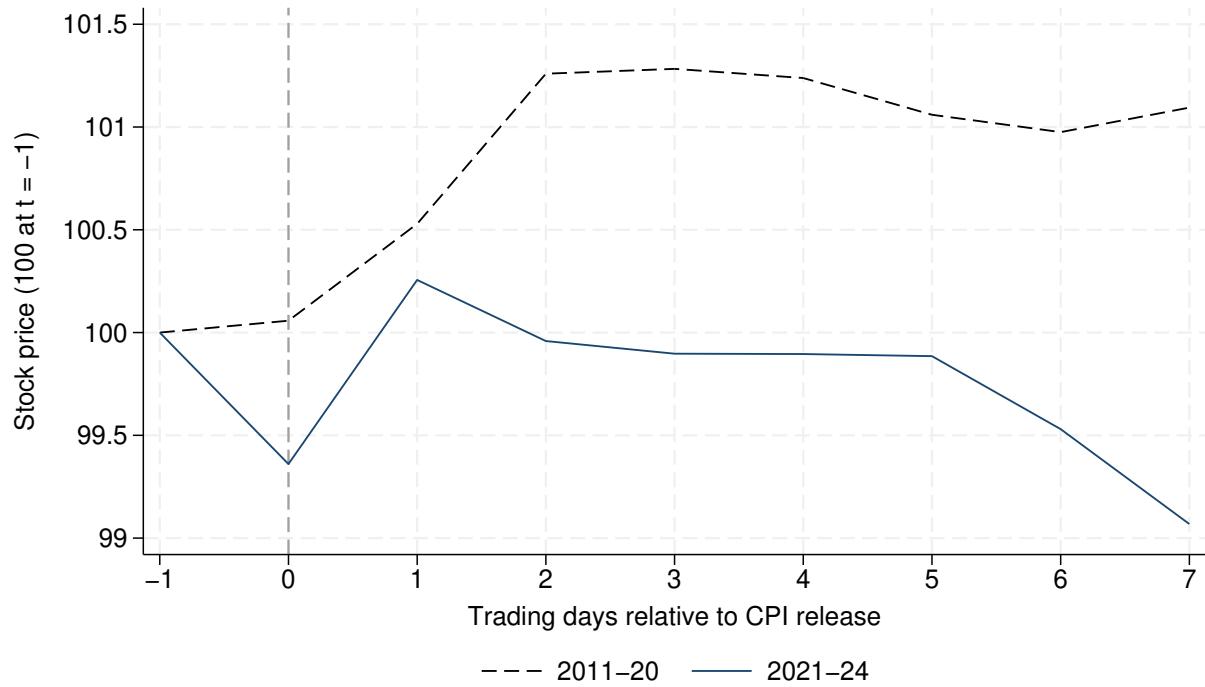
The figure shows the shares of daily online transactions of consumption goods by primary category for customers located in Germany. The sample consists of 212 firms that report transactions in at least 33 of the 39 months covered, yielding a total of 117,028,544 transactions over the period from 1 January 2021 to 31 March 2024.

**Figure 8**  
**Firms' stock prices and our Twitter-based inflation expectations index**



This figure plots the seven-day moving average of our inflation expectations index (blue dashed line, right y-axis) from January 2021 to June 2024, along with average stock prices of firms in the discretionary sector (light grey line) and all other firms (dark grey line). The industry classification “Consumer Discretionary” comes from LSEG. Both price series are normalized to 100 on January 1, 2021.

**Figure 9**  
**Stock price dynamics of a consumer discretionary firm around CPI announcements**



The figure plots the stock price of a leading European online fashion retailer (Consumer Discretionary) in the days surrounding CPI releases (day 0), conditional on inflation expectations increasing—that is, when the daily change in our index is positive on the release day. The gray line shows the average path during the low-inflation period (2011–2020), while the blue line shows the corresponding path during the high-inflation period (2021–June 2024).

## B Tables

**Table 1**  
**Summary statistics**

(a) Time series	<i>N</i>	Mean	Std. dev.	Min	Max
Index $E_t[\pi^{Twitter}]$	4,909	64.33	122.60	-196	1,278
Inflation CPI $_t$ in (%)	162	2.21	2.07	-0.60	8.80
BOP-HH Infl. exp. $E_t^{BOP}[\pi_{t,t+12}]$ (%)	54	4.25	1.66	2.31	7.88
BOP-HH Infl. perc. $E_t^{BOP}[\pi_{t-12,t}]$ (%)	47	5.16	2.08	2.21	8.43
GfK Infl. exp. $E_t^{GfK}[\pi_{t,t+12}]$	149	0.16	0.08	0.06	0.53
EA 1Y infl. swap rate $E_t[\pi_{Swap}^{EA,1Y}]$ (%)	3496	1.61	1.39	-0.85	8.63
EA 5Y infl. swap rate $E_t[\pi_{Swap}^{EA,5Y}]$ (%)	3499	1.51	0.64	0.09	3.44
German 10Y break. infl. $E_t[\pi_{BE}^{DEU,10Y}]$ (%)	3494	1.46	0.51	0.23	2.98
Mon. Pol. shocks $\epsilon_t$ (basis points)	290	-0.06	4.84	-22.80	18.70
Transactions $C_t^i$	207,661	563.56	2,143	1.00	53,902
Change in 10-year German Bund yield	1,154	0.00	0.05	-0.31	0.28
$\Delta$ Bund10Y $_t$ (%)					
Change in 1Y EA yield $\Delta$ Yields1YEAt (%)	1,154	0.00	0.04	-0.40	0.20
Dax returns $r_t^{DAX}$ (log changes)	1,154	0.03	0.90	-4.41	7.92
CDAX returns $r_t^{CDAX}$ (log changes)	3,412	0.03	1.20	-13.13	9.72
(b) Firm panel	<i>N</i>	Mean	Std. dev.	Min	Max
Firm returns $r_{i,t}$ (log changes)	443,994	0.01	2.33	-95.53	52.18
Total assets (log)	9,430	7.70	2.36	3.22	13.51
Return on assets (%)	9,100	0.87	2.27	-9.91	7.93
Return on equity (%)	7,662	2.29	6.47	-31.52	22.08
Leverage (%)	8,168	24.91	16.02	0.49	68.58
Short-term debt share (%)	8,166	26.99	21.66	0.18	96.72

This table presents summary statistics for our main variables. A description of all variables can be found in Table 2.

**Table 2**  
**Variable definitions**

Variable	Definition	Source
$1_{Discretion}$	Dummy equal to one for firms in the “Consumer Discretionary” sector according to the Global Industry Classification Standard (GICS).	LSEG
$\Delta \text{Bund}10Y_t$	Daily change in the 10-year German Bund yield.	OECD
$\Delta \text{Yields}1\text{YEA}_t$	Daily change in the one-year Euro Area yield-curve spot rate.	ECB Statistical Data Warehouse
$\epsilon_t$	ECB monetary policy surprise measured as the intraday change in the 2-year OIS rate around the ECB press-release window.	Altavilla et al. (2019)
$C_t^i$	Number of online transactions for firm $i$ on day $t$ .	Grips Intelligence
$\text{CPI}_t$	German monthly CPI inflation.	Destatis
$\text{CPI release}_t$	Indicator for German CPI release days.	Destatis, Investing.com
Discretionary consumption $_{i,t}$	Number of transactions for firms classified as primarily selling discretionary goods/services.	Grips Intelligence, own classification
$E_t[\pi_t^{Twitter}]$	Daily Twitter-based inflation expectation index constructed as the difference between daily counts of tweets classified as indicating rising vs. falling inflation/prices: $\text{Index}_t = \text{Up Index}_t - \text{Down Index}_t$ .	Twitter API
$E_t^{BOP}[\pi_{t-12,t}]$	BOP-HH inflation perceptions over the past 12 months, trimmed below $-12\%$ and above $12\%$ .	Deutsche Bundesbank (RDSC)
$E_t^{BOP}[\pi_{t,t+12}]$	Bundesbank Online Panel - Households (BOP-HH) inflation expectations over the next 12 months.	Deutsche Bundesbank (RDSC)
$E_t^{GfK}[\pi_{t,t+12}]$	GfK inflation expectations (share answering “Prices will increase more” for the next 12 months).	GfK
$E_t[\pi_{Breakeven}^{DEU}]$	German 5 or 10-year breakeven inflation.	Bloomberg
$E_t[\pi_{Swap}^{EA}]$	Euro Area 1 or 5-year inflation swap rate.	Bloomberg
$\text{Leverage}_{i,t}$	Total debt divided by total assets for firm $i$ in quarter $t$ .	Compustat
$r_{i,t}$	Daily stock return of firm $i$ .	LSEG
$r_t^{CDAX}$	Daily return on the German CDAX index.	Bloomberg
$r_t^{DAX}$	Daily return on the German DAX index.	LSEG
Return on assets $_{i,t}$	Quarterly income before extraordinary items divided by average total assets (average of total assets in quarters $t$ and $t-1$ ).	Compustat
Return on equity $_{i,t}$	Quarterly income before extraordinary items divided by average common equity (average of common equity in quarters $t$ and $t-1$ ).	Compustat
Short-term debt share $_{i,t}$	Short-term debt divided by total debt for firm $i$ in quarter $t$ .	Compustat
Total assets $_{i,t}$	Firm size measured as the natural logarithm of total assets for firm $i$ in quarter $t$ .	Compustat

The table reports the names of the variables used in the paper, their exact definitions, and the data sources.

**Table 3**  
**Correlations of CPI, inflation expectations and perceptions with our Twitter-based index**

Sample	CPI	$E_t^{GfK} \pi_{t,t+12}^*$	$E_t^{BOP} \pi_{t,t+12}$	$E_t^{BOP} \pi_{t-12,t}$
2011 - 2019	0.63	0.66		
2020 - 2024 (*2023)	0.83	0.48	0.94	0.68
2020 - 2022	0.88	0.64	0.95	0.85
2023 - 2024	0.80		0.88	0.84

Sample	$\mathbb{E}_t[\pi_{t,Swap}^{EA,1y}]$	$\mathbb{E}_t[\pi_{t,Swap}^{EA,5y}]$	$\mathbb{E}_t[\pi_{t,Breakeven}^{DEU,5y}]$	$\mathbb{E}_t[\pi_{t,Breakeven}^{DEU,10y}]$
2011 - 2019	0.63	0.43	0.54	0.25
2020 - 2024	0.91	0.84		0.76
2020 - 2022	0.92	0.89	0.93	0.86
2023 - 2024	0.78	0.69		0.71

The table reports correlations between the monthly averages of our Twitter-based inflation expectations index and actual CPI data, survey-based inflation expectations and perceptions, and market-based inflation expectations for different time periods. CPI data for Germany (January 2011 until June 2024) are from the Federal Statistical Office of Germany (Destatis). Inflation expectations  $E_t^{GfK} \pi_{t,t+12}$  are from the GfK survey, available from January 2011 until May 2023, and measured as the share of respondents expecting prices to increase over the next 12 months. Inflation expectations  $E_t^{BOP} \pi_{t,t+12}$  and perceptions  $E_t^{BOP} \pi_{t-12,t}$  come from the Bundesbank Online Panel – Households (BOP-HH), obtained from the Research Data and Service Centre (RDSC) of Deutsche Bundesbank, and are available from April 2019 to June 2024, own calculations.  $\mathbb{E}_t[\pi_{t,Swap}^{EA,1y}]$ ,  $\mathbb{E}_t[\pi_{t,Swap}^{EA,5y}]$ ,  $\mathbb{E}_t[\pi_{t,Breakeven}^{DEU,5y}]$ , and  $\mathbb{E}_t[\pi_{t,Breakeven}^{DEU,10y}]$  denote Euro Area 1-year and 5-year inflation swap rates and German 5-year and 10-year break-even inflation expectations, respectively, obtained from Bloomberg. German 5-year break-even inflation expectations are available only until August 2022, while all other market-based series extend to June 2024.

**Table 4**  
**Can the Twitter-based index explain inflation expectations and actual CPI?**

	$E_t^{BOP} \pi_{t,t+12}$	$E_t^{GfK} \pi_{t,t+12}$	$CPI_t$	$CPI_t$
$Index_t$	0.29*** (0.04)	0.35*** (0.07)	0.14*** (0.03)	0.12* (0.07)
$E_{t-1}^{BOP} \pi_{t-1,t+11}$	0.52*** (0.08)			-0.02 (0.14)
$E_{t-1}^{GfK} \pi_{t-1,t+11}$		0.77*** (0.05)		0.04 (0.05)
$CPI_{t-1}$	0.11** (0.05)	-0.22*** (0.07)	0.86*** (0.03)	0.89*** (0.09)
Adj. $R^2$	0.96	0.81	0.96	0.96
$N$	50	148	161	38

The dependent variable  $E_t^{BOP} \pi_{t,t+12}$  represents monthly inflation expectations from the Bundesbank Online Panel – Households (BOP-HH) for the next 12 months, obtained from the Research Data and Service Centre (RDSC) of Deutsche Bundesbank, own calculations. The dependent variable  $E_t^{GfK} \pi_{t,t+12}$  represents monthly inflation expectations from the GfK survey, calculated as the share of respondents expecting an increase in prices over the next 12 months. The dependent variable  $CPI_t$  is the monthly CPI for Germany from the Federal Statistical Office of Germany (Destatis). For our Twitter-based inflation expectation  $index_t$  we take the averages of the first 16 days of a given month. Due to data availability, the time period is April 2020 until June 2024 for the first column, January 2011 until May 2023 for the second column, January 2011 until June 2024 for the third column, and April 2020 until May 2023 for the fourth column. All variables have been standardized and standard errors are reported in parentheses. \*\*\*, \*\*, and \* denote significance at the 1, 5, and 10 percent levels, respectively.

**Table 5**  
**Forecasting inflation expectations and actual CPI**

	$h = 1$	$h = 2$	$h = 3$
AR(1): $E_t^{BOP} \pi_{t,t+12}$ , RMSE	0.544	0.923	1.328
AR(1): $E_t^{GfK} \pi_{t,t+12}$ , RMSE	0.040	0.056	0.064
AR(1): $\pi_{t-12,t}$ , RMSE	0.457	0.698	0.925
incl. index : $E_t^{BOP} \pi_{t,t+12}$ , <u>relative</u> RMSE	1.015	0.779	0.818
incl. index : $E_t^{GfK} \pi_{t,t+12}$ , <u>relative</u> RMSE	1.094	0.968	0.961
incl. index : $\pi_{t-12,t}$ , <u>relative</u> RMSE	0.984	0.909	0.904

The table reports the RMSE of the benchmark AR(1) models to forecast BOP-HH (first row, April 2020 until June 2024), GfK inflation expectations (second row, January 2011 until May 2023), and actual CPI data (third row, January 2011 until June 2024) and the ratios of the RMSE of the respective model including the monthly averages of our inflation expectations index relative to the benchmark model (rows four to six) for horizon  $h$  from 1 to 3 months ahead. For the latter, values below 1 indicate that the competing model performs better than the benchmark one. A recursive estimation scheme is applied with a first in-sample of 12 observations. Inflation expectations,  $E_t^{BOP} \pi_{t,t+12}$ , are from the Bundesbank Online Panel – Households (BOP-HH), obtained from Research Data and Service Centre (RDSC) of Deutsche Bundesbank, own calculations.

**Table 6**  
**Effect of inflation expectations on discretionary consumption: 2SLS panel IV estimates**

Days after CPI release	log(Index <sub>t</sub> )	First stage		Second stage			
		$\Delta_h \log(\text{Discretionary consumption}_{t+h}^i)$		$h=0$	$h=3$	$h=6$	$h=9$
		$h=0$	$h=3$				
CPI release <sub>t</sub>	0.21*** (0.04)						
log(Index <sub>t</sub> )		0.01 (0.04)	0.03 (0.07)	-0.11** (0.05)	-0.16** (0.07)	-0.07 (0.07)	
$r_t^{DAX}$	-0.01* (0.01)	-0.00 (0.00)	0.00 (0.00)	0.01 (0.00)	0.00 (0.00)	0.01 (0.00)	
$\Delta \text{Bund}10Y_t$	-0.23* (0.14)	-0.01 (0.05)	0.02 (0.07)	-0.01 (0.08)	-0.10 (0.11)	-0.03 (0.08)	
$\Delta \text{Yields}1YEA_t$	0.76*** (0.19)	0.12 (0.08)	0.17 (0.11)	0.04 (0.11)	0.14 (0.18)	0.06 (0.11)	
Weekend FE	Y	Y	Y	Y	Y	Y	
Firm FE	Y	Y	Y	Y	Y	Y	
Year-Month FE	Y	Y	Y	Y	Y	Y	
Adj. $R^2$	0.89						
$F$ statistic first stage		21.88	21.77	21.90	21.94	21.95	
$N$	156,024	155,697	154,973	154,308	153,663	153,041	

The dependent variable in column (1) is the log of the daily inflation expectations index, and those in columns (2)–(6) are the log changes in daily transactions between  $t - 1$  and  $t + h$ , that is,  $\log(C_{t+h}^i) - \log(C_{t-1}^i)$ , where  $t$  denotes the day of the CPI release. Controls include one to seven lags of log transactions and of the log index (both not shown here), the daily DAX returns, daily changes in the 10-year Bund, and the one-year Euro Area yield curve spot rate, as well as weekend, month-of-year and firm fixed effects. The sample period runs from 1 January 2021 to 31 March 2024. Standard errors are two-way clustered by firm and date and reported in parentheses. \*\*\*, \*\*, and \* denote significance at the 1, 5, and 10 percent levels, respectively.

**Table 7**  
**Effect of inflation expectations on discretionary consumption: 2SLS panel IV estimates by individual consumption category**

	Lifestyle	E-commerce and Shopping	Travel and Tourism	Computers, Electronics & Techn.	Gambling
<i>h</i> =6	-0.136** (0.069)	-0.082 (0.081)	-0.103 (0.069)	-0.261** (0.133)	0.158 (0.150)
<i>h</i> =9	-0.184** (0.083)	-0.155* (0.088)	-0.095 (0.089)	-0.134 (0.153)	-0.104 (0.157)
<i>F</i> first stage	21.71	24.35	22.60	23.16	17.65
<i>N</i>	67,579	12,595	13,955	4,498	8,091
	Home and Garden	Sports	Health	Games	Business and Consumer Services
<i>h</i> =6	-0.122 (0.084)	-0.201* (0.110)	-0.154* (0.088)	-0.101 (0.102)	0.065 (0.185)
<i>h</i> =9	-0.059 (0.094)	-0.122 (0.110)	-0.315*** (0.122)	-0.345** (0.136)	-0.113 (0.183)
<i>F</i> first stage	24.86	22.06	19.98	19.07	11.02
<i>N</i>	10,953	8,284	5,040	7,750	1,549
	Vehicles	Adult	Arts and Entertainment	Science and Education	Hobbies and Leisure
<i>h</i> =6	-0.122 (0.177)	-0.040 (0.102)	-0.082 (0.097)	0.252 (0.214)	-0.291 (0.250)
<i>h</i> =9	0.067 (0.162)	-0.099 (0.099)	-0.087 (0.109)	0.072 (0.230)	0.112 (0.317)
<i>F</i> first stage	22.27	26.16	22.34	20.77	18.62
<i>N</i>	2,167	4,795	3,268	2,108	1,082

The dependent variables are the log changes in daily transactions between  $t - 1$  and  $t + h$  (for  $h = 6, 9$ ), defined as  $\log(C_{t+h}^i) - \log(C_{t-1}^i)$ , where  $t$  denotes the day of the CPI release. Controls include one to seven lags of log transactions and the log index (not shown), daily DAX returns, daily changes in the 10-year Bund yield, and the one-year Euro Area yield curve spot rate, as well as weekend, month-of-year, and firm fixed effects. The sample period runs from 1 January 2021 to 31 March 2024. Standard errors are clustered by date only, as the number of firms in some categories is small, and are reported in parentheses. \*\*\*, \*\*, and \* denote significance at the 1, 5, and 10 percent levels, respectively.

**Table 8**  
**Effect of changes in inflation expectations on stock returns on CPI release days**

	$r_t^{CDAX}$							
	2011–2024				2021–2024			
$\Delta\mathbb{E}_t[\pi_t^{Twitter}]$	-0.104 (0.06)				-0.140** (0.06)			
$\Delta\mathbb{E}_t[\pi_{t,Swap}^{EA,1y}]$		0.081 (0.07)				-0.016 (0.08)		
$\Delta\mathbb{E}_t[\pi_{t,Swap}^{EA,5y}]$			0.282*** (0.08)				0.051 (0.10)	
$\Delta\mathbb{E}_t[\pi_{t,Breakeven}^{DEU,10y}]$				0.316*** (0.08)				0.084 (0.11)
<i>Adj.R</i> <sup>2</sup>	0.01	0.00	0.04	0.05	0.05	-0.01	-0.01	-0.01
<i>N</i>	301	300	300	299	83	83	83	82

The dependent variable consists of daily stock returns (in percentage points) of the German CDAX index.  $\Delta\mathbb{E}_t[\pi_t^{Twitter}]$  denotes daily changes in our Twitter-based inflation expectations index; for Mondays, we use the average of Monday, Sunday, and Saturday.  $\Delta\mathbb{E}_t[\pi_{t,Swap}^{EA,1y}]$ ,  $\Delta\mathbb{E}_t[\pi_{t,Swap}^{EA,5y}]$ , and  $\Delta\mathbb{E}_t[\pi_{t,Breakeven}^{DEU,10y}]$  denote daily changes in the Euro Area 1-year and 5-year inflation swap rates and in German 10-year breakeven inflation expectations, respectively. All variables are divided by their standard deviation on CPI release days. The sample period spans 3 January 2011 to 7 June 2024 in columns (1)–(4), and 4 January 2021 to 7 June 2024 in columns (5)–(8). Standard errors are reported in parentheses, and \*\*\*, \*\*, and \* denote significance at the 1, 5, and 10 percent levels, respectively.

Table 9

Effect of changes in inflation expectations on discretionary stock returns on CPI release days during periods of elevated inflation

	$r_{i,t}$			
	Inflation > sample average		CPI-Interaction	
$\Delta\mathbb{E}_t[\pi_t^{Twitter}]$	-0.156** (0.06)		0.114 (0.15)	
$\Delta\mathbb{E}_t[\pi_t^{Twitter}] \times \mathbf{1}_{Discretion}$	-0.115*** (0.04)	-0.112** (0.04)	0.162** (0.07)	0.147* (0.08)
$\Delta\mathbb{E}_t[\pi_t^{Twitter}] \times \mathbf{1}_{Discretion} \times CPI_{m(t)}$			-0.042*** (0.01)	-0.041*** (0.01)
$\Delta\mathbb{E}_t[\pi_t^{Twitter}] \times CPI_{m(t)}$			-0.042** (0.02)	
$CPI_{m(t)}$			0.041 (0.03)	
$\mathbf{1}_{Discretion} \times CPI_{m(t)}$			0.032** (0.02)	0.033* (0.02)
Firm controls	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y
Time FE	N	Y	N	Y
<i>Adj.R</i> <sup>2</sup>	0.02	0.19	0.01	0.17
<i>N</i>	10,915	10,915	39,386	39,386

The dependent variable consists of stock returns of firm  $i$  on day  $t$ .  $\Delta\mathbb{E}_t[\pi_t^{Twitter}]$  denotes daily changes in our Twitter-based inflation expectations index, divided by its standard deviation on CPI release days; for Mondays, we use the average of Monday, Sunday, and Saturday.  $\mathbf{1}_{Discretion}$  is a dummy variable that takes the value 1 if the firm's industry is "Consumer Discretionary".  $CPI_{m(t)}$  is the German CPI in month  $m$  of day  $t$ . Firm controls include a firm's size, RoA, RoE, leverage, and short-term-debt-share in the previous quarter. Columns (1) and (2) include only days on which CPI inflation was above its sample average of 2.21%. Columns (3) and (4) include all days. The sample period spans 3 January 2011 to 7 June 2024. Standard errors are two-way clustered by firm and date and reported in parentheses. \*\*\*, \*\*, and \* denote significance at the 1, 5, and 10 percent levels, respectively.

# Online Appendix

## OA.1 Details on the index creation

### OA.1.1 Downloaded Tweets—Details

We use the following German keywords to select the tweets that we download for our initial data set: preis OR Lebenshaltungskosten OR (hohe rechnung) OR inflation OR teure OR teuer OR benzinpreis OR (hohe miete) OR (niedrige miete) OR energiekosten OR deflation OR disinflation OR schlussverkauf OR abverkauf OR (niedrige rechnung) OR (niedrige kosten) OR billig. These correspond roughly to the English keywords: price, cost of living, high bill, inflation, expensive, gasoline price, high rent, low rent, energy costs, deflation, disinflation, clearance sale, sell-off, low bill, low costs, and cheap.

### OA.1.2 Details on the topic modeling approach to filter out inflation tweets

To filter out the relevant tweets that are about inflation, we use the topic modeling technique BERTopic (Grootendorst, 2022) that is based on pre-trained transformer-based language models.<sup>31</sup> Besides creating more coherent and better interpretable topics than other topic models such as, e.g., Latent Dirichlet Allocation (LDA), BERTopic has the advantages that it can handle large volumes of text data and is computationally efficient and flexibly adjustable to a specific setting.<sup>32</sup> The pipeline within BERTopic consists of three main steps<sup>33</sup>: First, the text has to be converted to a vector representation—the embedding—for which pre-trained transformer-based models are used. One of the most commonly used models and the basis for many other state-of-the-art models is the language model BERT, which was developed by Google, uses artificial neural networks and has been trained on very large data sets, including more than 50 languages.<sup>34</sup> By default, BERTopic uses the SBERT model, which is a modification of the BERT architecture that is specifically designed to generate high-quality sentence embeddings. However, given our specific application context, we utilize the TwHIN-BERT model introduced by Zhang et al. (2023) for the embedding generation.<sup>35</sup> As this model is specifically trained on seven billion tweets covering over 100 distinct languages, it should be well suited to analyze our German-language inflation tweets.

In the second step, a dimensionality reduction is performed, and the reduced embeddings are

---

<sup>31</sup>Transformer-based language models are neural network architectures that use attention mechanisms to selectively process different parts of input text and have achieved state-of-the-art performance on a wide range of NLP tasks.

<sup>32</sup>For example, in contrast to more traditional models like LDA, BERTopic uses graphics processing units (GPU) instead of central processing units (CPU), making it much more time efficient, which is a relevant advantage given our relatively large data set.

<sup>33</sup>For a more detailed description of the single steps involved in BERTopic, see Grootendorst (2022).

<sup>34</sup>BERT stands for Bidirectional Encoder Representations from Transformers.

<sup>35</sup>TwHIN-BERT is based on the BERT architecture and part of the Hugging Face’s Transformers library that provides a range of pre-trained transformer-based models for various natural language processing tasks. Hugging Face is a company and open-source community that develops and maintains NLP tools and frameworks (see <https://huggingface.co/>).

clustered.<sup>36</sup> Subsequently, the documents within each cluster are aggregated, and a bag-of-words representation is generated using a count vectorizer model.

Finally, to assign a topic to each cluster, we extract topic representations using the class-based term frequency-inverse document frequency (c-TF-IDF) algorithm implemented in BERTopic. The c-TF-IDF method weights each term according to how frequent it is within a given cluster relative to its frequency across all clusters, thereby enhancing interpretability. Specifically, the weight  $W_{t,c}$  for term  $t$  in cluster  $c$  is calculated as

$$W_{t,c} = tf_{t,c} \cdot \log \left( 1 + \frac{A}{f_t} \right),$$

where  $tf_{t,c}$  denotes the frequency of term  $t$  within cluster  $c$ ,  $f_t$  represents the total frequency of term  $t$  across all clusters, and  $A$  is the average number of words per cluster. This approach, introduced in BERTopic, computes term importance at the topic level rather than across individual documents, thereby facilitating the human interpretation of the resulting topics.

In general, the brevity of tweets could pose a challenge for topic models. Therefore, to train our model, we create documents consisting of five tweets from the same user in chronological order, which leads to around 2 million documents. We clean the documents to remove web addresses, user mentions, hashtag symbols, and extra white spaces. Although the model can, in principle, handle text without this cleaning process, we still perform this step as it simplifies the interpretation of the resulting topics. For the embeddings, we use ngrams between 1 and 2 words. To decide on the number of topics, we explored different numbers and finally set the number of topics to 150, as these seemed the best to interpret and to be able to separate different topics from each other. Such a large number of topics further helps disentangle bot activity from the tweets that we are interested in.<sup>37</sup> We manually go through the most important tokens for these topics and select 19 topics that deal with inflation.

### OA.1.3 Details on tweet classification

To classify our inflation-related tweets according to whether they indicate increasing, decreasing, or stable inflation, we use the pre-trained neural network language model TwHIN-BERT (Zhang et al., 2023), which we fine-tune for our specific task by training additional layers on top of the pre-trained model. For this fine-tuning step, we require a training dataset consisting of tweets and their corresponding labels in one of three categories: *increasing prices or increasing inflation*, *decreasing prices or decreasing inflation*, and *other*.<sup>38</sup>

<sup>36</sup>For the dimensionality reduction, a uniform manifold approximation and projection (UMAP) technique is applied (see McInnes et al., 2018). The embeddings are clustered using HDBSCAN, a hierarchical density-based clustering approach developed by McInnes et al. (2017).

<sup>37</sup>For example, there is a bot tweeting about cheap gasoline prices that is now assigned to its own topic and not to a broader topic about energy and gasoline prices.

<sup>38</sup>In German, these are *steigende Preise oder steigende Inflation*, *sinkende Preise oder sinkende Inflation*, and *andere*. We do not distinguish between changes in prices and inflation, as these concepts are often used interchangeably in public discussions.

Labeling such a dataset manually would be highly time-consuming and could introduce idiosyncratic biases across annotators. We therefore employ ChatGPT, an autoregressive language model developed by OpenAI based on the GPT (Generative Pretrained Transformer) architecture and trained on massive text corpora using unsupervised learning techniques. Specifically, we use the *gpt-3.5-turbo-0301* model via the OpenAI API to label a selected set of tweets until we obtain 6,500 examples per category.<sup>39</sup> The exact prompt is the following:

*Classify this German tweet into one of the three categories: 'steigende Preise oder steigende Inflation oder hohe Inflation', 'sinkende Preise oder sinkende Inflation oder niedrige Inflation' or 'andere'. Confidence of prediction (COP): low, medium, high. The tweet is: '[tweet]' Output expected in the form of - Label: xxx, Explanation: xxx, COP: xxx.*

Consider, for example, the following tweet:

*#inflation steigt weiter in die Höhe und dann verkündet #Habeck diese Woche auch noch Alarmstufe Gas. Wir müssen das bezahlen was die Politiker vergeigt haben. Sie mussten unbedingt Sanktionen gegen Russland verhängen. Quelle des Screenshots: [https:...](https://...)*

This tweet leads to the following response: Label: *steigende preise oder steigende inflation*, Explanation: *The tweet mentions inflation rising and the need to pay for mistakes made by politicians*, COP: *medium*.

We set the system role to *You are an economist* and transform the responses in a parsable manner. We add the requirement of giving a reason for the choice of the class for a specific tweet to better judge the ability of ChatGPT to classify the inflation tweets. We manually browse randomly through a subset of the labelled tweets and can confirm that ChatGPT is able to understand the main messages of the tweets—at least comparable to a human, as sometimes even for humans, it is not easy to decide on the correct category of a tweet.

We aim at generating a balanced training data set with equal amounts of tweets for each class. However, since in reality the three classes are heavily unbalanced with the low inflation class being much smaller than the others, this is not straightforward to achieve.

Therefore, we first perform a zero-shot classification method, for which we do not have to provide any annotated training data, but only the three classification labels. Specifically, we use the machine learning algorithm "mDeBERTa-v3-base-mnli-xnli" by Laurer et al. (2024), which aims at understanding many different languages and performing any kind of classification task. This algorithm is a fine-tuned version of the "DeBERTaV3-base" transformer by He et al. (2021) from Microsoft, which is an improved version of the original BERT language model. The idea of this algorithm is that it is pre-trained so well that it does not need any or only very few additional annotated data to learn how to perform a given task.<sup>40</sup> Since the model was fine-tuned using

---

<sup>39</sup>For further details on the GPT-3 model, see Brown et al. (2020).

<sup>40</sup>The original BERT model was trained on 16 gigabytes of books and Wikipedia texts, added by 145 gigabytes data on news articles, links on Reddit, and story-like texts for the "DeBERTaV3-base" transformer by Microsoft. In addition, for the "mDeBERTa-v3-base-mnli-xnli" model, the previous model's pre-training is further fine-tuned by using more than a million classification examples from different Natural Language Inference (NLI) data sets—specifically, the English "MNLI" data set and the "XNLI" data set, containing 15 different languages, including German. For further details, see Laurer et al. (2024).

different languages, we can directly apply it to our German-language tweets. To evaluate the quality of this approach, we calculate the accuracy score based on the manually labeled set of tweets, which is 0.46—hence, significantly lower than the one for our final fine-tuned model.

In addition to the respective label for each tweet, the algorithm also returns a score which indicates how certain the decision for a specific task is, *i.e.* how large the prediction probability is. We use this score to sort the tweets within each category in descending order based on the model’s certainty that a specific tweets belongs to a certain class. Besides having equally distributed classes in the training data set, it is also important to include both easy and hard examples, where the former can be expected to have a high zero-shot prediction probability and the latter a low one. Therefore, after first choosing randomly 2,000 tweets from each class, we add another 2,000 tweets per class from both easy and hard examples in a 1:3 ratio, resulting in 12,000 tweets. We feed more difficult classification problems to ChatGPT to not waste API calls on those tweets that are very easy to classify, anyway. Since this approach still returns much more tweets belonging to the increasing or other class, we add more tweets to the ChatGPT prompt in a different ratio. We now add 8,000 tweets from the increasing and other class based on the zero-shot algorithm’s classification and 16,000 from the decreasing class. Afterwards, we have 6,500 tweets belonging to the decreasing class and take as many tweets from the the other two classes, as well, to have a balanced final training data set.

We use 90% of this labeled dataset to fine-tune the pre-trained TwHIN-BERT model by Zhang et al. (2023), which we also employ in the topic-modeling stage. For classification, we use the original (uncleaned) tweets, as the model can process special characters, emojis, and other non-standard tokens that may carry relevant contextual information.

During training, the model iteratively learns from the labeled examples and adjusts its parameters to minimize the loss function—the difference between predicted and true labels.<sup>41</sup> After every 250 training steps, we compute the accuracy score, defined as the share of correctly classified tweets in the validation set. If the accuracy improves relative to the previously saved model, the current model is stored. We repeat this process for three full passes through the dataset (three epochs) and select the best-performing model.

The final model achieves an accuracy of 0.76 on the validation set, reached during the second epoch.<sup>42</sup> When evaluated against an independent set of around 2,000 manually labeled tweets, the model attains an accuracy of 0.64.<sup>43</sup>

Applying this final model to our entire tweet corpus yields 405,184 tweets that discuss increasing inflation or prices, 86,666 tweets referring to decreasing inflation or prices, and 865,759 tweets classified as *other*.

---

<sup>41</sup>We use the cross-entropy loss function, defined as  $L_{CE} = -\sum_{i=1}^n t_i \log(p_i)$ , where  $t_i$  is the true label  $\in [0, 1]$  and  $p_i$  is the predicted probability distribution.

<sup>42</sup>We also track precision, which evolves similarly to the accuracy score and reaches values in a comparable range.

<sup>43</sup>Manual classification is itself imperfect, as even human annotators sometimes disagree on the appropriate label.

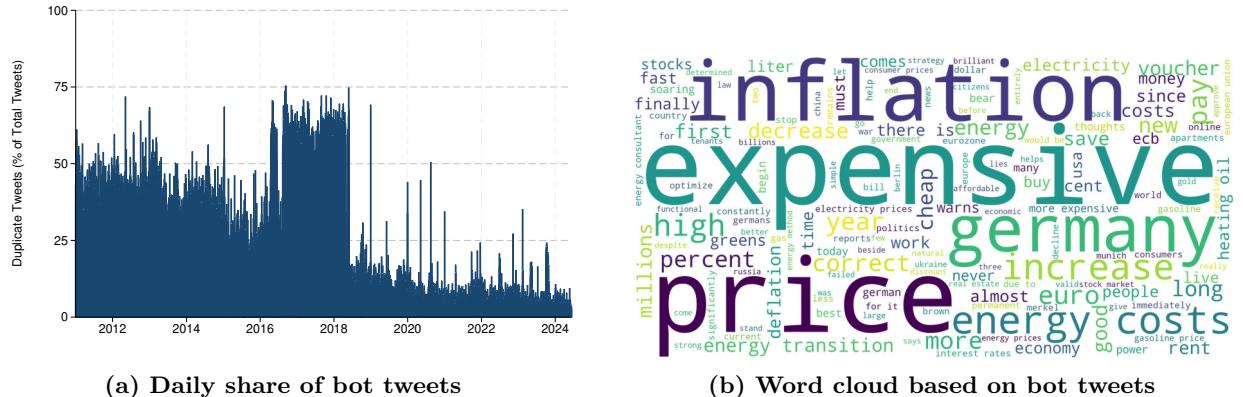
## OA.2 Bots

One important feature of Twitter data is the substantial presence of automated bot activity. A common characteristic of bots is that they repeatedly post the same tweet multiple times, often with only minor modifications—for instance, a changed URL. Retaining such tweets in the sample would artificially inflate our index and potentially bias our measure.

To address this concern, we identify and remove bot activity in two steps. First, we clean each tweet by removing URLs, user mentions, special characters, numbers, and white spaces. We then discard exact duplicates based on the cleaned text, retaining only one instance of any repeated message. This procedure removes more than two million tweets, thereby preserving the informational content while preventing automated repetitions from distorting our index. Second, we exclude tweets that are not directly related to inflation using topic modeling. This step eliminates a considerable share of advertising content, where automation is particularly prevalent.

Figure OA.1a plots the daily share of bot tweets relative to total tweets (duplicate tweets + unique tweets) and shows that automated activity is present throughout the sample, with elevated levels during 2017 and early 2018. Twitter strengthened its efforts to restrict automated posting in May 2018, which accounts for the sharp and persistent decline in the duplicate share around mid-2018.

**Figure OA.1**  
Daily share of bot tweets and the most important words



Panel (a) plots the daily share of duplicate tweets—our proxy for bot activity, calculated as  $\text{duplicates}_t / (\text{duplicates}_t + \text{unique}_t)$  for each day from January 2011 to June 2024. Panel (b) displays a word cloud of these bot tweets, with larger words representing more frequently occurring terms. The figure is based on the original German tweets, with all words translated into English.

To examine the content of tweets identified as bot-generated, Figure OA.1b presents a word cloud based on these posts. The most frequent tokens are clearly inflation-related—“inflation” (inflation), “preis/preise” (price[s]), “teuer” (expensive), and “energiekosten/energie” (energy costs/energy). This indicates that a substantial share of automated tweets concerns inflation-relevant topics.

Excluding bot-generated tweets may therefore raise the concern that we discard information relevant for households’ inflation expectations. Indeed, several of the most active “bot” accounts in

our sample belong to major news outlets *e.g.*, *FAZ Finanzen*, *FAZ Wirtschaft*, *Süddeutsche Zeitung*, *Börsen-News*) that frequently post about ECB policy decisions, energy and grocery prices, and developments in inflation expectations. Prior research has shown that such exposure can shape households' attention and inflation expectations (see, *e.g.*, Binder et al., 2025).

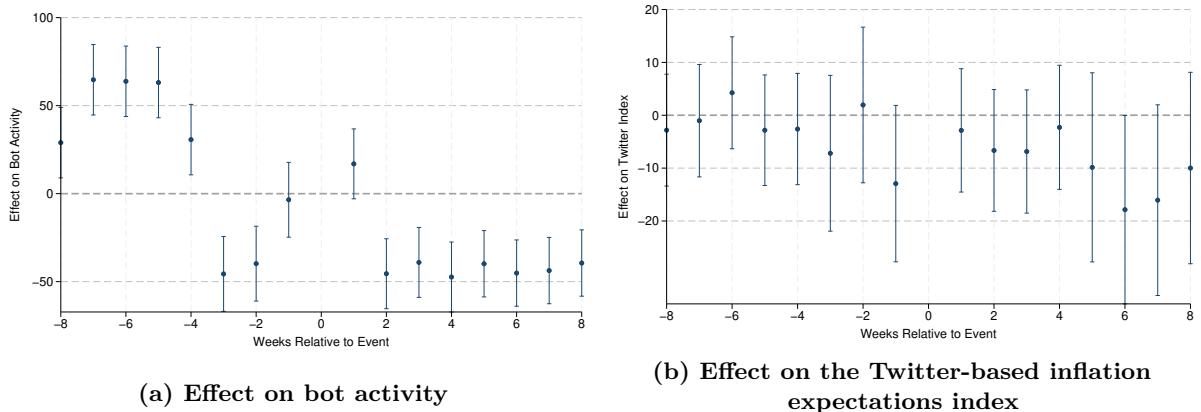
To assess whether our filtering procedure removes economically relevant information, we conduct an event study around Twitter's intensified efforts to limit automated posting in May 2018, which marked the sharpest drop in bot activity in our sample. Specifically, we estimate the following regression at the weekly frequency:

$$y_t = \sum_{k \neq 0} \beta_k 1\{\text{Relative Week}_t = k\} + X_t' \gamma + \eta_{y(t)} + u_t, \quad (8)$$

where  $y_t$  denotes either the number of duplicate tweets or the Twitter-based inflation expectations index in week  $t$ .  $X_t$  is a vector of macroeconomic controls, and  $\eta_{y(t)}$  represents year fixed effects. We report the coefficients  $\beta_k$  with 95% confidence intervals for eight weeks before and after the event.

Figure OA.2 summarizes the results. Panel (a) shows a pronounced and persistent decline in duplicate tweets beginning just before the event and deepening thereafter, consistent with Twitter's removal of suspicious accounts. By contrast, Panel (b) shows no statistically significant change in our index and substantially smaller point estimates. Hence, while Twitter's intervention markedly reduced automated posting, our index remained stable, suggesting that filtering bot activity effectively mitigates amplification without discarding relevant information for inflation sentiment.

**Figure OA.2**  
Actions against automation and spam



The figure reports coefficients from a weekly event study centered on Twitter's actions against automated and coordinated accounts in mid-May 2018. Coefficients are estimated from Equation 8, which includes macroeconomic controls, year fixed effects, and—in the case of the index regression—the lagged index. The points represent coefficient estimates, and bars denote 95% confidence intervals. Panel (a) shows results for bot tweets, and Panel (b) for our Twitter-based inflation expectations index.

### OA.3 Classification of different users

To classify the 364,330 user accounts that tweet about inflation into private individuals, media organisations and other possible user categories, we use a pre-trained machine learning model (specifically, the *twinbert-large* model) which we now fine-tune for this specific task. We provide six different labels, *private individuals*, *individual journalist*, *influencer*, *media organization*, *business organisation*, and *other organisation*. To generate training data for this fine-tuning process, we again use the OpenAI API to label almost 45,000 users, which we then use to fine-tune the *twinbert-large* model.<sup>44</sup> After running inference on all users in our sample, it became clear that the model's performance varies across classes, as indicated by the precision scores in Table OA.1.<sup>45</sup> The table shows that the model performs best for *private individuals*, which is the most important category for our purpose. We identify 194,422 private user accounts, representing more than 50% of all accounts in our sample.<sup>46</sup>

**Table OA.1**  
Precision scores of the different user categories

User class	Precision
private individuals	0.89
individual journalist	0.61
influencer	0.62
media organization	0.84
business organisation	0.67
other organisation	0.74

<sup>44</sup>We need this large amount of labelled user accounts to obtain enough training data for all the different categories as we need an equal amount of training data for each class.

<sup>45</sup>We rely primarily on the precision score, complemented by manual checks, as it is the most suitable measure for our propose. It is especially useful when the emphasis is on minimizing false positives, ensuring that when the model predicts a positive outcome, it is highly likely to be correct—*i.e.*, when the model predicts a user to be a household, the probability that this user is indeed a household is very high.

<sup>46</sup>Some media accounts were likely already removed during our bot-filtering and sample-cleaning steps.

## OA.4 Additional results

### OA.4.1 The Twitter-based index around CPI announcement days

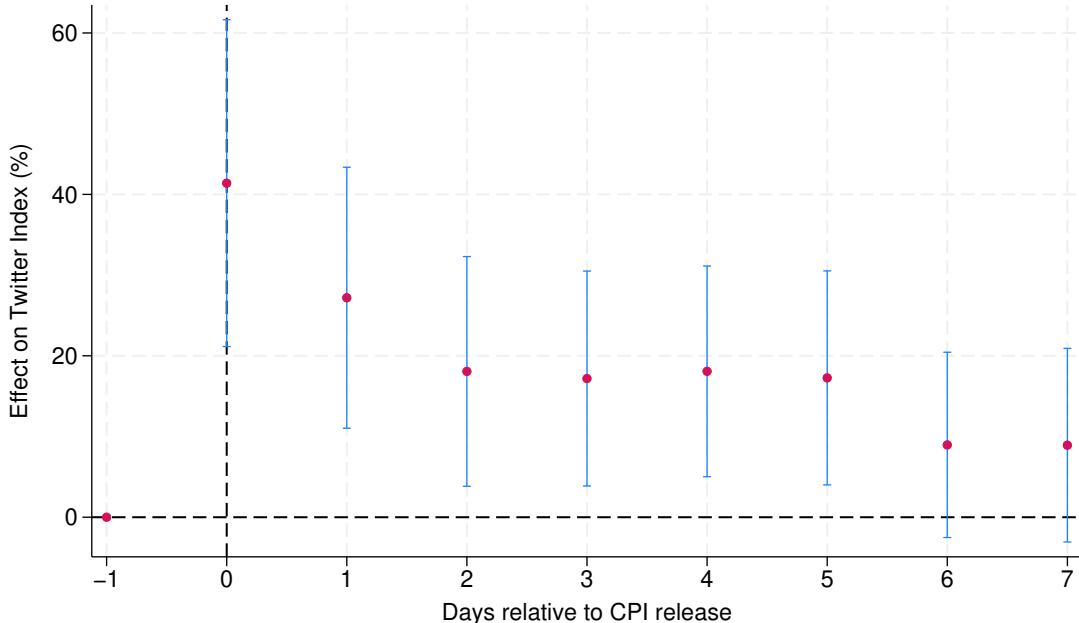
To examine whether CPI releases coincide with systematic movements in our Twitter-based inflation expectations index, we conduct an event study around CPI announcement days from January 2021 to June 2024. Specifically, we estimate the following regression at the daily frequency:

$$y_{e,t} = \sum_{k \neq -1} \beta_k \mathbf{1}\{\text{Event Day}_{e,t} = k\} + \eta_{q(t)} + \delta_e + u_{e,t}, \quad (9)$$

where  $y_{e,t}$  denotes the percentage change in the Twitter-based inflation expectations index on day  $t$  for event  $e$  relative to the trading day preceding the announcement.  $\eta_{q(t)}$  and  $\delta_e$  denote year-quarter and event fixed effects, respectively. Event time is measured in trading days, where day 0 is the CPI release day.

Figure OA.3 plots the evolution of our Twitter index around CPI releases. The coefficient for day  $-1$  is normalized to zero, allowing all estimates to be interpreted as percentage changes relative to the trading day before the announcement. The index increases sharply on the CPI release day, by roughly 40%, and remains elevated for about five trading days before gradually reverting to its pre-release level within the following week.

**Figure OA.3**  
Our Twitter-based inflation expectations index around CPI release days



The figure shows results from a daily event study around CPI releases (Jan 2021 - Jun 2024), based on Equation (9). The model includes year-quarter and event fixed effects. Dots indicate point estimates and bars represent 95% confidence intervals. Event time  $t = 0$  corresponds to the release day. Estimates are expressed as percentage deviations from the pre-announcement level at  $t = -1$ .

#### OA.4.2 Consumer behavior

**Table OA.2**  
**Effect of inflation expectations on discretionary consumption: OLS**

Days after CPI release	$\Delta_h \log(\text{Discretionary consumption}_{t+h}^i)$				
	$h = 0$	$h = 3$	$h = 6$	$h = 9$	$h = 12$
log(Index <sub>t</sub> ) (OLS)	0.04*** (0.01)	0.01 (0.02)	-0.02 (0.01)	-0.03** (0.02)	-0.04** (0.02)
$r_t^{DAX}$	-0.00 (0.00)	0.00 (0.00)	0.01** (0.00)	0.00 (0.00)	0.01* (0.00)
$\Delta \text{Bund10Y}_t$	-0.00 (0.05)	0.01 (0.07)	0.02 (0.07)	-0.06 (0.11)	-0.02 (0.08)
$\Delta \text{Yields1YEA}_t$	0.10	0.19**	-0.03	0.04	0.03
Weekend FE	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y
Year-Month FE	Y	Y	Y	Y	Y
Adj. $R^2$	0.28	0.35	0.24	0.35	0.32
N	155,697	154,973	154,308	153,663	153,041

The dependent variable in columns (1)–(5) are the log changes in daily transactions between  $t - 1$  and  $t + h$ , that is,  $\log(C_{t+h}^i) - \log(C_{t-1}^i)$ , where  $t$  denotes the day of the CPI release. Controls include one to seven lags of log transactions and of the log index (both not shown here), the daily DAX returns, daily changes in the 10-year Bund, and the one-year Euro Area yield curve spot rate, as well as weekend, month-of-year and firm fixed effects. The sample period runs from 1 January 2021 to 31 March 2024. Standard errors are two-way clustered by firm and date and reported in parentheses. \*\*\*, \*\*, and \* denote significance at the 1, 5, and 10 percent levels, respectively.

**Table OA.3**  
**Effect of inflation expectations on consumption: 2SLS panel IV estimates - All firms**

Days after CPI release	log(Index <sub>t</sub> )	First stage		Second stage			
		<i>h</i> = 0	<i>h</i> = 3	<i>h</i> = 6	<i>h</i> = 9	<i>h</i> = 12	
		$\Delta_h \log(\text{Consumption}_{t+h}^i)$					
CPI release <sub>t</sub>	0.25*** (0.05)						
log(Index <sub>t</sub> )		0.02 (0.03)	-0.01 (0.06)	-0.07* (0.04)	-0.13** (0.05)	-0.07 (0.06)	
$r_t^{DAX}$	-0.02* (0.01)	-0.00 (0.00)	-0.00 (0.00)	0.01** (0.00)	0.00 (0.00)	0.01 (0.00)	
$\Delta\text{Bund10Y}_t$	-0.11 (0.14)	0.01 (0.04)	0.00 (0.07)	-0.04 (0.07)	-0.05 (0.12)	-0.11 (0.08)	
$\Delta\text{Yields1YEA}_t$	0.72*** (0.20)	0.12* (0.07)	0.20* (0.11)	0.02 (0.09)	0.06 (0.18)	0.11 (0.13)	
Weekend FE	Y	Y	Y	Y	Y	Y	
Firm FE	Y	Y	Y	Y	Y	Y	
Year-Month FE	Y	Y	Y	Y	Y	Y	
Adj. $R^2$	0.90						
<i>F</i> statistic first stage		27.89	27.90	27.82	27.81	27.75	
<i>N</i>	956,249	953,653	948,304	944,011	939,792	935,803	

The dependent variable in column (1) is the log of the daily inflation expectations index, and those in columns (2)–(6) are the log changes in daily transactions between  $t - 1$  and  $t + h$ , that is,  $\log(C_{t+h}^i) - \log(C_{t-1}^i)$ , where  $t$  denotes the day of the CPI release. Controls include one to seven lags of log transactions and of the log index (not shown here), the daily DAX returns, daily changes in the 10-year Bund, and the one-year Euro Area yield curve spot rate, as well as weekend, month-of-year and firm fixed effects. The sample comprises the full universe of firms in the dataset, irrespective of reporting regularity or discretionary classification, covering the period from 1 January 2021 to 31 March 2024. Standard errors are two-way clustered by firm and date and reported in parentheses. \*\*\*, \*\*, and \* denote significance at the 1, 5, and 10 percent levels, respectively.