Different Newspapers—Different Inflation Perceptions*

Sarah Arndt
Heidelberg University

June 4, 2024

Abstract

In this paper, I investigate how inflation signals from different types of newspapers influence household inflation expectations in Germany. Using text data and the large language model GPT-3.5-turbo-1106, I construct newspaper-specific indicators and find significant heterogeneity in their informativeness based on the genre—tabloid versus reputable sources. The tabloid's indicator is more effective for predicting perceived inflation among low-income and lower-education households, while reputable newspapers better predict higher-income and more educated households' expectations. Local projections reveal that tabloid sentiment shows an immediate decrease following a monetary policy shock, whereas responses from reputable newspapers are smaller and delayed. Household expectations also vary depending on the type of newspaper affected by the sentiment shock and the socioeconomic background of the household. These findings underscore the differentiated impact of media on inflation expectations across various segments of society, providing valuable insights for policymakers to tailor communication strategies effectively.

JEL-Codes: E31, E58, C53, C55

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

1 Introduction

Households observe various inflation signals ranging from grocery prices to newspaper articles and social media posts. They use these sources of information to form their expectations about future inflation. Sometimes, these expectations become de-anchored and exceed what actual levels of inflation would suggest. This happened for example in Germany, in mid 2021. Even though inflation was below target in 2020, inflation expectations measured by the European Commission's (EC) *Business and Consumer Survey* were already increasing and attained its highest levels since 2011 (see Figure 1). At the same time, Germany's yellow press suddenly increased their reporting on inflation, as shown by the blue shaded area in Figure 2, representing the absolute weekly number of articles related to the topic of inflation published by Germany's largest boulevard newspaper— *Die Bild.* This press is known for its polarizing style of reporting and published headlines such as *"The inflation monster is back!"* (May 8, 2021), even though inflation was barely exceeding 2% back then, and one month later: *"Inflation is eating away our savings"* (June 2, 2021), fueling households' fear of inflation. There is the potential risk that the public opinion had been distorted by this sudden increase of dramatized media coverage.

This paper takes a closer look at the different signals towards inflation emitted by the yellow press, represented by Germany's largest newspaper in-terms of sold units—*Die Bild*, and more reputable German daily newspapers like the *Frankfurter Allgemeine Zeitung* (FAZ) and *Die Süddeutsche Zeitung* (SZ), being the two most popular representatives. I collect inflation related articles of these three newspapers and analyze their sentiment towards inflation with the large language model gpt-3.5-turbo-1106 developed by OpenAI. Out of the raw sentiment scores per article, I use different methods to compute overall sentiment indices per newspaper.

As it is well established that some socioeconomic groups of households have persistently different inflation expectations (Meeks and Monti, 2023), I test whether different newspaper sentiment indices can account for this heterogeneity. To do so, I compare the newspaper indices with household survey perspectives on inflation, differentiated by age, income, education, and gender of the survey participants. The analysis reveals heterogeneity in contemporaneous correlations, suggesting that expectations derived from the yellow press resonate more with lower education households, while the opposite holds true for FAZ and SZ.

In a second step, I analyze the forecasting power of the three news indices by targeting different household sub-categories' survey expectations and perceived inflation through an out-of-sample forecasting exercise. The results reveal heterogeneity among the indices, suggesting that the Bild index is more suitable for predicting perceived inflation among lower income and lower education survey participants. Conversely, the FAZ and SZ indices outperform the Bild index in predicting overall inflation in Germany. In general, all newspaper indices carry some degree of informativeness regarding household survey

measures on inflation, as including them in a forecasting model reduces the root mean squared error in most cases. This makes these indices an attractive alternative measure of household inflation expectations and perceptions due to their high-frequency real-time availability.

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

Lastly, I explore another potential source of heterogeneity by comparing the responses of household inflation expectations and perceptions to shocks in the sentiment of the three newspapers. Survey participants with higher income and education levels are most responsive to shocks in FAZ and SZ sentiment. In contrast, shocks in Bild sentiment primarily affect households in the second income quartile, although overall heterogeneity is less pronounced for Bild. This indicates that some of the heterogeneous informativeness of the indicators can be attributed to their varying ability to influence household perspectives on inflation.

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

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

but to test the epidemiological model introduced by Carroll (2003) to see if tone and news coverage bring household expectations closer to those of professional forecasters.

Furthermore, this paper contributes to the literature on inflation expectation elicitation by offering a real-time proxy to measure heterogeneous household expectations. For an overview of recent literature on inflation expectations, readers are referred to D'Acunto et al. (2023). Weber et al. (2023) find that households are better informed about inflation and more attentive to news about it during times of rising prices. Consistent with their results, I find that the correlation between newspaper indices and EC household survey expectations from April 2020 to August 2023—which overlaps with the recent high inflation period in Europe—is higher than in previous years (see Figure D-2 in Appendix D). Additionally, Andre et al. (2023) emphasize the importance of media in shaping and creating macroeconomic narratives that ultimately affect household expectations.

The paper is structured in the following way: Section 2 describes the newspaper and survey data set. Section 3 outlines construction of the newspaper inflation expectations indices, followed by an analysis of their informativeness in Section 4. I present the effects of a monetary policy shock on the newspaper indices in Section 5 and the responses of survey expectations after a shock in news sentiment in Section 6. Finally, Section 7 provides some robustness checks and Section 8 concludes.
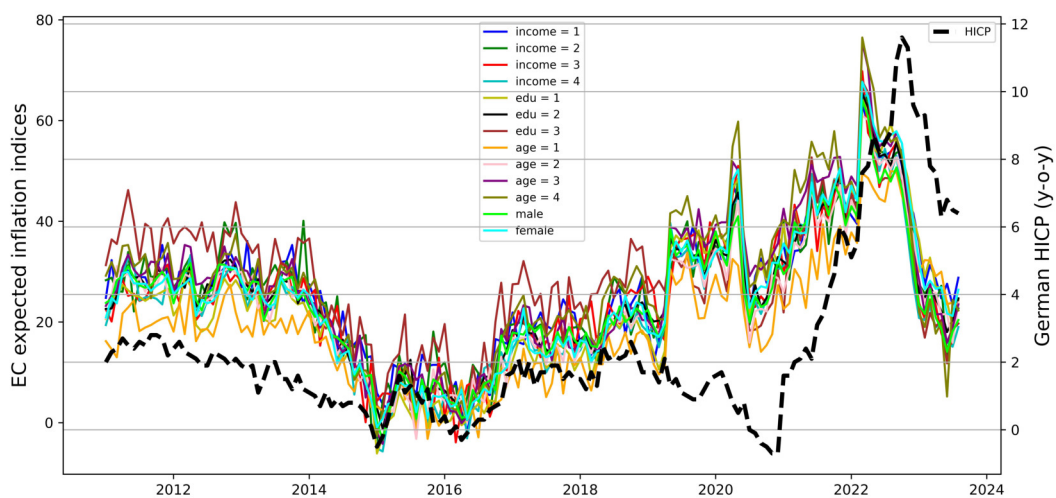

Figure 1: Left axis: EC survey expectation indices for different socioeconomic groups. Right axis: monthly year-on-year growth rates of German HICP (black dashed line). Sample span: January 2011 to August 2023.

2 Data sources

2.1 Newspaper data

In Germany, there is only one popular daily nationwide yellow press newspaper—*Die Bild* (Bild), which recorded 1,146,616 sold units of its Sunday edition in the third quarter of 2023 and about 500 million page visits each month in 2023 on its online news portal *Bild.de* (source: Informationsgemeinschaft zur Feststellung der Verbreitung von Werbeträgern e. V. (IVW) (2024)). It is widely known for its polarizing style of reporting, characterized by exaggerations, neologisms, and an oversimplified language. In contrast, the two most popular nationwide daily newspapers considered serious news sources are *Die Süddeutsche Zeitung* (SZ) and the *Frankfurter Allgemeine Zeitung* (FAZ). Each recorded 454,319 and 498,730 sold units respectively in the third quarter of 2023, with around 55 and 65 million monthly page visits on their news portals *Süddeutsche.de* and *FAZ.net* in 2023 (source: IVW, 2024). I use articles from both SZ and FAZ to increase the sample of "serious" news articles and cover a larger segment of the German population.

I collect SZ and FAZ articles from their respective news archives using search queries that filter articles related to inflation in Germany. The queries employ German terms such as *central bank*, *high prices*, *deflation*, *inflation*, *price index*, *wage-price-spiral*, or *monetary policy* (detailed search queries for all newspapers can be found in Appendix A). Both archives include print and online articles. Bild newspaper articles are obtained from the online database *LexisNexis*. Since Bild articles are only available until 2017 there, I additionally use the open-source platform *mediacloud.org* (hereafter mediacloud) to collect URLs directing to inflation-related Bild online articles. I scrape all collected web pages to obtain the article data. This procedure results in a total dataset of 13,365 SZ, 16,157 FAZ, and 16,002 Bild articles published from January 2010 to August 2023. Despite using specific search queries, not all collected articles are actually concerned with inflation in Germany; this is particularly true for Bild articles. Therefore, Figure 2 depicts the distribution of articles over time that are genuinely related to inflation in Germany against the year-on-year growth rates of the German harmonized index of consumer prices (HICP). At the beginning of the sample period, there are very few Bild articles on inflation due to limited data availability in the respective online databases. The number of articles increases significantly starting from 2021. In contrast, a substantial number of inflation-related articles is available throughout the entire sample period for SZ and FAZ, with a noticeable increase starting in 2021 that coincides with rising year-on-year growth rates of German HICP.

2.2 Survey data

To assess the informativeness of the newspaper indicators regarding household's perceived inflation and inflation expectations, I compare them to the EC's *Business and Consumer*
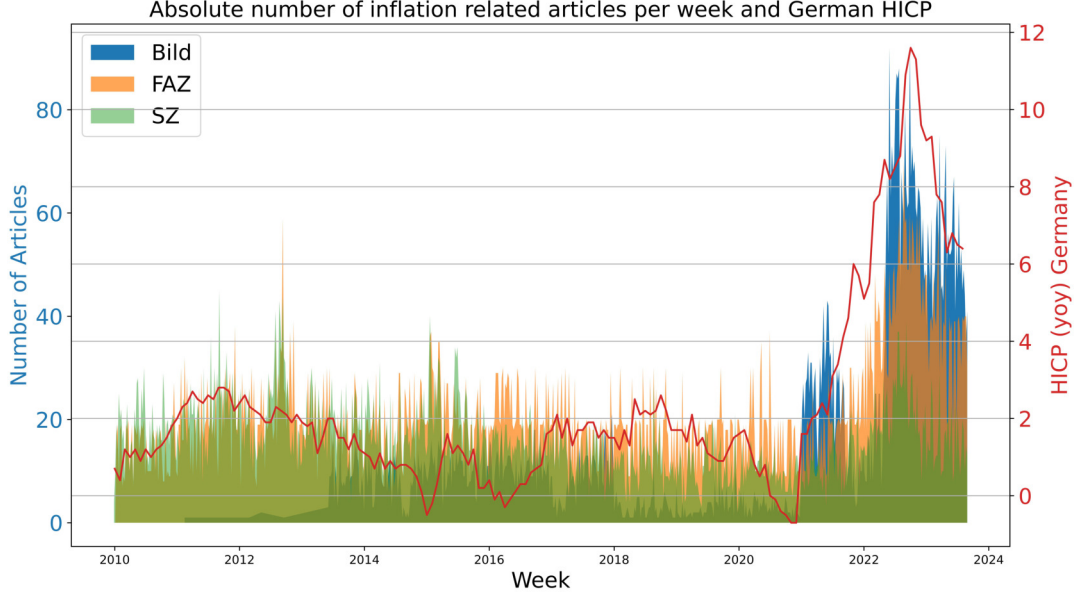
Figure 2: Left axis: weekly absolute numbers of SZ (green), FAZ (orange), and Bild (blue) newspaper articles dealing with inflation in Germany (selection based on classification by ChatGPT). Right axis: monthly year-on-year growth rates of German HICP (red line). Sample span: January 2010 to August 2023.

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

$$\mathbb{E}_t^{\pi_{t+12}} = PP + \frac{P}{2} - \frac{M}{2} - MM. \tag{1}$$

Similarly, to measure the household's perceived inflation, the survey asks: *"How do you think that consumer prices have developed over the last 12 months? They have...* with the following possible answers: *"risen a lot" (PP), "risen moderately" (P), "risen slightly", "stayed about the same (M), fallen (MM), or "don't know."* The EC provides seasonally adjusted data, disaggregated at the following subcategories: age (four bins), education (primary, secondary, and further), gender, income quartiles and occupation status.
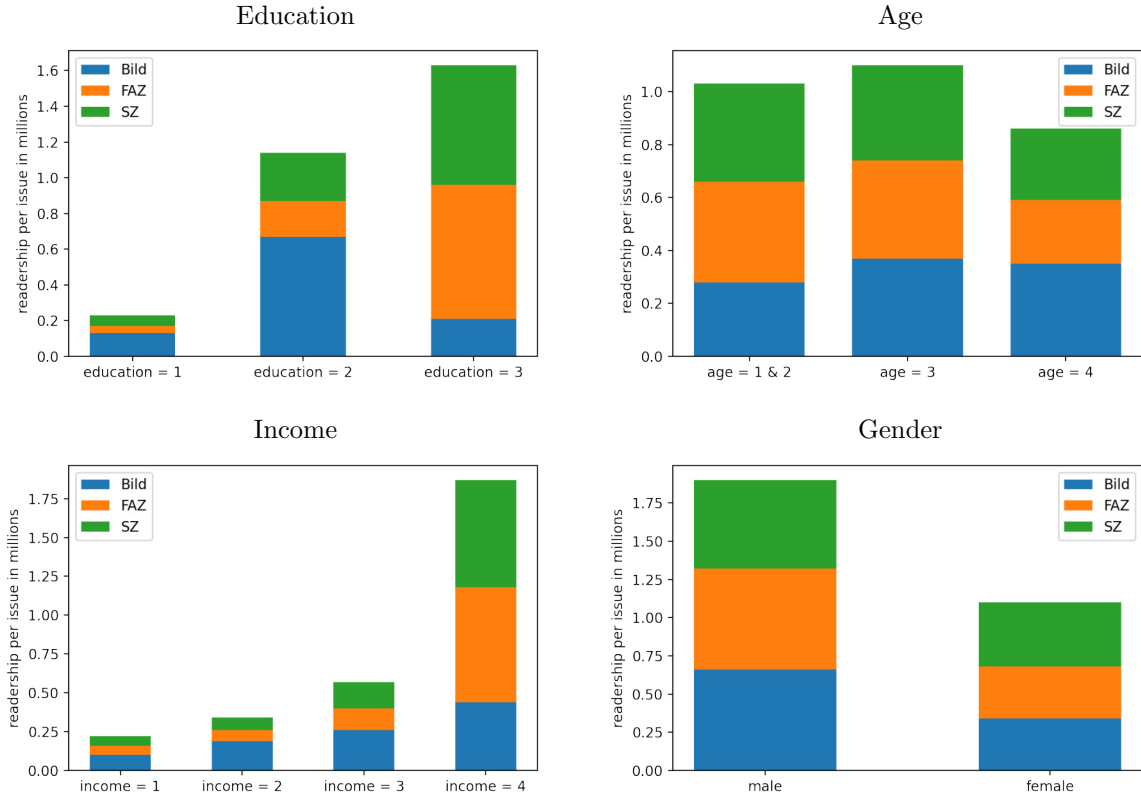
Figure 3: Readership per Bild (blue), FAZ (orange), or SZ (green) issue in 2022, normalized to 1 million readers per newspaper. Disaggregated per education and income level, age cohort and gender such that the data matches the EC survey household subcategories. Data source: agma (2022).

2.3 Readership data

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

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

3 Newspaper-based inflation expectation indices

3.1 Applying ChatGPT for sentiment analysis

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

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

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

Figure 4: Prompt for generating inflation expectation sentiments with ChatGPT-3.5. "%title" and "%article" are place holders for single articles and their headlines, respectively.

In comparison to simpler methods for text-based sentiment measurement, like the widely-used dictionary approach (see for example Shapiro et al. (2022) or Angelico et al. (2022)), an advantage of ChatGPT is its advanced ability to understand human language, capturing also nuances in the tone of the text. This is especially important when analyzing yellow press articles like those of the Bild newspaper, as they are characterized by a polemic language and constant creation of new expressions. The highly adaptive nature of the Bild language makes creating a Bild suited dictionary impractical as it would require to expand constantly. Table 1 presents examples of Bild expressions and tentative English translations thereof. As shown in Appendix B, ChatGPT is indeed capable of understanding the expressions highlighted in Table 1 and captures the headline's intended tone.

Besides ChatGPT's extended ability to understand human language, the model additionally proves to perform well at analyzing the sentiments of texts. Belal et al. (2023) find that ChatGPT-3.5 significantly outperform two lexicon-based algorithms (VADER and TextBlob) in analyzing the sentiment of amazon product reviews and a soccer tweets

7

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

Table 1: Examples of Bild typical languange and their use of neologisms in the context of inflation in Germany and the euro zone.

data set. Comparing ChatGPT-3.5 to a fine-tuned BERT[1] model, Wang et al. (2023) find that ChatGPT scores equally well on zero-shot sentiment classification tasks and that it even beats the BERT model in recognizing sentiment polarity shifts.

3.2 Index construction

The raw data generated by ChatGPT consists of article-level sentiment scores from 1 to 10, or 'no' if the article does not cover German inflation. Since the frequency of articles is daily, it is feasible to aggregate the data into indices at any higher frequencies. In this analysis, I focus on monthly aggregates to align with the monthly frequency of the household expectation surveys. Similar to Angelico et al. (2022), I use multiple methods to aggregate the single scores to sentiment indices for each newspaper separately and only include those articles, which ChatGPT deems to be related to inflation in Germany:

1. **Index #1:** I take the mean over all articles published per month to obtain a monthly sentiment index.

2. **Index #2:** First, I obtain mean scores per day to then compute moving averages (MA) over a window of 10, 30 or 60 days (MA_10, MA_30 & MA_60). I finally aggregate them to monthly frequency by taking the average for each month.

3. **Index #3:** Analogously to Angelico et al. (2022), I compute exponentially weighted averages over the daily mean scores using three alternative smoothing parameters: 0.1, 0.3 and the one that minimizes the sum-of-squared in-sample forecast errors, which is close to 0.1 (indices are denoted by EWM_01, EWM_03 & EWM_opt).

---

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

4. **Index #4:** I reproduce all indices #1 to #3 including only articles that score 5 or above on ChatGPT's inflation sentiment scale.

I choose to construct the indices also on a sub-sample based on articles with higher sentiment scores as it might be the case that newspapers rather report on increasing inflation than on stable, low price levels. The increase in the number of inflation related articles shown in Figure 2 as well as the worsening of the average sentiment depicted in Figure 5 is supporting evidence for this hypothesis. Therefore, ChatGPT might also be more reliable when assigning high sentiment scores. Similarly, Weber et al. (2023) find that households also pay more attention to inflation, when inflation is increasing, therefore an index based on higher scores might be more successful in capturing their inflation expectations.

Figure 5 exemplary shows one index from each of the categories #1 to #3 as well as their versions based on articles with sentiment scores greater than 4 for each of the newspapers separately and the FAZ and SZ articles data combined. Overall, the EWM_01 and MA_60 indices seem to be less noisy than the mean indices. While the indices based on FAZ and SZ data are very well aligned with each other, the Bild inflation indices significantly differ. In all depicted cases, the Bild inflation index reaches higher sentiment scores than the other two newspapers, which is especially visible in the years around 2019. Actual inflation in Germany was mostly below the ECB's target during this period, so it is surprising that Bild's sentiment is significantly higher than that of the other two newspapers. A closer examination of the articles from that time reveals the cause of this spike in sentiment: In 2019, the ECB's key interest rate hit the zero lower bound, prompting banks to introduce negative interest rates on deposits. Bild reported very negatively on this development, spreading the narrative that it imposed substantial costs on "normal people" due to the lack of interest on their savings.[2] Additionally, Bild articles exaggerated the severity of inflation. Any positive level of inflation, even if below 2%, was portrayed very negatively, as it further eroded the purchasing power of deposits. Conversely, deflation was depicted as highly beneficial.

Due to the scarcity of Bild articles at the beginning of the sample, the Bild sentiment index is often inconclusive for the period before 2014. Therefore, I truncate the sample at 2014 in most of the subsequent analyses. Comparing the newspaper indices to actual year-on-year growth of German HICP reveals some correlation over time, with the newspaper indicators beginning to align with inflation trends at the end of 2021. When considering only articles with a sentiment score above 4, it becomes evident that the disparity between Bild and the FAZ and SZ newspapers is driven by Bild's abundance of articles suggesting high or increasing inflation.

---

[2]Mario Draghi, the ECB's president at that time, was dubbed "Graf Draghila" by Bild, a reference to Count Dracula (see, for example, articles from 02/01/2020, 10/24/2019, and 09/12/2019).
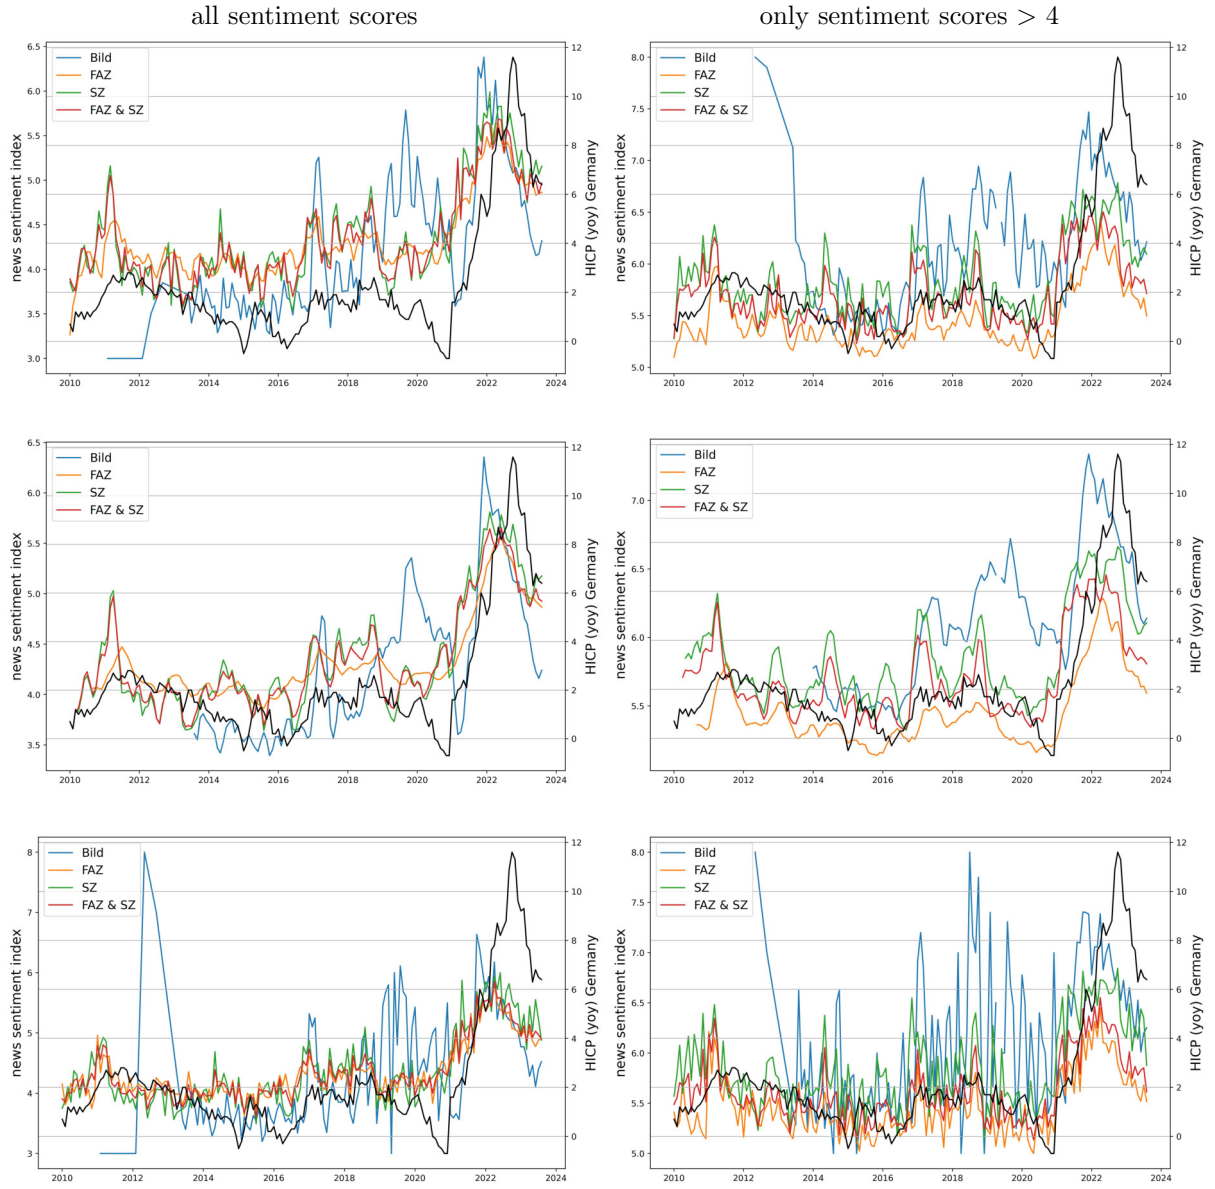
Figure 5: Left column: shows the EWM_01, MA_60 and mean (top to bottom) sentiment indices (left axis) for each newspaper on a monthly frequency against German year-on-year HICP growth rates (black line, right axis). Right column: corresponding versions of the indices using only articles with sentiment scores above 4.

4 Informativeness of the newspaper indices

4.1 Heterogeneous correlations

First, I compare the sample correlations between the inflation sentiment indices of different newspapers and the various subgroups of the EC household survey measures of perceived and expected inflation. This analysis reveals initial heterogeneity in the strength of the co-movement between each newspaper index and each household subgroup. Specifically, I calculate the sample correlation with a one-month lag for the newspaper indices, as

households are surveyed in the first weeks of a month, ensuring that the newspaper indices include only the information available to households at that time.

Figure 6 presents the resulting correlation heat maps between the 14 different newspaper indices (on the x-axis) and the EC survey's measures of inflation expectations (lower left tile) and German inflation (last row). I report results for FAZ and SZ articles combined, as they yield very similar correlations individually. Generally, when considering the full sample starting in January 2014, Bild achieves the highest sample correlations with household inflation expectations (ranging between 0.53 and 0.82) among the three newspapers. Indices based on scores higher than 4, shown in the left half of the lower left tile of the heat map, produce lower correlations than those based on all sentiment scores. While the Bild indices are only weakly correlated with German HICP (last row in the heat map), FAZ and SZ indices attain higher correlations with inflation than the survey expectations themselves. For Bild, there is visible heterogeneity in correlations between the different household groups, with the largest correlations seen among lower income and lower education households. A similar pattern is observed for FAZ and SZ sentiments, though the differences are much smaller.

In contrast, as shown in the lower left tile of the heat maps in Figure 7, the correlation of Bild indices is overall smaller with EC survey measures of perceived inflation (y-axis), while it is larger for FAZ and SZ inflation sentiments. Upon closer examination, one can detect some heterogeneity between the correlations with the indices of different household categories. For Bild, correlation is highest with lower education and middle income households. For FAZ and SZ, this heterogeneity is less pronounced, but correlations with higher education and higher income households are slightly larger compared to other household groups
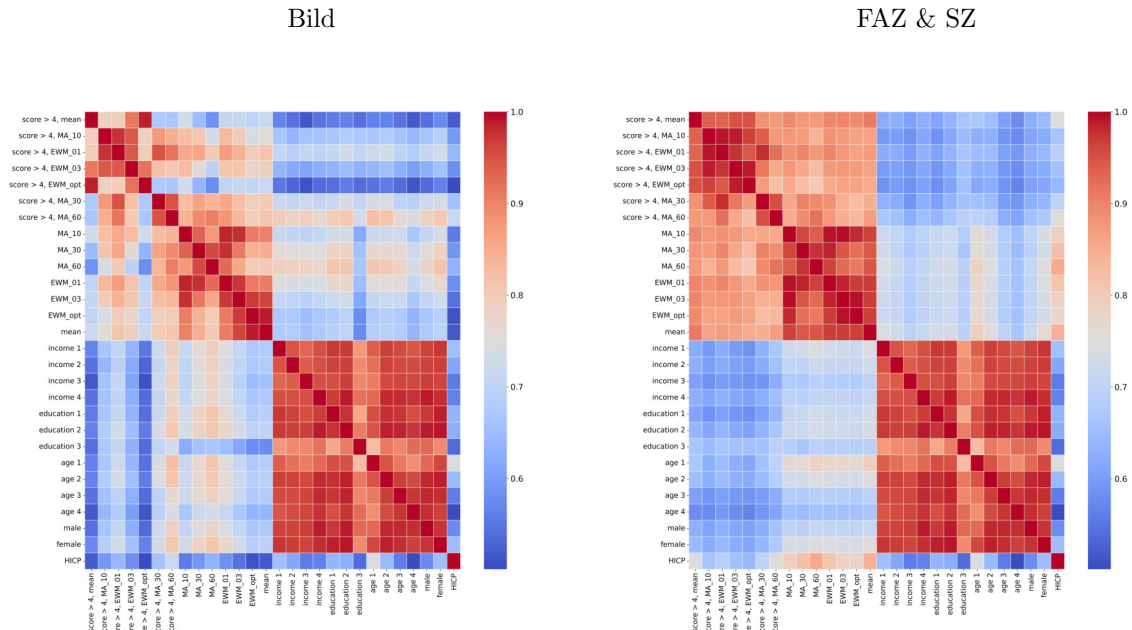

Figure 6: Sample correlations between EC inflation expectations and Bild (left) and FAZ and SZ combined (right) inflation indices based on the whole sample period (2014M01-2023M08). Red indicates high and blue low correlations.

11

Figure 7: Sample correlations between EC perceived inflation measures and Bild (left) and FAZ and SZ combined (right) inflation indices based on the whole sample period (2014M01-2023M08). Red indicates high and blue low correlations.
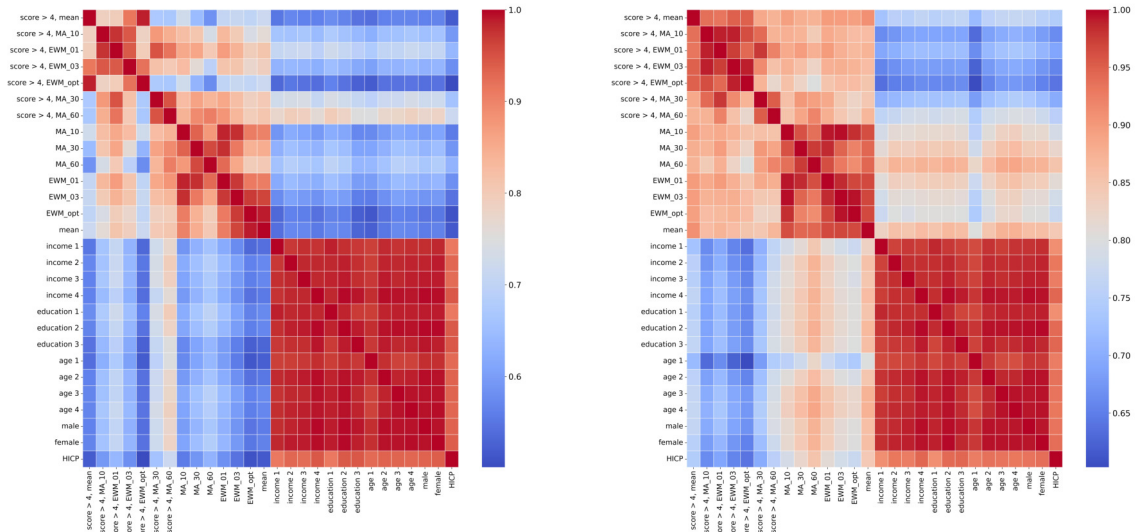
4.2 Heterogeneous forecasting powers

In a second step, I run an out-of-sample forecasting exercise and compare the ability of the different newspaper indices to predict the survey expectations and perceived inflation measures of the various groups of households. As benchmark, I estimate an autoregressive model, $AR(p)$, for each of the sub-aggregated survey expectation measures separately, determining the optimal lag order p for each model and forecast iteration with the BIC criterion. The in-sample ends in December 2016, translating into a training sample length of 84 months for SZ and FAZ newspapers and 36 months for the Bild, due to the shorter data availability of the latter. I then augment each benchmark model with one of the 14 different inflation expectation indices for each newspaper separately and also a combined version of FAZ and SZ data. To be conservative, as the BCS participants are surveyed in the first weeks of a month, all newspaper indices enter with one lag. The forecast horizon is six months and all horizons are predicted recursively.

I compare the forecasting powers by calculating the root mean squared error (RMSE) for each model. I divide the augmented model's RMSE by the benchmark's RMSE such that a value smaller than one indicates an improvement in forecast accuracy when including the newspaper indices. In general, the RMSE ratios are very similar across the 14 aggregation methods for each newspaper with minor differences. Therefore, and to make the results more tractable, I only report their average values here. test of equal forecast accuracy.

Table 2 presents the relative RMSEs based on EC survey data on expected inflation. The results indicate that, in most cases, including the newspaper expectation indices

improves the benchmark model. Overall, and despite the smaller training sample, the Bild indices achieve the lowest RMSE ratios, significantly outperforming the benchmark model. There is notable heterogeneity between the newspapers. While the Bild indices are more effective in predicting expectations for lower income quartiles, FAZ and SZ indices also provide valuable information for the third income quartile. Additionally, when comparing education levels, FAZ and SZ offer the largest information gains for the third education level, even though these results do not pass the DM-test. For longer horizons, FAZ and SZ do not improve predictions of gender-specific expectations or expectations of age classes 3 and 4 (ages 50 to 65+). In these cases, starting from horizon 2, the Bild indices can lead to improvements. Nonetheless, FAZ and SZ have a greater informativeness advantage when predicting inflation in Germany compared to Bild.

Table 3 reports the results of the out-of-sample forecasting exercise for predicting perceived inflation measures. Consistent with the sample correlations, the FAZ and SZ indices outperform the Bild indices, with all their RMSE ratios being smaller than one and generally lower than those of Bild indices. Within each newspaper, a similar pattern of heterogeneity is observed. Across income and education levels, the Bild indices are most informative for the lowest levels in each category. In contrast, FAZ and SZ indices outperform the benchmark model when predicting perceived inflation for higher education households. Additionally, they achieve the lowest average RMSEs for households in the third income quartile.

In accordance with the readership characteristics depicted in Figure 3, these findings suggest a pattern wherein the inflation expectation indicators from Bild more accurately predict expectations of households with lower income and education levels. Conversely, the sentiment reflected in newspapers like FAZ and SZ tends to resonate more with higher-income, more educated households, and aligns more closely with actual inflation figures. Moreover, it appears that Bild's sentiment serves as a better predictor of household inflation expectations, whereas FAZ and SZ articles offer more reliable signals for predicting perceived inflation.

News	income 1	income 2	income 3	income 4	edu 1	edu 2	edu 3	age 1	age 2	age 3	age 4	male	female	HICP
Bild	**0.99**	**0.99**	1.01	1.03	**1.00**	1.00	1.02	1.00	1.02	1.00	1.02	1.00	1.01	**0.99**
	0.96*	0.97*	1.00	1.01	0.98	0.98	0.98	0.94***	0.99	0.99	1.00	0.98	0.99	0.98*
	0.94**	0.95**	0.98*	0.99	0.96*	0.97**	0.96**	0.91***	0.98	0.97	0.99	0.97*	0.97*	0.97**
	0.91***	0.92***	0.96**	0.97*	0.94**	0.95***	0.93***	0.88***	0.96**	0.96**	0.99	0.97**	0.95**	0.95***
	0.89***	0.89***	0.94***	0.95**	0.91***	0.93***	0.90***	0.85***	0.94***	0.94**	0.97**	0.95***	0.93***	0.94***
	0.86***	0.86***	0.91***	0.92***	0.89***	0.91***	0.87***	0.83***	0.91***	0.92***	0.96***	0.93***	0.91***	0.93***
FAZ	0.94*	0.92*	0.98	0.99	0.98	1.00	0.95	0.90***	0.99	1.00	1.00	0.99	1.00	0.88**
	0.95	0.92*	0.96	1.00	0.99	1.01	0.95	0.86***	0.97	1.01	1.02	1.01	1.01	0.81***
	0.94*	0.92**	0.96	1.00	0.99	1.02	0.94	0.84***	0.97	1.02	1.03	1.01	1.01	0.77***
	0.91**	0.92**	0.94	1.01	0.99	1.02	0.94	0.82***	0.96	1.03	1.03	1.01	1.01	0.71***
	0.90**	0.91**	0.94*	1.01	0.98	1.01	0.93	0.80***	0.96	1.03	1.03	1.01	1.00	0.66***
	0.90**	0.92**	0.94*	1.02	0.99	1.01	0.93	0.80***	0.96	1.03	1.03	1.02	1.00	0.63***
SZ	0.95*	0.93*	0.99	0.99	0.98	1.01	0.96	0.91**	0.99	1.00	0.99	0.99	1.00	0.90**
	0.96	0.92*	0.96	0.99	0.98	1.01	0.95	0.86***	0.97	1.00	1.01	**1.00**	1.00	0.86***
	0.95*	0.92**	0.96	1.00	0.98	1.01	0.95	0.84***	0.97	1.01	1.02	1.00	1.01	0.83***
	0.93**	0.92*	0.95	1.01	0.98	1.02	0.95	0.82***	0.97	1.02	1.03	1.01	1.01	0.80***
	0.93**	0.91**	0.95	1.02	0.98	1.02	0.94	0.81***	0.97	1.03	1.03	1.01	1.01	0.77***
	0.94*	0.92*	0.94	1.02	0.99	1.01	0.95	0.80***	0.98	1.03	1.03	1.02	1.01	0.75***
FAZ & SZ	0.94**	0.92**	0.98	0.99	0.98	1.00	0.96	0.90**	0.98	0.99	0.99	0.99	1.00	0.89**
	0.95*	0.91**	0.95	0.98	0.98	0.99	0.93	0.85***	0.96	1.00	1.00	0.99	0.99	0.84***
	0.93**	0.91**	0.94*	0.99	0.98	0.99	0.92	0.83***	0.95	1.00	1.01	0.99	1.00	0.80***
	0.91**	0.91**	0.92*	1.00	0.98	0.99	0.91	0.81***	0.95	1.01	1.01	0.99	0.99	0.75***
	0.90**	0.90**	0.92**	1.00	0.97	0.99	0.90*	0.79***	0.94	1.01	1.01	0.99	0.99	0.72***
	0.90**	0.91**	0.91**	1.00	0.98	0.98	0.90*	0.78***	0.94	1.01	1.01	0.99	0.98	0.70***

Table 2: Relative root mean squared errors of the averaged newspaper indicators of the FAZ, SZ, Bild and FAZ & SZ combined compared to EC survey data. The symbols ***, **, and * represent significance levels of 1%, 5%, and 10%, respectively, in the Diebold–Mariano test for equal forecast accuracy. The rows in each block mark the different forecast horizons (one to six months ahead).

News	income 1	income 2	income 3	income 4	edu 1	edu 2	edu 3	age 1	age 2	age 3	age 4	male	female	HICP
	0.97**	1.02	1.02	1.00	1.00	1.02	1.00	0.97**	1.00	1.03	1.01	1.02	0.99	0.99
	0.95***	1.03	1.03	1.00	0.99	1.01	0.99*	0.95***	0.99	1.03	0.99	1.00	0.98	0.98*
Bild	0.93***	1.02	1.01	0.99*	0.96**	1.00	0.99*	0.93***	0.98**	1.01	0.98**	0.99*	0.97**	0.97**
	0.91***	0.99*	0.97**	0.97***	0.94***	0.97***	0.98**	0.91***	0.96***	0.98***	0.96***	0.97***	0.94***	0.95***
	0.90***	0.96***	0.94***	0.95***	0.91***	0.94***	0.98***	0.90***	0.94***	0.95***	0.94***	0.95***	0.92***	0.94***
	0.88***	0.95***	0.91***	0.93***	0.89***	0.92***	0.97***	0.89***	0.92***	0.93***	0.92***	0.93***	0.89***	0.93***
	0.91**	0.93*	0.92**	0.90***	0.93**	0.91**	0.83***	0.92**	0.92**	0.91**	0.92**	0.89***	0.91***	0.88**
	0.83***	0.87**	0.80***	0.83***	0.88***	0.84***	0.76***	0.83**	0.85***	0.86***	0.86***	0.83***	0.85***	0.81***
FAZ	0.73***	0.80***	0.72***	0.78***	0.78***	0.78***	0.72***	0.78***	0.76***	0.80***	0.80***	0.77***	0.78***	0.77***
	0.66***	0.75***	0.66***	0.75***	0.70***	0.76***	0.70***	0.74***	0.74***	0.77***	0.76***	0.75***	0.73***	0.71***
	0.60***	0.69***	0.61***	0.72***	0.68***	0.73***	0.68***	0.70***	0.71***	0.74***	0.73***	0.73***	0.69***	0.66***
	0.58***	0.67***	0.58***	0.70***	0.67***	0.71***	0.66***	0.68***	0.71***	0.73***	0.71***	0.72***	0.67***	0.63***
	0.93**	0.94*	0.92**	0.91**	0.95*	0.91**	0.84***	0.94*	0.93*	0.91**	0.94*	0.89**	0.93**	0.90**
	0.86***	0.89**	0.81***	0.84***	0.89**	0.83***	0.76***	0.85**	0.85***	0.86***	0.87***	0.82***	0.87***	0.86***
SZ	0.79***	0.83***	0.72***	0.79***	0.81***	0.77***	0.70***	0.81***	0.77***	0.80***	0.80***	0.76***	0.81***	0.83***
	0.72***	0.78***	0.66***	0.77***	0.74***	0.75***	0.68***	0.78***	0.74***	0.77***	0.75***	0.74***	0.76***	0.80***
	0.68***	0.71***	0.60***	0.74***	0.72***	0.72***	0.65***	0.74***	0.71***	0.75***	0.71***	0.72***	0.72***	0.77***
	0.67***	0.70***	0.57***	0.72***	0.71***	0.70***	0.63***	0.73***	0.70***	0.73***	0.68***	0.70***	0.70***	0.75***
	0.92**	0.94	0.91**	0.89***	0.94*	0.90**	0.83***	0.93*	0.93*	0.90**	0.93**	0.88***	0.92**	0.89**
	0.84***	0.88**	0.78***	0.81***	0.88***	0.81***	0.75***	0.82***	0.84***	0.84***	0.85***	0.80***	0.85***	0.84***
FAZ & SZ	0.75***	0.81***	0.68***	0.75***	0.78***	0.74***	0.68***	0.78***	0.73***	0.77***	0.77***	0.73***	0.77***	0.80***
	0.67***	0.75***	0.62***	0.72***	0.69***	0.71***	0.65***	0.74***	0.70**	0.73***	0.71***	0.70***	0.71***	0.75***
	0.61***	0.67***	0.56***	0.69***	0.66***	0.67***	0.62***	0.70***	0.66***	0.70***	0.67***	0.67***	0.66***	0.72***
	0.60***	0.65***	0.52***	0.67***	0.65***	0.65***	0.60***	0.69***	0.65***	0.68***	0.64***	0.66***	0.64***	0.70***

Table 3: Relative root mean squared errors of the averaged newspaper indicators of the FAZ, SZ, Bild and FAZ & SZ combined compared to EC survey data on perceived inflation. The symbols ***, **, and * represent significance levels of 1%, 5%, and 10%, respectively, in the Diebold–Mariano test for equal forecast accuracy. The rows in each block mark the different forecast horizons (one to six months ahead).

Following Angelico et al. (2022), another way to evaluate the forecast performance of competing models against each other and further identify heterogeneity in the predictability of household expectations by the different newspaper indices is by computing the cumulative sum of squared forecasting errors differences (CSSED), defined as:

$$CSSED_{m,t} = \sum_{t=1}^{T} (\hat{e}_{bm,t}^2 - \hat{e}_{m,t}^2), \tag{2}$$

where $\hat{e}_{bm}$ and $\hat{e}_m$ correspondingly denote the forecasting errors of the benchmark and newspaper indices augmented models in the out-of-sample. Similar to the previous analysis, I average the forecast errors across the 14 aggregation methods for the results I report here. If the CSSED falls below 0 at a particular juncture during out-of-sample prediction, it indicates that when we divide the in-sample and evaluation samples at that juncture, the benchmark model performs better, displaying a lower RMSE compared to the competing model. Conversely, if the CSSED exceeds 0, then the competing model surpasses the benchmark.

Figure 8 shows the one month ahead CSSEDs of the averaged newspaper indices targeting EC survey inflation expectations (plots of the CSSEDs of the remaining forecast horizons can be found in Appendix E). Again, one can see heterogeneity in which newspaper index is more useful in predicting the inflation expectations of different household groups. In general, the Bild index appears to be more informative for most groups of households than the other two newspapers as the CSSED exceeds zero already early on in the out-of-sample, while for FAZ and SZ this is mainly the case starting in 2021. The Bild index generates the largest CSSED for male, lower income and lower education, as well as middle aged respondent's expectations and is at par with the benchmark when predicting HICP. In contrast, for the FAZ and SZ indices the CSSEDs are especially large when predicting inflation expectations of higher educated participants, or those belonging to the second or third income quartile.

The CSSED results for perceived inflation, shown in Figure 9, are similar to those for expected inflation but exhibit slightly more heterogeneity across household sub-categories. Bild indices improve the benchmark models primarily for households in the first income quartile and for very young survey participants, particularly towards the end of the sample period. FAZ and SZ indices begin to outperform the benchmark only later in the sample, around 2020. The largest CSSEDs are achieved when predicting perceived inflation for households with the highest education level and those in the third income quartile.

5 Monetary policy shocks and news sentiments

To explore a possible source of the heterogeneity between the newspaper indices, I estimate the newspaper's sentiment reaction after a monetary policy shock. Therefore, I use a local
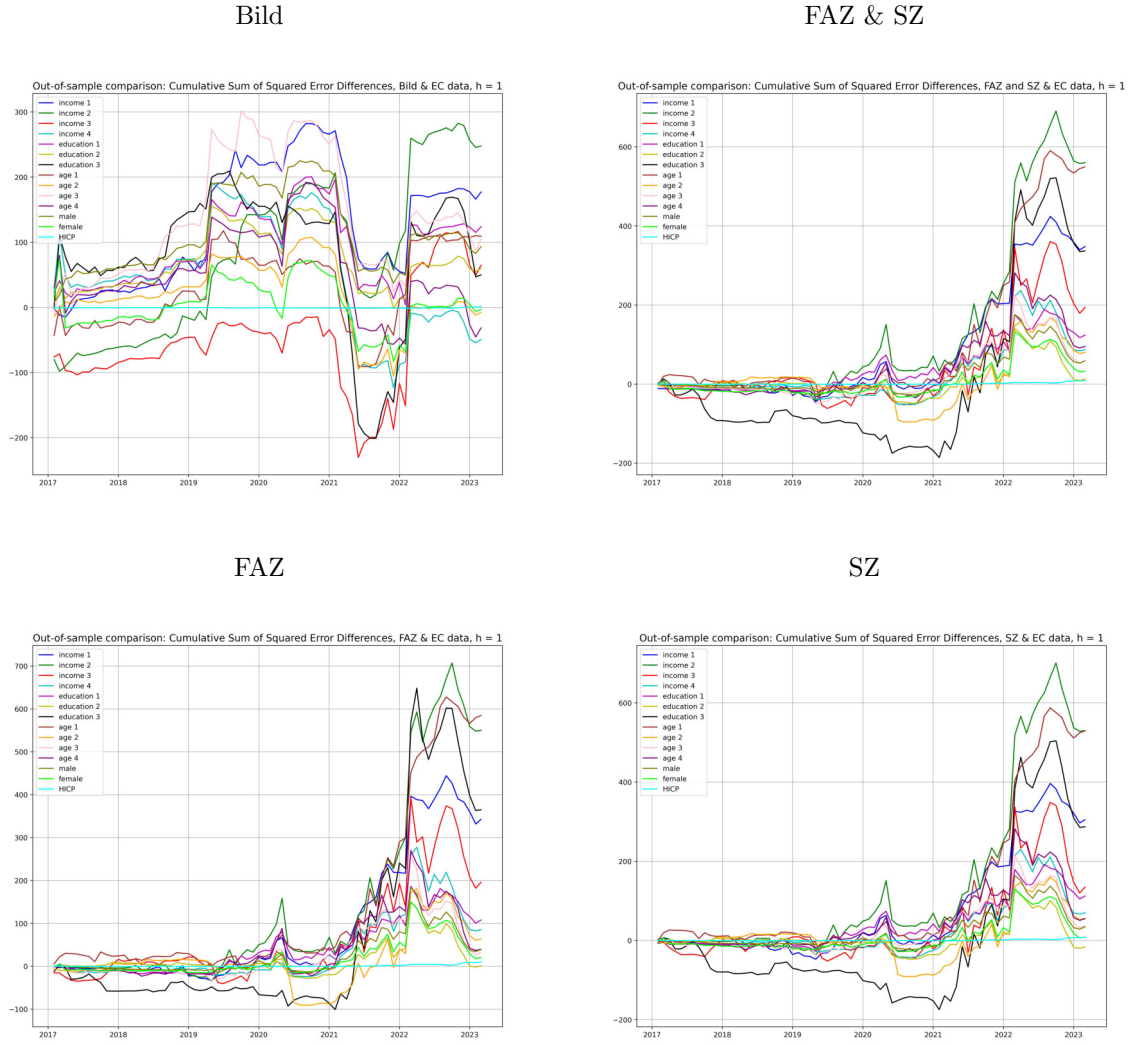
Bild	FAZ & SZ

FAZ	SZ

Figure 8: Cumulative sum of squared forecasting errors differences for the one month ahead out-of-sample forecasts of the averaged Bild, FAZ, SZ, and FAZ and SZ data combined newspaper sentiment index. If the CSSED is above 0, the competing model fares better than the benchmark $AR(p)$. The target variables are the inflation expectations of EC survey's sub-categories.

projections model as introduced by Jordà (2005):

$$y_{t+h} = \mu_h + \beta_h shock_t + \boldsymbol{\Delta}_h \sum_{l=0}^{n} W_{t-l} + u_{t+h}. \tag{3}$$

I estimate the model at daily frequency for each horizon h separately. The model includes a horizon-specific constant μ_h and the dependent variable y_t is in this case the log differences of one of the three newspaper sentiment indices. Specifically and similar to Born et al. (2023), I use the 7-day moving average of each newspaper's daily mean sentiment index to smooth out the days without any news reporting. The variable $shock_t$ is a shock series of ECB monetary policy surprises generated as in Jarociński and Karadi (2020) and provided online by Marek Jarociński[3]. For estimating euro area monetary policy shocks they rely on the *Euro Area Monetary Policy Event-Study Database* developed by Altavilla et al. (2019) which contains intra-day asset price changes around ECB
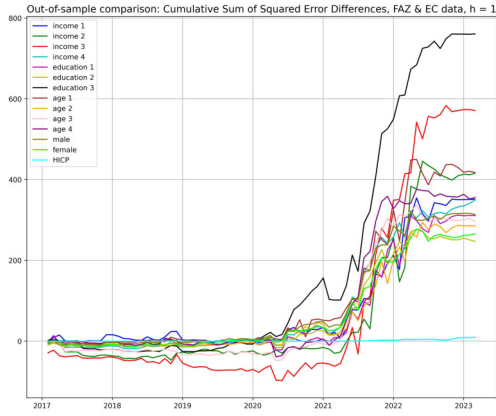
---

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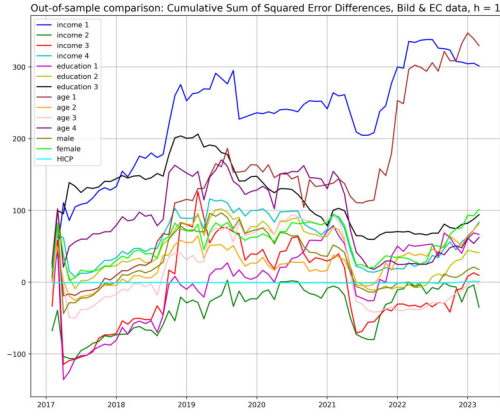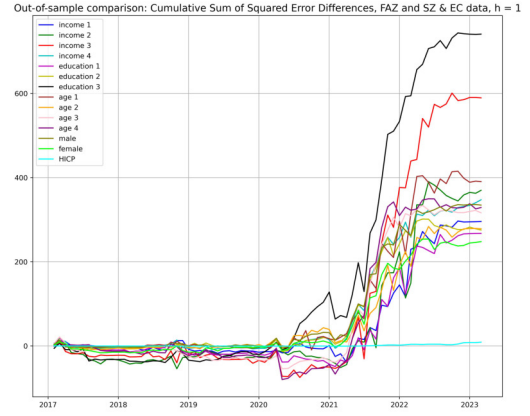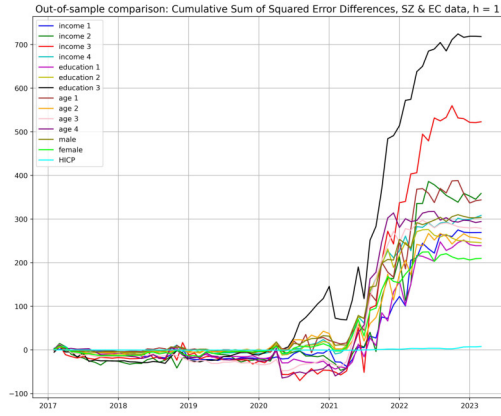
Bild

FAZ & SZ

FAZ

SZ

Figure 9: Cumulative sum of squared forecasting errors differences for the one month ahead out-of-sample forecasts of the averaged Bild, FAZ, SZ, and FAZ and SZ data combined newspaper sentiment index. If the CSSED is above 0, the competing model fares better than the benchmark $AR(p)$. The target variables are perceived inflation indices of EC survey's sub-categories.

policy announcements. The shock series is available in daily frequency, even though at maximum one monetary policy decision is taken per month. The controls W_t consist of the contemporaneous value and $n = 14$ lags (i.e., days) of the respective 7-day moving average newspaper sentiment index and German HICP. As inflation is only available at monthly frequency it enters the regression as a series that takes the same inflation value of a respective month for each day of that month. The model is estimated on a sample spanning from 2014M01 to 2023M06, due to the limited data availability of Bild articles at the beginning of the sample and the availability of the shock series. I follow Montiel Olea and Plagborg-Møller (2021) and compute Eicker-Huber-White (EHW) heteroskedasticity-robust standard errors which are sufficient when including lags of the dependent variable as controls. Finally, the impulse responses are given by the β_h coefficients for each horizon separately.
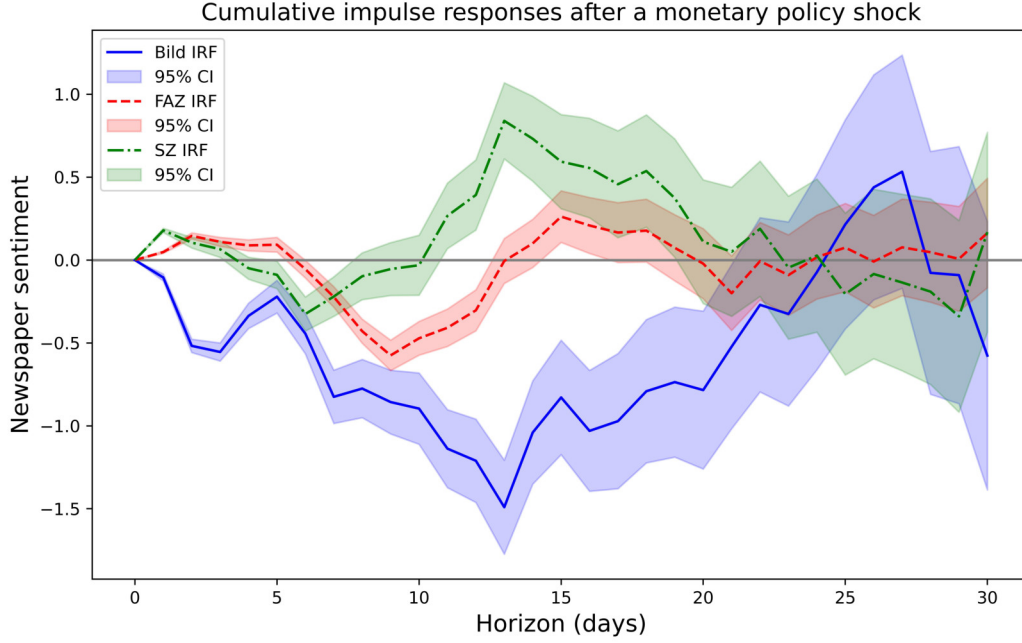
Figure 10: Impulse responses of the three newspaper sentiment indices (Bild in blue, solid line, FAZ in red, dashed line, and SZ in green, dotted-dashed line) to a contractionary monetary policy shock on a 30-days horizon. Shaded areas are 95% confidence intervals computed with EHW standard errors.

Figure 10 displays the cumulative impulse response functions over a horizon of 30 days. In the initial days following a contractionary monetary policy shock, the Bild sentiment immediately starts to decline, while the FAZ and SZ sentiments show significant decreases only after three to five days. The responses of SZ and FAZ appear quite similar, although the FAZ sentiment seems to lag behind SZ for a few days.

To examine these findings more closely, I restrict the sample to articles containing at least one of the German words for ECB, central bank, or key interest rate. This way, the articles considered are those that more explicitly address the ECB's policy. Figure 11 shows that in the initial periods after impact, the Bild sentiment again shows an immediate decrease, albeit slightly smaller than when considering all Bild articles, with an even stronger decrease after seven days. FAZ and SZ sentiments decrease after a few days' lag, and by 16 days post-impact, the effect has essentially dissipated, while the Bild sentiment reaction remains more persistent.

Overall, the signals towards inflation in the FAZ and SZ newspapers seem less strongly affected by a monetary policy surprise compared to those in the Bild newspaper. One possible reason for this could be the generally higher level and volatility of Bild sentiment compared to the other two newspapers. It seems that monetary policy effectively conveys its aim behind a contractionary shock, namely to decrease inflation.
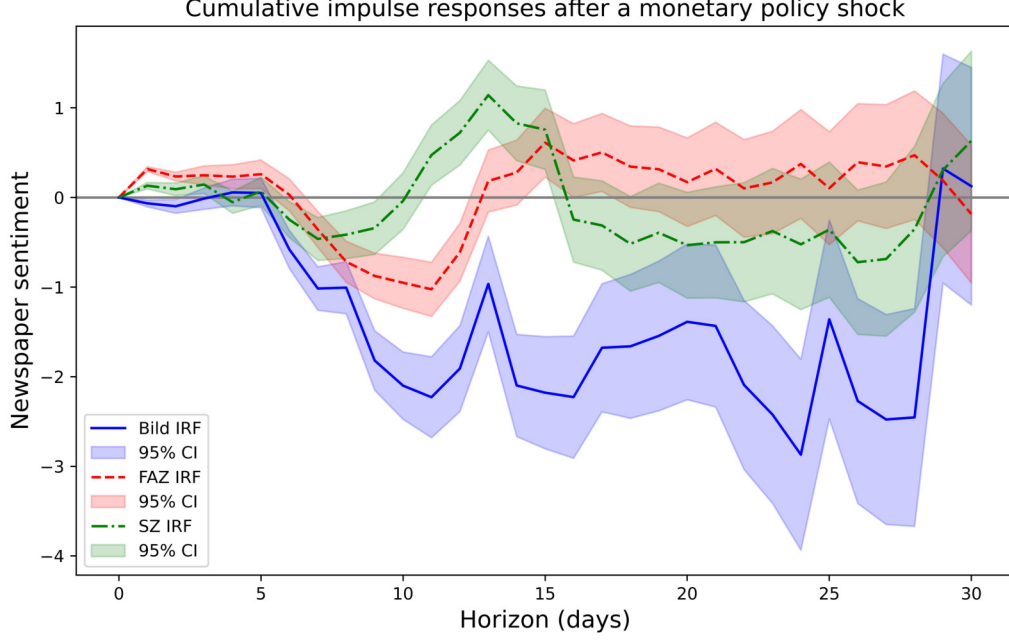
Figure 11: Impulse responses of the three newspaper sentiment indices (Bild in blue, solid line, FAZ in red, dashed line, and SZ in green, dotted-dashed line) to a contractionary monetary policy shock on a 30-days horizon. Shaded areas are 95% confidence intervals computed with EHW standard errors. Articles are restricted to those containing at least one of the words "EZB", "Zentralbank", or "Leitzins".

6 Household's reaction to news sentiment shocks

In a second step, I investigate whether the heterogeneous informativeness of the three newspaper indicators roots in their potential ability to affect household inflation expectations and perceptions differently. To this end, I closely follow Shapiro et al. (2022)'s modelling strategy using local projections, but make some adaptions to this specific use case. Shapiro et al. (2022) estimate the effect of a news sentiment shock on measures of economic activity as the component of the news sentiment index that is "*orthogonal to current and 4 lags of economic activity as well as 4 lags of itself.*" In contrast, the dependent variable in my model is one of the 13 sub-categories of the EC survey's household inflation expectation and perceived inflation indices, but I also estimate their impulse responses as their component which is orthogonal to the current value and $n = 4$ lags of the news sentiment index and indicators of economic activity, respectively. This way, I am able to capture that part of the household's reactions, that are not due to past news reporting on inflation or overall economic conditions. A separate regression is run for each forecast horizon $h = 1, ..., 6$ months:

$$y_{t+h}^i = \alpha_h^i + \beta_h^i NSI_t + \mathbf{\Gamma}_h^i \sum_{l=0}^{n} W_{t-l} + \sum_{l=1}^{n} \delta_{h,l}^i NSI_{t-l} + u_{t+h}^i, \qquad (4)$$

where y^i is one of the $i = 1, ..., 13$ survey measures, α_h^i a horizon specific constant, and NSI denotes one of the four newspaper sentiment indices (Bild, FAZ, SZ and FAZ & SZ

combined) as the average across all ways of aggregation excluding those that only take articles with sentiment scores greater than 4 into account. Both, the survey and sentiment measures are expressed in first differences. The set of controls W is also oriented towards Shapiro et al. (2022) and contains the contemporaneous value and 4 lags of the German industrial production index excluding energy and construction in log-differences (as a monthly proxy for output), year-on-year growth rates of German HICP, and first differences of a monthly measure of consumer confidence as a proxy for consumption. The first two series are obtained from the Federal Statistical Office of Germany while the latter can be retrieved from Eurostat and is compiled by the EC. The impulse responses are then given by estimates of the β_h^i coefficients which I report with 95% confidence bands constructed with Newey-West autocorrelation and heteroskedasticity robust standard errors. I estimate this model on 2014M01 to 2023M08 data with the availability of Bild sentiment scores being the limiting factor.

Figures 12 to 15 show the resulting cumulative impulse response functions to shocks in Bild, FAZ, and SZ (and the latter two combined) inflation sentiments, each figure representing one of the four EC survey sub-categories education and income. Results for the different age levels and gender can be found in Appendix F. In general, the heterogeneity in the responses is less pronounced for shocks to SZ and Bild sentiment and most significant for the FAZ. Nonetheless, there are some trends distinguishable. Figure 12 shows that FAZ and SZ have the strongest impact on the survey expectations of more educated households (education further than secondary school) while responses of education level 1 and 2 households react less and are mostly insignificantly different from each other. Impulse responses after a shock in Bild sentiment are not significantly different from each other and are very close to a zero response. The impulse responses of perceived inflation across different levels of education depicted in Figure 13 are very similar to those of inflation expectations.

When disaggregating among income quartiles (Figure 14), in the case of the Bild, the inflation expectations of second lowest income group are most responsive and significantly different from zero in the first four months after impact, while the response of expectations in the highest income group is even negative most of the time. This might reflect the circumstance that shocks in Bild sentiment do not necessarily reach wealthier households and explain some of the heterogeneity in household expectations. At the same time, shocks to FAZ and SZ sentiment induce the smallest response in the lowest income group (in the case of the SZ it is even mostly insignificant from zero), while the effects are most pronounced in the second and third income quartile. In the case of household's reaction in perceived inflation, a similar pattern emerges. Shocks to FAZ and SZ sentiment affect households of the third income quartile the most (Figure 15).
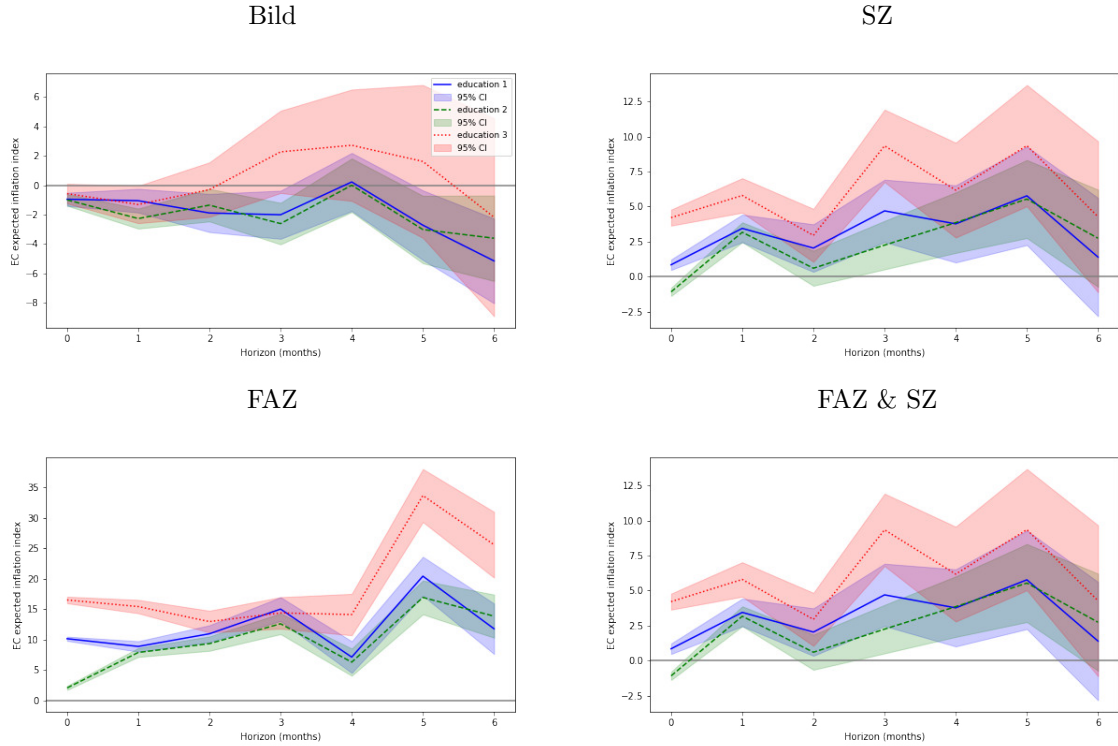
Figure 12: Cumulative impulse responses of EC survey expectations disaggregated across education levels to a shock in newspaper sentiment. Shaded areas represent 95% confidence bands constructed with Newey-West standard errors.

7 Robustness checks

7.1 Consistency of ChatGPT

Even though I set ChatGPT's model temperature to zero, which is intended to make its responses more deterministic, there is no guarantee that the model will always produce the same output given the exact same input prompt. However, this setting significantly increases the likelihood of consistency. To demonstrate that this does not affect the overall informativeness of the newspaper sentiment indices generated by this zero-shot approach, I had the model reevaluate the article's sentiments using a different prompt, closer to the one used in Bybee (2023). In addition to providing a sentiment score, I asked ChatGPT to assign a confidence probability between 0 and 1 to each score (see Figure 16. I then weighted the articles by this probability. Despite this additional weighting when aggregating the data, the final newspaper indices showed a sizable positive correlation with the baseline indices (Bild: 0.70, FAZ: 0.52, SZ: 0.46, average across all index aggregation methods). Without applying the probability weights, the correlations between the mean of all article sentiments per month were even higher: for Bild sentiments, the correlation was 0.82, and for FAZ and SZ articles combined, it was 0.78.
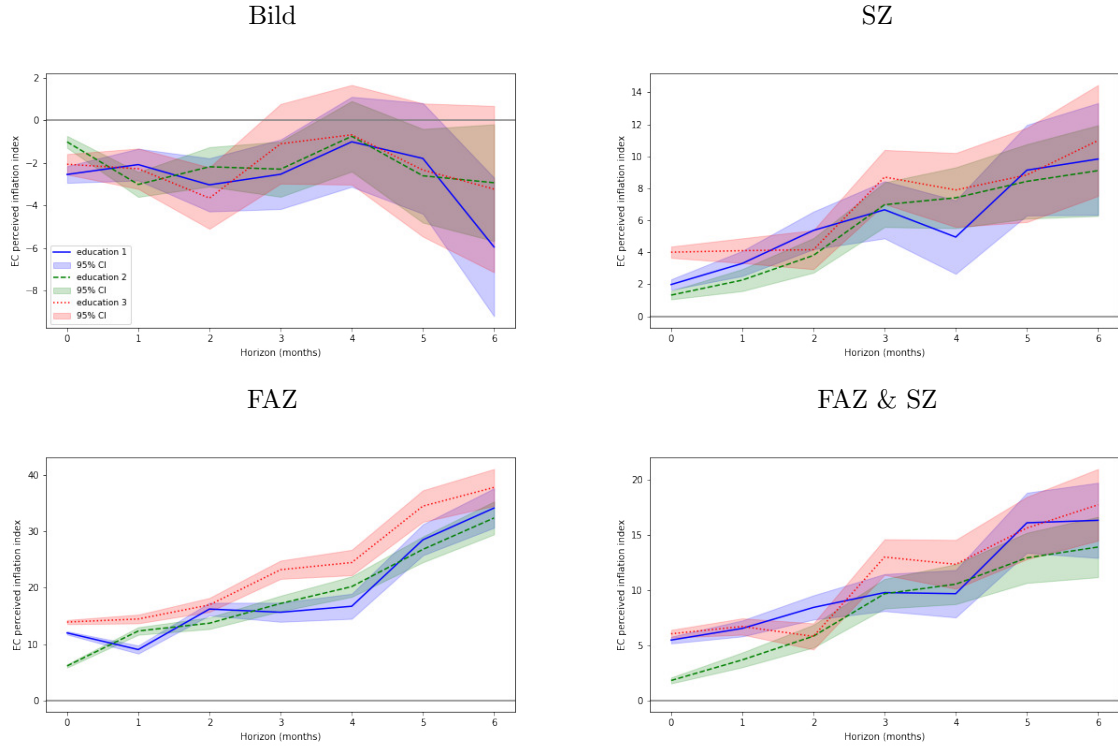
Figure 13: Cumulative impulse responses of perceived inflation disaggregated across education levels to a shock in newspaper sentiment. Shaded areas represent 95% confidence bands constructed with Newey-West standard errors.

7.2 Alternative survey expectations

As the EC survey data only provides information on the tendency whether households expect inflation to go up or down, I additionally use the ECB's *Consumer Expectations Survey* (CES), which runs since April 2020. On a monthly basis, the CES asks between 2.000 and 3.000 participants for their quantitative inflation expectations: *"How much higher (lower) do you think prices in general will be 12 months from now in the country you currently live in? Please give your best guess of the change in percentage terms. You can provide a number up to one decimal place."* The survey micro data is available online and provides background information on participant's age, gender, and education and income level. To correct for outliers I follow Bańkowska et al. (2021) and winsorize all quantitative observations at their 2nd and 98th percentile of their weighted distributions. In a second step, I compute weighted averages for each gender, age, education and income category.

To assess the informativeness of the newspaper sentiment indices in targeting quantitative inflation expectations measured by the CES, I conducted the same out-of-sample forecasting exercise as for the EC survey data, but with a forecast horizon of three months due to the short survey sample. Table 4 presents the results of this forecasting exercise. Almost all relative RMSEs are smaller than one, indicating information gains from including the newspaper indices. For Bild, there is minimal heterogeneity across household sub-categories. This contrasts with FAZ and, to some extent, SZ newspapers. Both FAZ and SZ perform best in predicting the expectations of the highest income and education
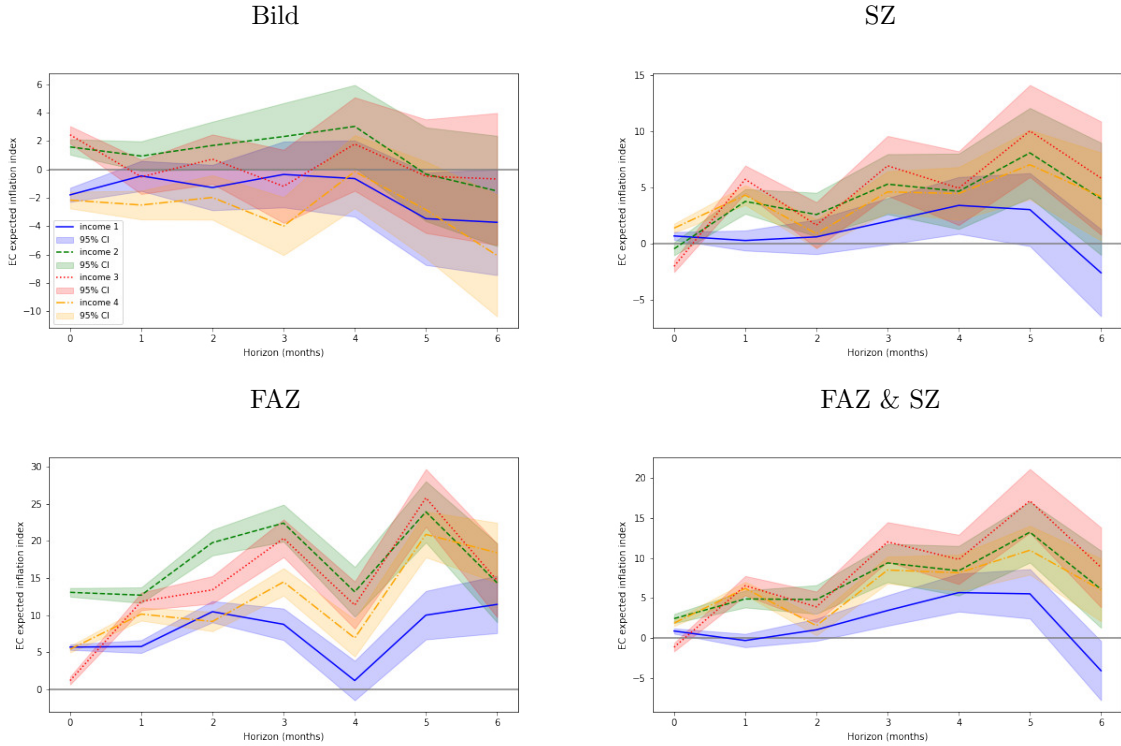
Figure 14: Cumulative impulse responses of EC survey expectations disaggregated across income quartiles to a shock in newspaper sentiment. Shaded areas represent 95% confidence bands constructed with Newey-West standard errors.

levels, as well as the age 35 to 49 cohort. Therefore, FAZ and SZ indices exhibit a similar heterogeneity pattern when compared to predicting qualitative EC survey expectations.

7.3   Sample splits

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

Table 6 presents the pre-inflation RMSE ratios on perceived inflation for the first sample split. For FAZ and SZ, the information gains appear to be primarily driven by the recent high inflation period, as the RMSEs are barely below 1 before April 2020 and decrease significantly thereafter (see Figure 6). In contrast, the Bild news sentiment shows the opposite trend. Here, the informativeness of the news indices is mainly driven by the pre-inflation period. The heterogeneity between household categories exhibits the same patterns as in the baseline: Bild indices best predict inflation perceptions for lower-income and lower-education households (with respect to significance levels). A similar pattern is observed when predicting inflation expectations (results can be found in Appendix G), where Bild RMSE ratios are even lower during the pre-inflation period.
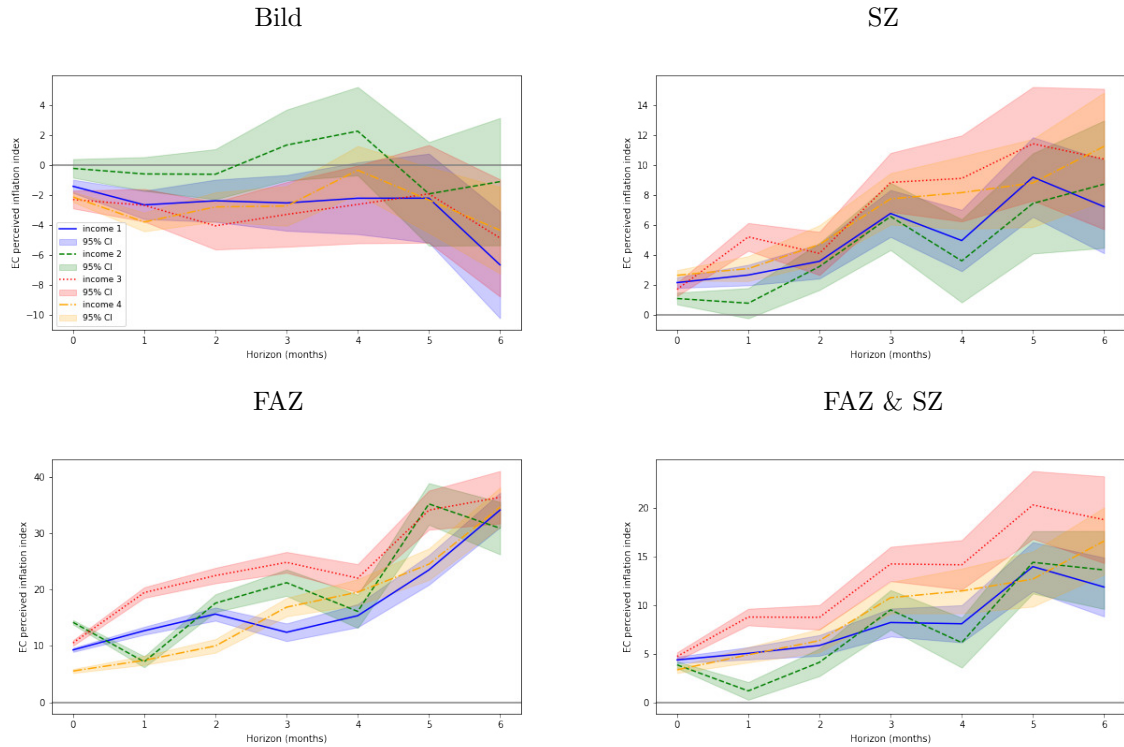
Figure 15: Cumulative impulse responses of perceived inflation disaggregated across income quartiles to a shock in newspaper sentiment. Shaded areas represent 95% confidence bands constructed with Newey-West standard errors.

```
System: You are an AI that rates inflation sentiment from 1 to 10.

User: Here is a newspaper article: "%title" + "%article"

Provide a sentiment score for inflation in this article on a scale of 1 to 10,
where 1 means inflation is likely to decrease, and 10 means inflation is
likely to increase. If the article does not address inflation in Germany or
the EU, respond with the word 'no'. Provide a probability between 0 and 1 how
confident you are about your assessment. Write your answer as:

sentiment score:
confidence:
```

Figure 16: Alternative prompt for generating inflation expectation sentiments with ChatGPT-3.5. "%title" and "%article" are place holders for single articles and their headlines, respectively.

News	age 1	age 2	age 3	age 4	income 1	income 2	income 3	income 4	income 5	edu 1	edu 2	edu 3	male	female
Bild	0.89***	0.92*	0.96	0.91*	1.12	0.97	0.97	0.91***	0.93*	0.96	0.93	0.90**	0.86***	0.99
	0.85**	0.92**	0.91**	0.90**	1.13	0.93	0.93	0.89***	0.92*	0.95	0.88**	0.89**	0.83**	0.94
	0.82**	0.88**	0.88**	0.90***	1.18	0.90**	0.92	0.86***	0.89*	0.94	0.85**	0.85**	0.79***	0.91*
FAZ	0.84***	0.77**	0.95	0.78**	0.86	0.85	0.90	0.83*	0.72**	0.86	0.83	0.73***	0.76**	0.86
	0.78***	0.73**	0.82**	0.75**	0.80	0.80*	0.83*	0.76***	0.65**	0.84**	0.76**	0.69***	0.70***	0.80**
	0.73***	0.67**	0.76**	0.73***	0.75	0.75**	0.79*	0.70***	0.56**	0.82**	0.71**	0.61***	0.63***	0.74**
SZ	0.88***	0.85*	0.97	0.88***	0.96	0.90	0.91	0.85***	0.81*	0.94	0.90	0.80***	0.82**	0.93
	0.84***	0.84**	0.88*	0.86***	0.93	0.87**	0.86**	0.83***	0.77*	0.89**	0.86**	0.78**	0.79***	0.88*
	0.81***	0.79**	0.82**	0.83***	0.91	0.84**	0.84**	0.80***	0.70*	0.88**	0.83**	0.73**	0.74***	0.85**
FAZ & SZ	0.88***	0.81**	0.92	0.86***	0.96	0.87*	0.89	0.82***	0.80**	0.91	0.86	0.77***	0.79***	0.88
	0.84***	0.79**	0.84**	0.83***	0.93*	0.84**	0.84**	0.79***	0.76**	0.88**	0.82**	0.76***	0.76***	0.84**
	0.81***	0.75**	0.78***	0.81***	0.91	0.80**	0.82**	0.75***	0.70*	0.86**	0.79**	0.71***	0.71***	0.81**

Table 4: Relative root mean squared errors of the averaged newspaper indicators of the FAZ, SZ, Bild and FAZ & SZ combined compared to ECB survey data. The symbols ***, **, and * represent significance levels of 1%, 5%, and 10%, respectively, in the Diebold–Mariano test for equal forecast accuracy.

News	income 1	income 2	income 3	income 4	edu 1	edu 2	edu 3	age 1	age 2	age 3	age 4	male	female	HICP
Bild	1.26	1.19	1.31	1.28	1.33	1.34	1.25	1.30	1.27	1.37	1.32	1.38	1.30	1.11
	1.42	1.23	1.43	1.50	1.51	1.48	1.49	1.36	1.45	1.49	1.33	1.65	1.39	1.03
	1.40	1.23	1.38	1.50	1.53	1.53	1.48	1.46	1.44	1.49	1.27	1.68	1.37	0.99*
FAZ	0.88***	0.87***	0.89***	0.90***	0.86***	0.95***	0.91***	0.90***	0.90***	0.95**	0.93***	0.89***	0.86***	0.92**
	0.84***	0.78***	0.79***	0.87***	0.83***	0.89***	0.90***	0.87***	0.85***	0.92***	0.88***	0.93***	0.77***	0.84***
	0.79***	0.70***	0.70***	0.74***	0.76***	0.74***	0.77***	0.68***	0.74***	0.80***	0.83***	0.86***	0.67***	0.76***
SZ	0.97***	0.94***	1.02	0.95***	0.95***	0.98**	0.99**	0.96***	1.00	0.97**	1.01	0.92***	0.98*	0.95**
	0.91***	0.85***	0.95**	0.84***	0.91***	0.85***	0.88***	0.97***	0.95**	0.83***	0.94***	0.79***	0.87***	0.92***
	0.88***	0.78***	0.83***	0.71***	0.87***	0.70***	0.77***	0.79***	0.86***	0.71***	0.88***	0.70***	0.77***	0.88***
FAZ & SZ	0.95***	0.93***	0.98**	0.93***	0.94***	0.96**	0.97**	0.95***	0.99**	0.95***	0.98**	0.90***	0.93**	0.94***
	0.87***	0.85***	0.87***	0.76***	0.89***	0.79***	0.84***	0.92***	0.89***	0.79***	0.91***	0.77***	0.79***	0.89***
	0.82***	0.78***	0.78***	0.63***	0.83***	0.62***	0.72***	0.74***	0.78***	0.67***	0.84***	0.69***	0.68***	0.84***

Table 5: Relative root mean squared errors of the averaged newspaper indicators of the FAZ, SZ, Bild and FAZ & SZ combined compared to EC survey data on perceived inflation. Sample length: 2020M04-2023M08. The symbols ***, **, and * represent significance levels of 1%, 5%, and 10%, respectively, in the Diebold–Mariano test for equal forecast accuracy.

News	income 1	income 2	income 3	income 4	edu 1	edu 2	edu 3	age 1	age 2	age 3	age 4	male	female	HICP
Bild	**0.90***	1.04	1.02	**0.98**	1.03	**0.98**	**0.94***	**0.94**	**0.99**	1.01	**0.96**	**0.98**	**0.98**	1.08
	0.91*	1.11	1.07	**0.99**	1.03	**1.00**	**0.91****	**0.89****	1.03	1.06	**0.94***	**0.97**	**0.99**	1.08
	0.91*	1.10	1.06	1.01	**0.97**	1.02	**0.93***	**0.87****	1.05	1.05	**0.92****	**0.95**	**0.97**	1.13
	0.89*	1.04	1.01	**0.98***	**0.94****	**0.97**	**0.94***	**0.82****	**1.00**	**0.97****	**0.90****	**0.93****	**0.92****	1.10
	0.87*	**0.97****	**0.94****	**0.91****	**0.91****	**0.89****	**0.93***	**0.78****	**0.95****	**0.88****	**0.87****	**0.89****	**0.86****	1.08
	0.85*	**0.94****	**0.87****	**0.85****	**0.84****	**0.83****	**0.94***	**0.78****	**0.91****	**0.84****	**0.83****	**0.86****	**0.81****	1.06
FAZ	1.01	1.02	1.04	1.02	1.01	1.02	1.02	1.00	1.02	1.04	1.01	**1.00**	1.00	1.03
	0.98	1.03	1.03	1.03	1.01	1.03	1.01	1.00	1.02	1.08	**1.00**	1.00	1.02	1.01
	0.96	1.04	1.01	1.05	1.02	1.02	1.01	1.00	**1.00**	1.15	**0.98**	0.99	1.03	1.00
	0.96	1.02	**0.99**	1.07	**1.00**	1.02	1.01	1.00	**0.98**	1.20	**0.97**	0.98	1.03	**0.98**
	0.94**	1.01	**0.97**	1.06	**0.98**	1.00	1.02	0.99	**0.94****	1.24	**0.95***	0.96	1.01	**0.95****
	0.92**	**0.99**	**0.91****	1.02	**0.98**	**0.97**	0.98	0.96	**0.90****	1.25	**0.94***	**0.92***	0.99	**0.94****
SZ	1.02	1.02	1.02	1.02	1.02	1.02	1.03	1.01	1.02	1.04	1.03	1.01	1.02	1.04
	1.02	1.03	1.01	1.04	1.03	1.03	1.06	1.01	1.02	1.08	1.03	1.00	1.05	1.05
	1.01	1.07	1.01	1.08	1.06	1.06	1.09	1.03	1.04	1.16	1.04	1.03	1.10	1.04
	1.01	1.07	**0.99**	1.12	1.05	1.09	1.11	1.03	1.04	1.22	1.03	1.03	1.14	1.03
	1.00	1.08	**0.97**	1.13	1.04	1.10	1.17	1.03	1.01	1.27	1.03	1.02	1.15	**0.98**
	0.99	1.07	**0.90***	1.11	1.04	1.07	1.15	1.01	**0.97***	1.28	1.04	**0.99**	1.14	**0.96***
FAZ & SZ	1.02	1.03	1.03	1.02	1.02	1.01	1.04	1.01	1.01	1.05	1.03	1.01	1.02	1.05
	1.00	1.04	1.01	1.05	1.04	1.03	1.05	1.01	1.03	1.10	1.03	1.00	1.06	1.07
	0.99	1.08	1.01	1.10	1.09	1.07	1.07	1.03	1.04	1.20	1.04	1.04	1.12	1.04
	1.00	1.08	**0.99**	1.14	1.07	1.10	1.08	1.04	1.04	1.28	1.03	1.03	1.16	1.03
	0.99	1.09	**0.97**	1.15	1.05	1.10	1.12	1.04	1.00	1.35	1.03	1.02	1.15	**0.97**
	0.97	1.06	**0.89****	1.11	1.05	1.06	1.09	1.00	**0.95***	1.34	1.02	**0.98**	1.13	**0.94***

Table 6: Relative root mean squared errors of the averaged newspaper indicators of the FAZ, SZ, Bild and FAZ & SZ combined compared to EC survey data on perceived inflation. The symbols ***, **, and * represent significance levels of 1%, 5%, and 10%, respectively, in the Diebold–Mariano test for equal forecast accuracy. Sample length: 2017M01-2020M03.

8   Conclusion

In this paper, I extract inflation sentiments at the article level from three major daily German newspapers—Bild, FAZ, and SZ—using the LLM ChatGPT-3.5 to construct inflation indices. I then analyze the informativeness of these indices regarding different socioeconomic groups' inflation expectations and perceptions to explore their role in influencing them. In terms of forecasting power, the Bild indices perform better in predicting expectations among lower education and lower income households, whereas the FAZ and SZ indices are more effective for higher income and more educated households. Additionally, FAZ and SZ sentiments are better suited to predict German HICP.

In a subsequent step, I run local projections to estimate the impulse responses of the three newspaper indices to a monetary policy shock. Again, heterogeneity in responses is evident: while Bild sentiment shows an immediate and sizable decrease—suggesting a less negative tone towards the inflation environment—the responses of FAZ and SZ indices are delayed by a few days and more muted. This may reflect Bild's volatile and mood-driven reporting style, as visible in the sentiment indices shown in Figure 5. On the other hand, it also indicates that monetary policy can effectively influence yellow press sentiment. The main policy implication from this observation is that central banks should focus more on targeting tabloids in their communication strategy to prevent unfounded negative reporting, as seen in the 2019 episode, and to leverage their policy impact on tabloids.

The final source of heterogeneity I explore is the varying responsiveness of different types of households to shocks in one of the three newspaper sentiment indices. Consistent with my previous findings, expectations and inflation perceptions of higher income and more educated households are most affected by shocks in FAZ and SZ sentiment. For Bild, this heterogeneity is less pronounced but still significant for participants belonging to the second income quartile.

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

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

References

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

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

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

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

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

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

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

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

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

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

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

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

Informationsgemeinschaft zur Feststellung der Verbreitung von Werbeträgern e. V. (IVW) (2024). Ivw digital. *https://www.ivw.de/*.

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

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

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

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

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

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

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

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

Weber, M., Candia, B., Ropele, T., Lluberas, R., Frache, S., Meyer, B. H., Kumar, S., Gorodnichenko, Y., Georgarakos, D., Coibion, O., et al. (2023). Tell me something i don't already know: Learning in low and high-inflation settings. Technical report, National Bureau of Economic Research.

Appendix

A Search queries

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

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

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

LexisNexis:
Zins-Hammer or steigende Preise or fallende Preise or Preischaos or Bezinpreis* or Teuerschock or Inflation* or niedrige Preise or hohe Preise or Disinflation or Deflation or Geld-

politik or Lohn-Preis-Spirale or Zentralbank or Stagflation or Energiepreis* or Geldsteuerung or EZB or Kaufkraftverlust or Notenb* or Leitzins or Kostenanstieg or Preissenkung* or Zinswende or Preisexplosion or Preisanstieg or Lebenshaltungskosten or Teuerung or Preisindex and not Verlag* and not Preisträger* and not Preissieger* and not *Film* and not Oscar and not Preisverleihung* and not Preisgeld* and not Billigung and not Trainer* and not Autor* and not Fußball* and not preisge* and not Trump

mediacloud:

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

B   ChatGPT's interpretation of Bild expressions

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

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

Term	ChatGPT Explanation	Interpretation of tone
Madame Inflation	In the context of the euro zone, "Madame Inflation" could potentially refer to Christine Lagarde, the current President of the European Central Bank (ECB). Christine Lagarde holds a significant role in shaping and implementing monetary policies that influence economic conditions, including inflation rates, within the euro zone.	When using a term like "Madame Inflation" to refer to Christine Lagarde, it suggests a personification of responsibility for inflationary pressures or challenges within the euro zone. The term implies a critical or at least skeptical stance toward her policies or actions, holding her accountable for the economic conditions that may lead to inflation. The use of "Madame" adds a formal and somewhat sardonic tone, suggesting a level of authority and influence associated with Christine Lagarde's position. Overall, it implies a critical perspective on the individual in charge of monetary policies in the euro zone, particularly in the context of concerns or shocks related to inflation.
Luxus-Lagarde	In the context of the newspaper headline "Madame Inflation; Luxus-Lagarde macht Sparer und Rentner arm," the term "Luxus-Lagarde" is a metaphorical expression that combines "Luxus" (luxury) with the surname "Lagarde," referring to Christine Lagarde, the President of the European Central Bank (ECB).	The term "Luxus-Lagarde" suggests a critical tone, implying that Christine Lagarde is associated with a lifestyle of luxury or extravagance. By combining her name with "Luxus," it conveys the idea that Lagarde is perceived as leading a lifestyle that may be considered opulent or out of touch with the financial challenges faced by the general population, specifically savers and retirees mentioned in the headline. The tone of "Luxus-Lagarde" in this context is likely critical, implying a disapproval of Lagarde's perceived lifestyle in connection with the economic challenges faced by savers and retirees. It may be used as a form of criticism, suggesting a disconnect between Lagarde's perceived luxurious lifestyle and the economic struggles of certain groups in society. This kind of expression is often employed in headlines or articles to convey a strong opinion or perspective on a public figure's actions or policies.

Term	ChatGPT Explanation	Interpretation of tone
Zins-Hammer	In the context of the newspaper headline "Zins-Hammer! EZB erhöht auf 1,25 Prozent" (Interest Rate Hammer! ECB Increases to 1.25 Percent), the term "Zins-Hammer" is a metaphorical expression. It combines "Zins" (interest rate) with "Hammer," suggesting a forceful and impactful event related to interest rates.	The term "Zins-Hammer" implies a tone of surprise, emphasis, and potentially a sense of urgency or significance. It suggests that the increase in interest rates is not just a routine adjustment but a substantial and noteworthy event that is likely to have important consequences or effects. The use of "Hammer" adds a dramatic and attention-grabbing element to the headline, indicating that the interest rate change is being portrayed as a powerful or forceful action. It might also imply that the change in interest rates could have considerable impacts on various stakeholders, such as borrowers, savers, and the overall economy.

Table B-1: Explanations and interpretations of Bild-typical expressions shown in Table 1 generated by ChatGPT in response to the following prompt: *"Explain the term [...] in the context of the newspaper headline [...] and describe the tone this expression implies."*

C   Matching readership data with EC survey bins

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

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

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

D Further correlation heat maps

This section presents sample correlation heat maps of the SZ expectation indices and the EC and ECB surveys (Figure D-1), as well as different sample splits for the EC survey data (2014M01-2020M03 and 2020M04-2023M08) and the three newspapers in Figure D-2. Overall the correlation in the latter sample split is higher. The pattern of heterogeneity between different groups of participants remains consistent for the Bild indices, while it somewhat reverses for the FAZ and SZ indices.
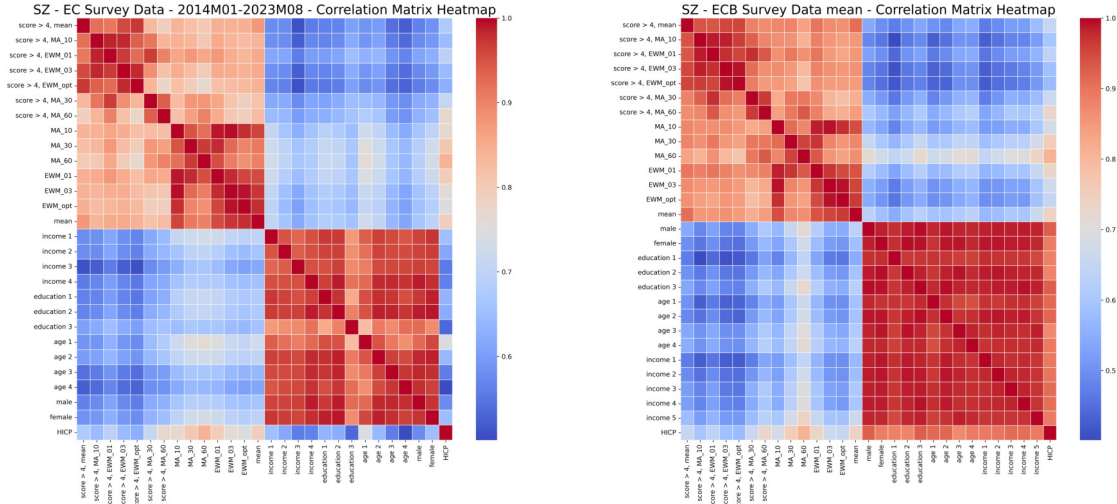
Figure D-1: Sample correlations between SZ inflation indices and EC data (left) and ECB survey data (right) based on the whole sample period; 2014M01-2023M08 for EC data and 2020M04-2023M08 for ECB data. Red indicates high and blue low correlations.
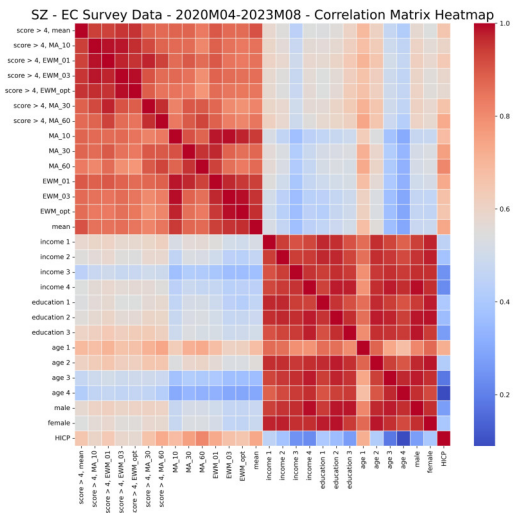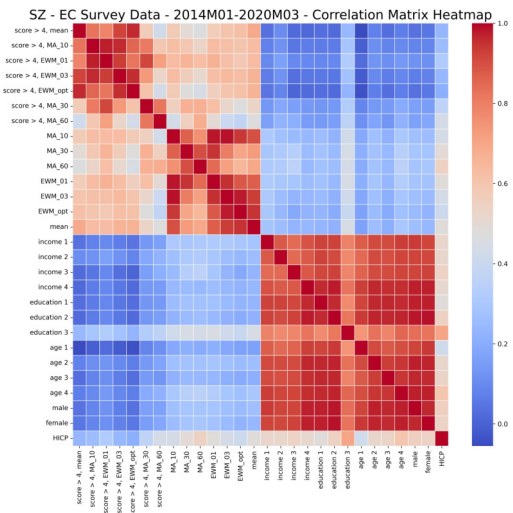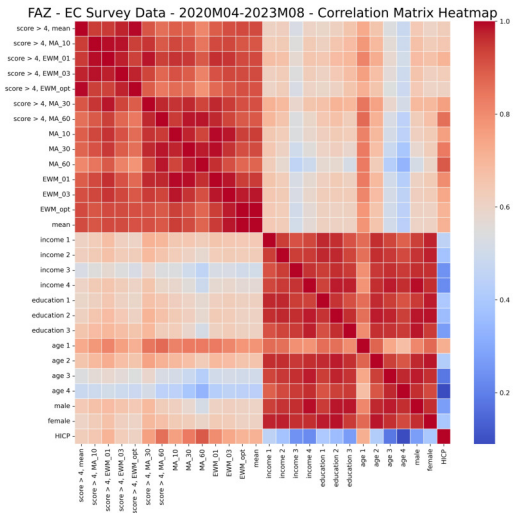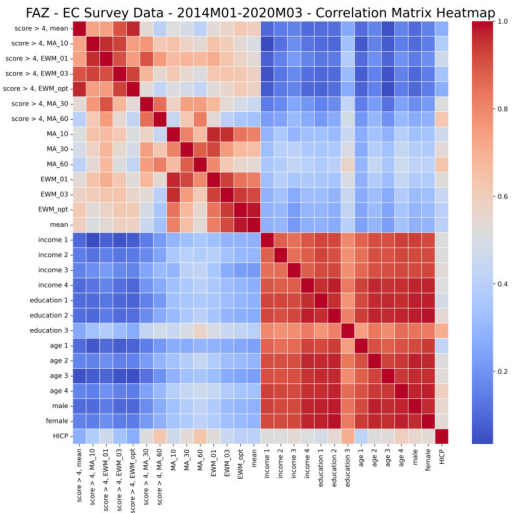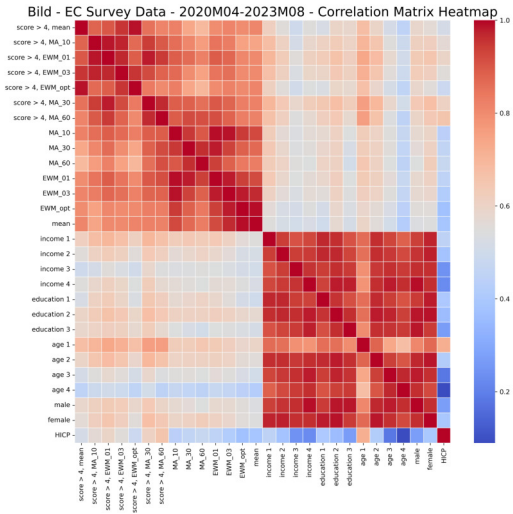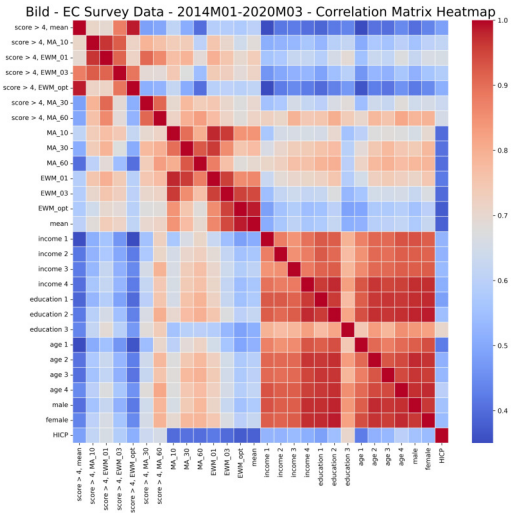
Figure D-2: Sample correlations between Bild, FAZ and SZ (top to bottom) inflation indices and EC data on the 2014M01-2020M03 (left) and the 2020M04-2023M08 sample period (right). Red indicates high and blue low correlations.

E   CSSED plots of horizons $h > 1$

This section contains the CSSED graphs of forecasting horizons larger than 1. Figure E-3 shows the results when predicting the ECB survey data and Figure E-4 when targeting the EC survey. At forecast horizons further in the future, the advantage of including the newspaper indices in a forecasting model becomes more evident as the CSSED increases over time.
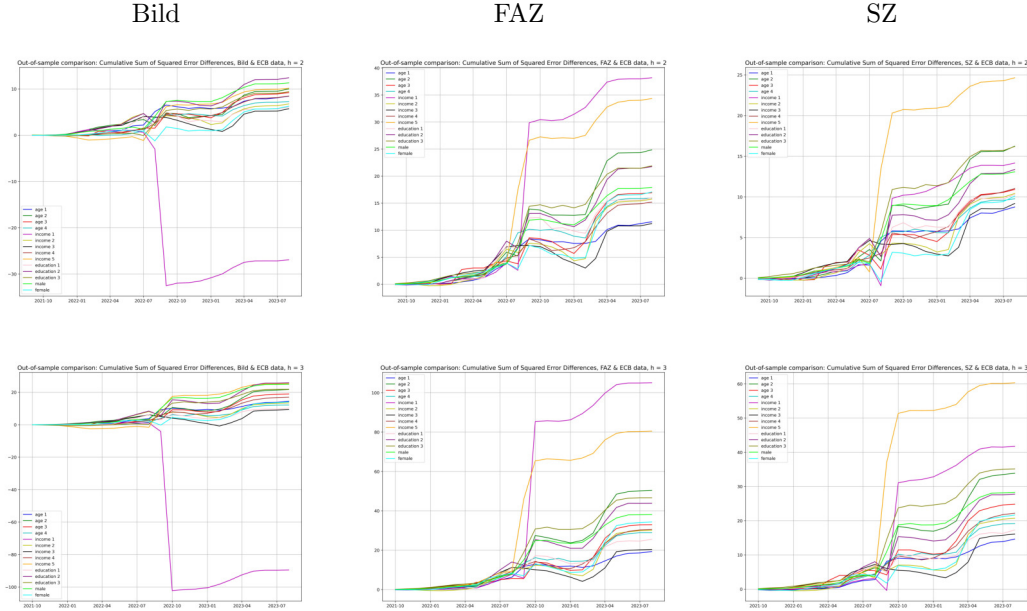

Figure E-3: Cumulative sum of squared forecasting errors differences for the two to six months ahead out-of-sample forecasts of the averaged Bild, SZ and FAZ newspaper expectation index. If the CSSED is above 0, the competing model fares better than the benchmark AR(p). The target variables are the ECB survey's sub-categories.

F   IRFs across gender and age after a news sentiment shock

There is no clear heterogeneity in responses of inflation expectations disaggregated across age after a shock to SZ and Bild sentiment (see Figure F-5), except for the FAZ, where the third age group (age 50 to 64) is the least responsive. Finally, Figure F-8 shows that shocks to different newspapers basically do not affect female or male survey participant's expectations differently. Only the overall effect differs which is largest for the FAZ and smallest or mostly insignificant for the Bild.

In the case of tha FAZ and SZ, perceived inflation of older survey participants are most affected by a news sentiment shock (Figure F-6). FAZ and SZ individually have a stronger effect on female's perceived inflation (Figure F-5).
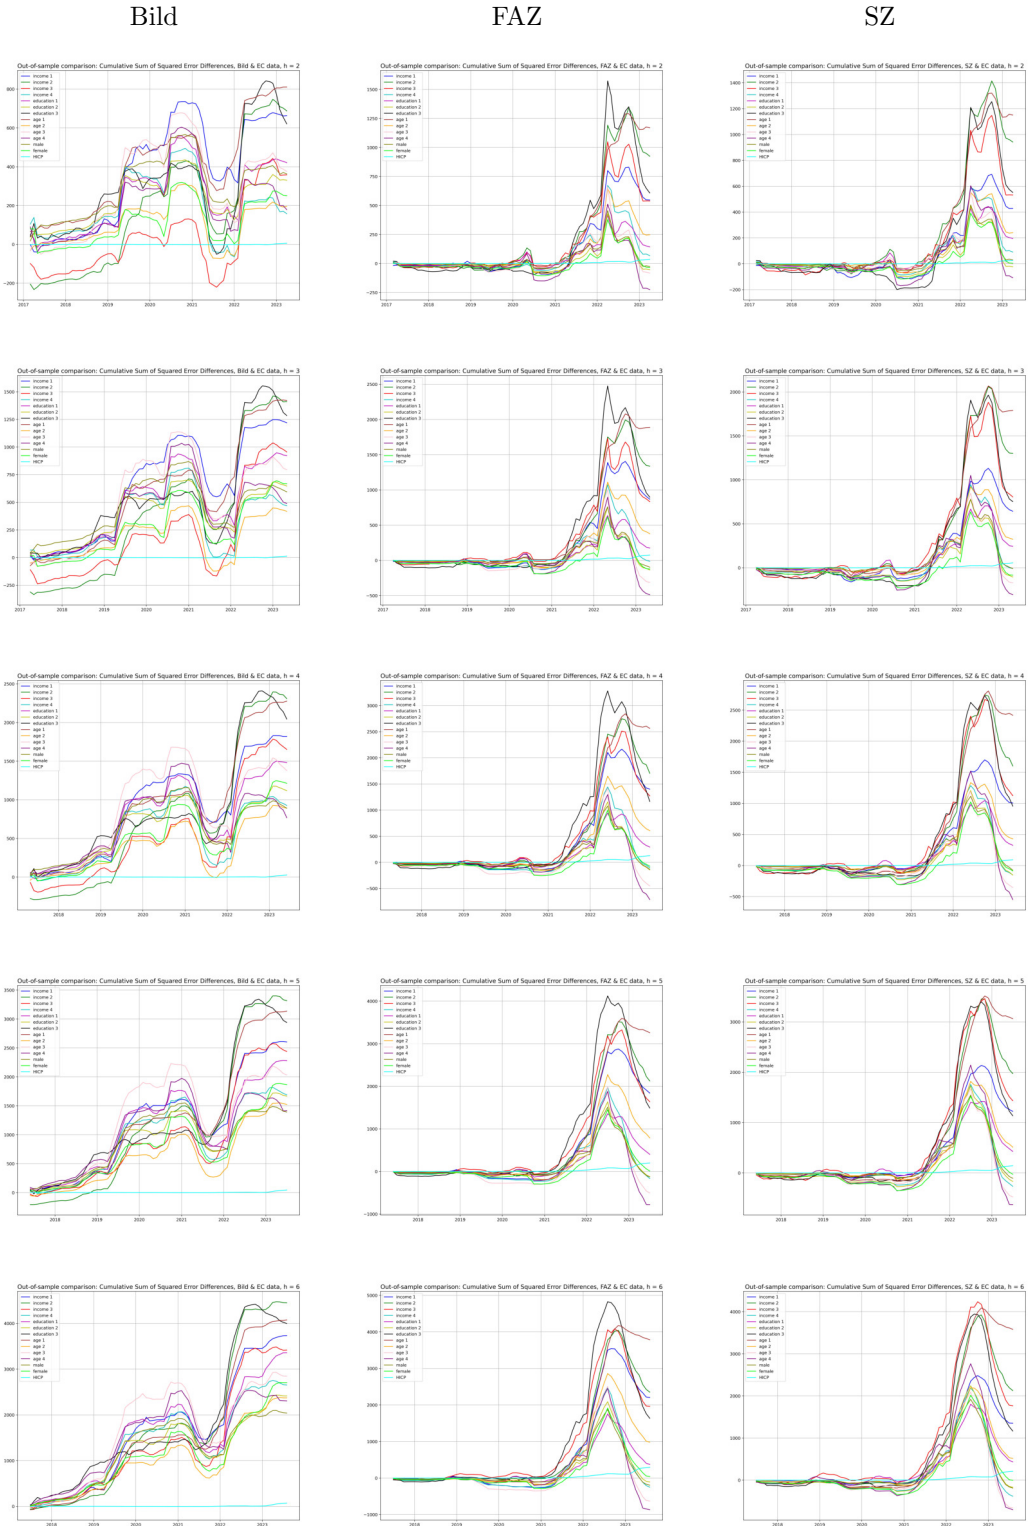
Figure E-4: Cumulative sum of squared forecasting errors differences for the two to six months ahead out-of-sample forecasts of the averaged Bild, SZ and FAZ newspaper index. If the CSSED is above 0, the competing model fares better than the benchmark AR(p). The target variables are the EC survey's sub-categories on inflation expectations.

Figure F-5: Cumulative impulse responses of EC survey expectations disaggregated across age groups to a shock in newspaper sentiment. Shaded areas represent 95% confidence bands constructed with Newey-West standard errors.
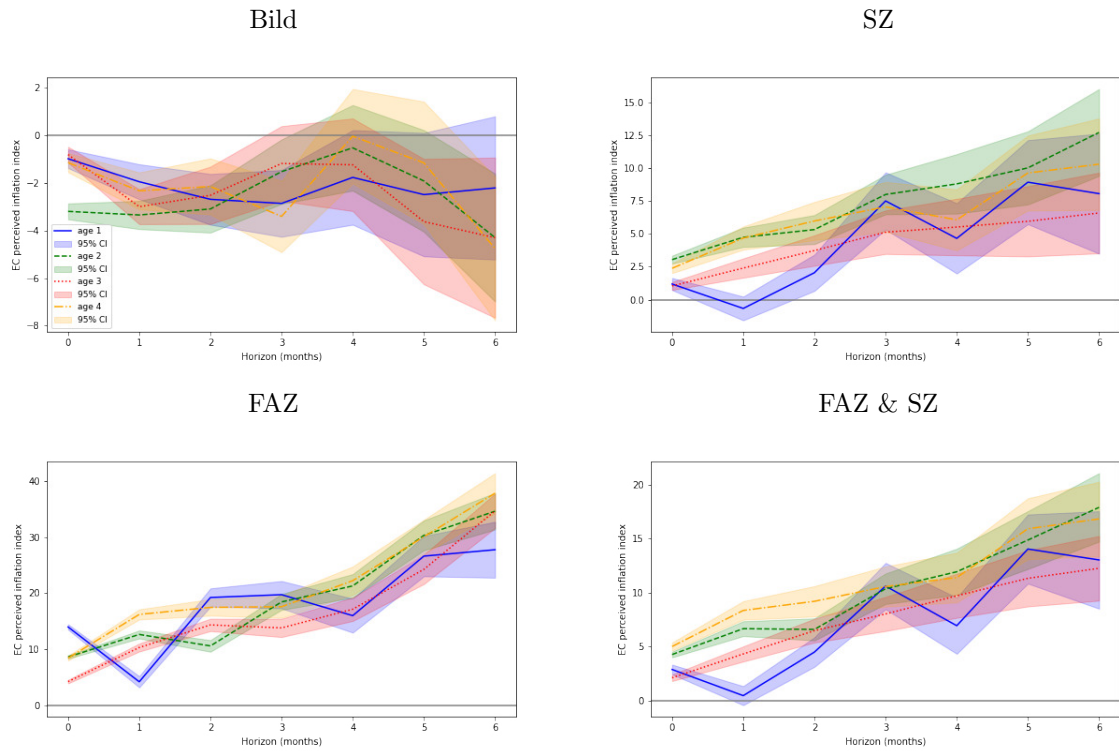
Figure F-6: Cumulative impulse responses of perceived inflation disaggregated across age groups to a shock in newspaper sentiment. Shaded areas represent 95% confidence bands constructed with Newey-West standard errors.
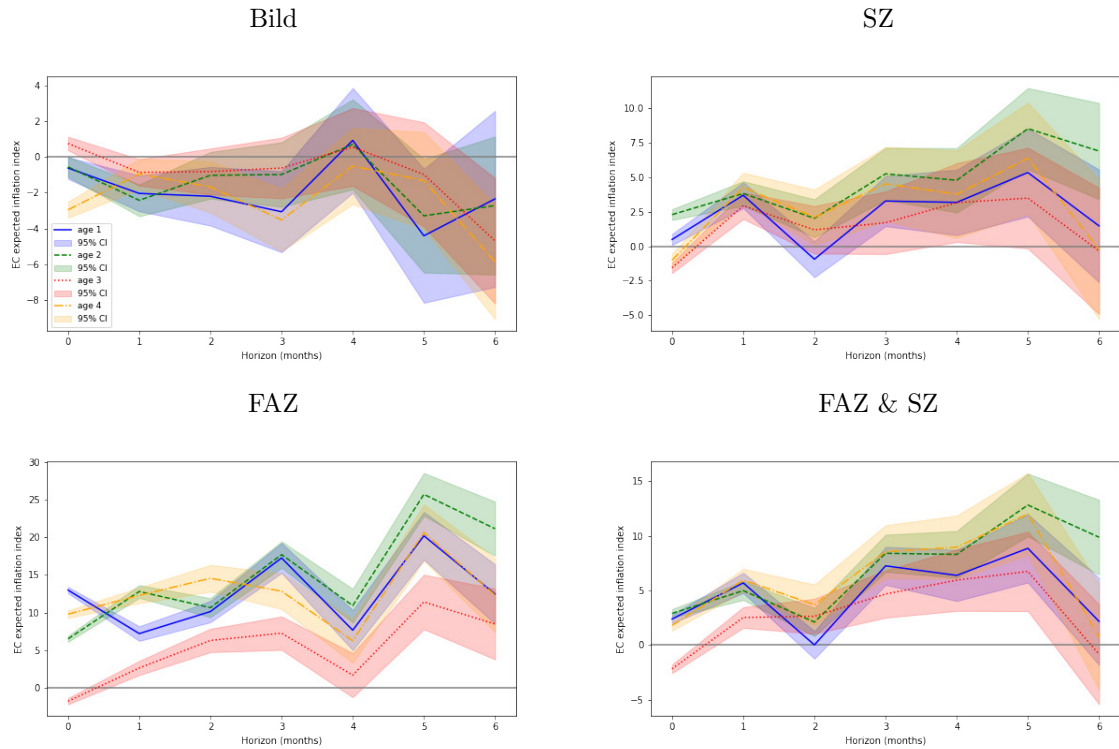
Figure F-7: Cumulative impulse responses of EC survey expectations disaggregated across gender to a shock in newspaper sentiment. Shaded areas represent 95% confidence bands constructed with Newey-West standard errors.

Figure F-8: Cumulative impulse responses of perceived inflation disaggregated across gender to a shock in newspaper sentiment. Shaded areas represent 95% confidence bands constructed with Newey-West standard errors.

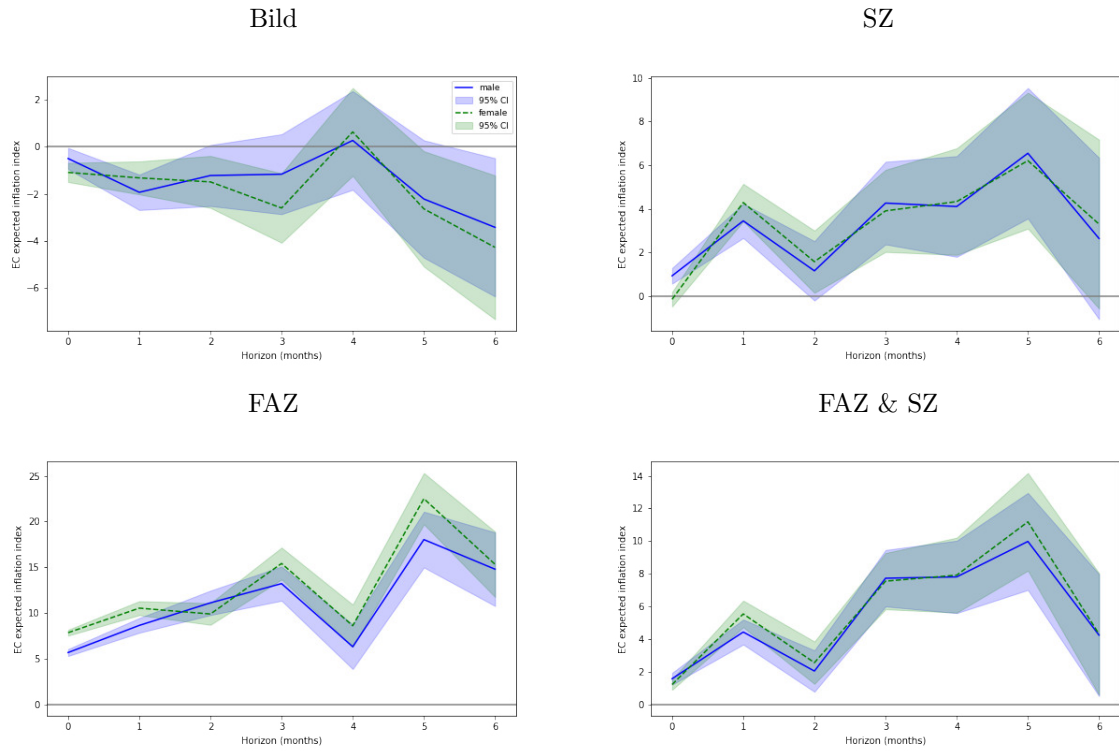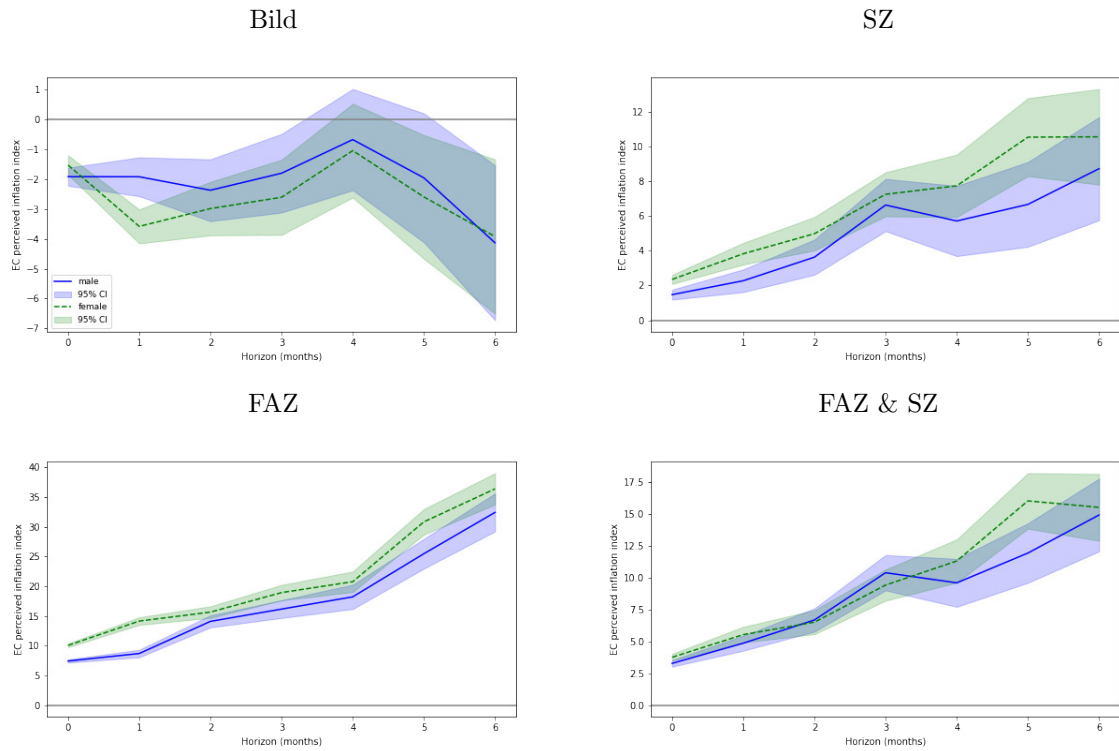# G RMSE results of sample splits: expected inflation

This section contains the RMSE ratios when predicting household's inflation expectations measured by the EC survey and reports the two sample splits: 2017M01 - 2020M03 (Table G-4) and 2020M04 - 2023M08 (Table G-3).

News	income 1	income 2	income 3	income 4	edu 1	edu 2	edu 3	age 1	age 2	age 3	age 4	male	female	HICP
	1.17	1.16	1.09	1.05	1.21	1.08	1.07	1.25	1.12	1.05	1.05	1.06	1.14	1.11
Bild	1.12	1.09	1.08	1.02	1.13	1.04	1.05	1.23	1.10	1.01	1.03	1.02	1.10	1.03
	1.09	1.05	1.03	**1.00**	1.07	1.02	1.02	1.16	1.07	**1.00**	1.01	1.01	1.06	0.99*
	1.00	1.02	1.11	1.10	**1.00**	1.08	1.07	**0.88****	1.06	1.08	1.08	1.09	1.04	0.92**
FAZ	0.95*	0.95*	1.08	1.07	0.95*	1.03	1.08	**0.73*****	1.03	1.05	1.09	1.07	**0.99**	0.84***
	0.89**	0.90***	**0.98**	1.03	0.89***	**0.96**	1.03	**0.65*****	0.94*	1.02	1.06	1.03	0.92**	0.76***
	0.99	1.02	1.08	1.12	**0.99**	1.07	1.08	**0.96**	1.06	1.06	1.07	1.08	1.04	0.95**
SZ	0.95*	**0.97**	1.05	1.07	0.95*	1.02	1.05	0.81***	1.03	1.02	1.06	1.05	**0.99**	0.92***
	0.91**	0.94**	**0.99**	1.03	0.92***	**0.97**	1.02	0.77***	**0.97**	1.01	1.04	1.02	0.95*	0.88***
	0.95**	**0.99**	1.04	1.08	**0.96**	1.03	1.05	0.92*	1.02	1.03	1.04	1.05	1.00	0.94***
FAZ & SZ	0.90***	0.93**	**0.99**	1.01	0.91***	0.95**	**0.99**	0.76***	0.95**	**0.98**	1.02	**0.99**	0.93**	0.89***
	0.85***	0.89***	0.93**	**0.97**	0.87***	0.90***	0.94**	0.70***	0.88***	0.96*	**0.99**	0.96*	0.88***	0.84***

Table G-3: Relative root mean squared errors of the averaged newspaper indicators of the FAZ, SZ, Bild and FAZ & SZ combined compared to EC survey expectation data on ECB survey data sample length. The symbols ***, **, and * represent significance levels of 1%, 5%, and 10%, respectively, in the Diebold–Mariano test for equal forecast accuracy.

News	income 1	income 2	income 3	income 4	edu 1	edu 2	edu 3	age 1	age 2	age 3	age 4	male	female	HICP
Bild	**0.91****	1.01	1.05	**0.94**	0.94*	0.92**	0.97*	**0.98**	0.96*	0.90**	**0.97**	0.90**	**1.00**	1.08
	0.89***	1.01	1.04	0.91*	0.93**	0.88***	0.95**	0.90***	0.93**	0.88**	0.96*	0.86***	**0.97**	1.08
	0.87***	**0.99**	1.03	0.90**	0.90***	0.87***	0.96**	0.90***	0.93**	0.85***	0.93***	0.85***	0.95**	1.13
	0.84***	0.95**	0.98**	0.86***	0.87***	0.84***	0.91***	0.84***	0.89***	0.80***	0.89***	0.83***	0.91***	1.10
	0.81***	0.90***	0.94***	0.82***	0.85***	0.82***	0.86***	0.81***	0.87***	0.78***	0.86***	0.80***	0.87***	1.08
	0.79***	0.86***	0.90***	0.78***	0.83***	0.79***	0.79***	0.80***	0.83***	0.75***	0.83***	0.78***	0.84***	1.06
FAZ	**1.00**	**1.00**	1.02	1.02	**1.00**	**1.00**	1.04	**0.99**	1.01	1.02	**1.00**	1.01	1.01	1.03
	1.02	1.01	1.00	1.02	1.01	1.02	1.04	1.02	1.01	1.03	1.01	1.01	1.03	1.01
	1.03	1.02	1.01	1.04	1.01	1.03	1.05	1.01	1.04	1.05	1.01	1.01	1.04	1.00
	1.04	1.03	1.01	1.05	1.02	1.04	1.05	1.03	1.04	1.06	1.02	1.02	1.04	**0.98**
	1.05	1.02	1.01	1.06	1.02	1.04	1.05	1.03	1.04	1.07	1.01	1.01	1.04	0.95**
	1.05	1.02	**0.99**	1.06	1.02	1.03	1.05	1.03	1.04	1.07	1.01	1.02	1.04	0.94**
SZ	**1.00**	**1.00**	1.03	1.02	**1.00**	1.01	1.05	**1.00**	1.01	1.02	1.00	1.01	1.02	1.04
	1.02	1.01	1.01	1.02	1.01	1.03	1.05	1.02	1.01	1.02	1.02	1.01	1.03	1.05
	1.03	1.02	1.03	1.05	1.01	1.06	1.08	1.03	1.03	1.03	1.04	1.01	1.06	1.04
	1.03	1.03	1.03	1.06	1.01	1.07	1.07	1.05	1.03	1.05	1.04	1.02	1.07	1.03
	1.02	1.02	1.03	1.06	1.01	1.07	1.07	1.05	1.03	1.04	1.04	1.02	1.07	**0.98**
	1.02	1.02	1.01	1.05	1.01	1.06	1.07	1.05	1.02	1.04	1.04	1.02	1.06	0.96*
FAZ & SZ	**1.00**	1.01	1.04	1.03	**1.00**	1.01	1.06	1.00	1.01	1.04	1.01	1.02	1.02	1.05
	1.03	1.02	1.02	1.04	1.01	1.04	1.07	1.03	1.03	1.04	1.03	1.02	1.05	1.07
	1.04	1.03	1.04	1.07	1.02	1.08	1.11	1.04	1.05	1.07	1.05	1.03	1.09	1.04
	1.05	1.04	1.04	1.09	1.03	1.09	1.11	1.07	1.06	1.08	1.06	1.05	1.09	1.03
	1.05	1.03	1.04	1.09	1.02	1.10	1.12	1.07	1.06	1.08	1.06	1.05	1.09	**0.97**
	1.06	1.03	1.03	1.09	1.02	1.08	1.13	1.08	1.05	1.08	1.06	1.05	1.08	0.94*

Table G-4: Relative root mean squared errors of the averaged newspaper indicators of the FAZ, SZ, Bild and FAZ & SZ combined compared to EC survey data on expected inflation. The symbols ***, **, and * represent significance levels of 1%, 5%, and 10%, respectively, in the Diebold–Mariano test for equal forecast accuracy. Sample length: 2017M01 to 2020M03.