

JUSTUS LIEBIG UNIVERSITY GIESSEN

Department of Economics and Business Studies

Doctoral Thesis

**Essays on Macroeconomic
Uncertainty, Expectations and
Monetary Policy**

1st Supervisor: Prof. Dr. Peter Tillmann

2nd Supervisor: Prof. Dr. Andreas Walter

Moritz Grebe

May 18, 2026

Preface

This dissertation consists of six essays that I worked on during my time at the University of Giessen and shortly afterwards, after joining the Deutsche Bundesbank. I ordered them not strictly chronologically, but rather chronologically within broader topics, since not all of them are closely related.

The first two essays study the effects of uncertainty surrounding the war in Ukraine. Specifically, Essay I, *Uncertainty about the War in Ukraine: Measurement and Effects on the German Economy*, co-authored with Sinem Kandemir and Peter Tillmann, constructs an uncertainty measure using more than eight million German-language tweets. We quantify uncertainty using text-based machine-learning methods and document its effects on daily financial and macroeconomic variables.

In Essay II, *Geopolitics in the Boardroom: How German Managers Respond to the War in Ukraine*, co-authored with the same team, we again use our Ukraine Uncertainty Index and study the responses of German managers. Using data from the German Business Panel, we document that managers respond strongly to uncertainty surrounding the war. They plan higher prices and expect declines in dividends, revenues, and profits after an uncertainty shock.

In Essay III, Salah Hassanin accompanied me on a journey into the metaverse. In the paper *Monetary Metanomics: The Effect of Monetary Policy on Virtual Real Estate Prices*, we create a price index from blockchain-based transactions of virtual land on the platform Decentraland. In a second step, we estimate the effects of the Federal Reserve's monetary policy. We show that prices in this virtual world respond strongly, but only briefly, to real-life monetary policy shocks.

Essays IV, V, and VI broadly investigate different influences on households' inflation beliefs.

In Essay IV, *Household Expectations and Dissent Among Policymakers*, Peter Tillmann and I study the effect of dissent in the ECB's Governing Council on inflation expectations. We conduct a randomized controlled trial in the Bundesbank Online Panel Households (BOP-HH), mimicking the communication of voting results of the Governing Council during ECB press conferences. This allows us to estimate the causal effect of this specific piece of information. We find that information about dissenting views especially affects the individual uncertainty of inflation expectations. In particular, households with anchored expectations become more uncertain when they learn that the Governing Council did not vote unanimously.

Essay V, *Inflation Expectations and a Conservative Central Banker: A Natural Experiment* (with Peter Tillmann), studies the effect of Jens Weidmann's resignation as President of the Deutsche Bundesbank on households' expectations. Weidmann is

a known policy hawk, and his resignation was unexpected. This allows us to study the effects of a conservative central banker in a natural experiment. Conveniently for our analysis, the resignation falls into wave 22 of the BOP-HH, allowing us to compare expectations before and after this event. We find that the announcement of the resignation increases uncertainty about inflation expectations and reduces trust in the ECB.

The final and most recent paper, Essay VI, *Prices at Prime Time: How TV News Shape Inflation Attention in Germany* (with Brenton Bruns), is about households' attention to inflation and the role of the *Tagesschau*. We measure *Tagesschau* coverage of inflation using data from the NDR archive and video material, and show that TV news plays an important role in shaping households' attention to inflation and their beliefs.

I would not have been able to finish even a single paper without the help of others. I want to thank them in the following.

First, I am incredibly thankful to my first supervisor, Prof. Dr. Peter Tillmann. I still remember the day Peter asked me whether I could imagine starting a doctorate. It was during one of our VWL5 BBQ sessions, and I was the one in charge of the barbecue. I was so surprised and happy that I completely forgot about the sausages, which resulted in them being badly burned. I would never have thought in that moment that, about six years later, I would be submitting this dissertation. I thank Peter for giving me this opportunity, for his incredible support, and for his patience throughout this time. I learned so much from him, and working with him is simply extremely fun. I also benefited greatly from his advice on papers on which he is not a co-author.

I also thank my second supervisor, Prof. Dr. Andreas Walter, for his support and his help on the job market. I hope he enjoys reading this dissertation, even though my most recent (and more finance-related) paper is not yet advanced enough to be included.

Before starting the doctorate, I was a student assistant at the chair VWL5. During this time, I learned a lot from Paul Rudel, Jörg Schmidt and Lucas Hafemann, who were doctoral students at the time. David Finck, also a doctoral student then, supervised my bachelor's thesis and opened the world of macroeconomic expectations to me, which I now find so fascinating. I thank him for becoming a fantastic colleague after I joined the chair as a research assistant, and I am happy that we still see each other often at the Bundesbank. He has so much expertise, and I benefited a lot from his local projection toolbox.

I started this journey together with Sinem Kandemir in the same month. I thank her

for being the best office buddy one could hope for and a great co-author. Together with Anisa Tiza Mimun, who joined us a few months later, we had so much fun at the chair. I thank both of them for creating such a great working atmosphere, from which I benefited very much. I further thank Mahmadou Lamine Barry and Jinyeong Yun for being such kind and supportive colleagues.

I thank all student assistants who were employed during my time at the chair for their help: Svea Lembcke, Marlene Alt, Svenja Küchler, Laura Eich, Oskar Langert, Laurenz Burggraf, Nicole Panz, Friederike Schimmel, Justin Berndt, Sara Joneidi, Murodjon Berdimurodov and Brenton Joey Bruns. I am very happy that Brenton became a co-author. The *Tagesschau* paper would not have been possible without his outstanding data science expertise.

Further, I am thankful to Cornelia Strack, the secretary of the chair, for all her patience and help with simply everything.

I also thank the last co-author I have not named yet: Salah Hassanin. I thank him for the idea of looking at the metaverse. It was so much fun exploring this world together and writing this paper.

I thank Tobias Schmidt and all my other new colleagues at the Bundesbank for welcoming me so warmly and for the flexibility that helped me a lot in the final phase of this dissertation.

Lastly, I thank my family: Sabine, Markus, Charlotte, Maria, Benjamin, Finn, Frida and Sem, for all they have done for me. They even sometimes provided feedback on the papers, or at least gave a thumbs-up in the family WhatsApp chat when I posted a new version there.

I dedicate this dissertation to my wife Jana. She always supports me and believes in me. I could not have done this without you! ♡

Contents

Essay I: Uncertainty about the War in Ukraine: Measurement and Effects on the German Economy	6
Essay II: Geopolitics in the Boardroom: How German Managers Respond to the War in Ukraine	49
Essay III: Monetary Metanomics: The Effect of Monetary Policy on Virtual Real Estate Prices	82
Essay IV: Household Expectations and Dissent Among Policymakers	132
Essay V: Inflation Expectations and a Conservative Central Banker: A Natural Experiment	183
Essay VI: Prices at Prime Time: How TV News Shape Inflation Attention in Germany	224
Affidavit	274
Submitted Papers	275

Essay I

Uncertainty about the War in Ukraine: Measurement and Effects on the German Economy

This paper has been published as:

Grebe, M., Kandemir, S. and Tillmann, P. (2024). Uncertainty about the war in Ukraine: Measurement and effects on the German economy, *Journal of Economic Behavior and Organization* **217**: 493-506. <https://doi.org/10.1016/j.jebo.2023.11.015>

This paper was presented at the following workshops and international conferences:

- i. Deutsche Bundesbank Research Seminar (Frankfurt)
Date: July 2023
Presenter: Peter Tillmann
- ii. FernUniversität Hagen
Date: August 2023
Presenter: Peter Tillmann
- iii. 6th Monash-Warwick-Zurich Text-As-Data (online)
Date: September 2023
Presenter: Sinem Kandemir
- iv. 3rd MORSE Conference (Maastricht)
Date: October 2023
Presenter: Sinem Kandemir
- v. Studying Central-Eastern European Politics in the Age of AI – A workshop on data and methods (Budapest)
Date: November 2023
Presenter: Sinem Kandemir

- vi. 17th South-Eastern European Economic Research Workshop (Tirana)
Date: December 2023
Presenter: Sinem Kandemir
- vii. 24th IWH-CIREQ-GW-BOKERI Macroeconometric Workshop: Commodity
Prices and Macroeconomic Developments (Halle)
Date: December 2023
Presenter: Moritz Grebe (Poster session)
- viii. International Conference of the School of Economics and Business in Sarajevo
(Sarajevo)
Date: October 2024
Presenter: Moritz Grebe

Acknowledgement

We are grateful to the editor of the JEBO and two anonymous reviewers for helpful comments and suggestions. We also thank David Finck, Salah Hassanin, Albina Latifi, Vivien Lewis, Viktoriia Naboka-Krell, Matthias Neuenkirch, Nicolas Pröllochs, Matthias Rottner and seminar participants at Deutsche Bundesbank, FernUniversität Hagen, the 6th Monash-Warwick-Zurich Text-As-Data Workshop, the 3rd MORSE Conference Maastricht and the University of Giessen for feedback on an earlier draft. Financial support from Fritz Thyssen Stiftung is gratefully acknowledged.

Uncertainty about the War in Ukraine: Measurement and Effects on the German Economy

Moritz Grebe* Sinem Kandemir[†] Peter Tillmann[‡]

Abstract

We assemble a data set of more than eight million German Twitter posts related to the war in Ukraine. Based on state-of-the-art methods of text analysis, we construct a daily index of uncertainty about the war as perceived by German Twitter. We then estimate a VAR model with daily financial and macroeconomic data and identify an exogenous uncertainty shock. The increase in uncertainty has strong effects on financial markets and causes a significant decline in economic activity as well as an increase in expected inflation. We find the effects of uncertainty to be particularly strong in the first months of the war.

Keywords: war, Twitter, geopolitical risk, machine learning, business cycle

JEL Classification: D8, E3, G1

*Justus Liebig University Giessen, Email: moritz.grebe@wirtschaft.uni-giessen.de

[†]Justus Liebig University Giessen, Email: sinem.kandemir@wirtschaft.uni-giessen.de

[‡]Justus Liebig University Giessen, Email: peter.tillmann@wirtschaft.uni-giessen.de

1 Introduction

The Russian attack on Ukraine in February 2022 was a huge blow to the world economy. As a consequence of the war, many countries in Europe and beyond have to adjust their military and security policies in light of new geopolitical realities. In addition, countries in Europe, most notably Germany, have to revise their energy policy and, with an eye on future geopolitical conflicts, their economic interdependencies with other countries. These adjustments entail large macroeconomic costs (Garicano et al., 2022).

One key economic dimension of the Russian attack on Ukraine is uncertainty: households and firms are uncertain about a potential escalation of the war, the duration of the conflict and the post-war order. Fluctuations in this uncertainty should be a driver of financial markets, real economic activity and inflation expectations.

In principle, uncertainty can impact the economy through three main channels. First, irreversible investment in light of fixed costs implies that uncertainty prompts firms to rethink their investment plans and wait until the uncertainty is resolved. Second, risk-averse households respond to an increase in uncertainty by reducing consumption expenditure. Third, uncertainty increases risk premia, which makes investment and consumption more expensive. All three channels imply that an increase in uncertainty has contractionary effects on the real economy.

In this paper, we study the consequences of the uncertainty about the war in Ukraine on the Germany economy. There are two main obstacles to this analysis. First, we need to measure the uncertainty perceived by households and firms about the evolution of the war. Second, we need to trace the causal effect of a change in uncertainty on the economy. We make two contributions that address these problems. As a first contribution, we assemble a data set of more than eight million German-speaking tweets, i.e. short messages on Twitter.¹ These tweets are drawn from the Twitter API based on a list of specific keywords, i.e. "Ukraine", "Putin" etc. Twitter is a particularly useful source of data on perceived uncertainty because users can express their views as a response to geopolitical events and can do so in real-time. Both aspects are clear advantages compared to newspaper-based measures of uncertainty. The drawback is that the content of Twitter posts is probably more noisy than curated newspaper reporting.²

We then apply a state-of-the-art technique from the field of Natural Language Processing (NLP) to explore the public perception of the war. In this regard, we

¹Twitter has been re-branded to X in 2023. Nevertheless, we use the old and well-established name of the service.

²Twitter is also affected by Russian propaganda as investigated by Geissler et al. (2023).

use zero-shot classification with German Bidirectional Encoder Representations from Transformers (GBERT), i.e. a transformer-based language model. For each tweet, we obtain a probability of being classified as expressing uncertainty. In contrast to dictionary-based uncertainty indicators such as Baker et al. (2016) and others, this method does not rely on identifying words in a given tweet that are included in a pre-specified dictionary. We then compute the daily averages of the uncertainty probabilities. As a result, we obtain a daily indicator of uncertainty, which clearly reflects the evolution of the war and peaks on February 24, 2022, the day of the attack.

The second contribution of this paper is an identified vector autoregression (VAR) model that allows us to estimate the causal effect of uncertainty shocks on the German economy. We include daily variables such as stock prices, the volatility of the stock market, the price of natural gas, a daily measure of real economic activity and break-even inflation expectations from the bond market for the period January 2022 to February 2023. As a matter of fact, uncertainty is an endogenous variable and responds to the state of the economy and the financial market. To achieve identification, we exploit the fact that Twitter is active around the clock, while the stock market is subject to trading hours and most economic activity takes place during the day. Specifically, we separate uncertainty expressed on Twitter during the night from uncertainty during the day and order uncertainty at night first in the VAR model. A recursive ordering of the variables implies that on a given day, uncertainty at night, i.e. before the opening of the stock market, can drive financial markets and economic activity, while the reverse effect takes place after at least one day. In contrast to other papers in the literature relying on a recursive identification of uncertainty shocks, e.g. Bachmann et al. (2013), Baker et al. (2016), Altig et al. (2020) and Jurado et al. (2015), the recursiveness assumptions hold by construction and allow us to obtain a clean identification. We also report results from an alternative identification, i.e. a "hybrid VAR" following Romer and Romer (2004), Coibion (2012) and Caldara and Herbst (2019), among others.

We find that an uncertainty shock has strong adverse effects: stock prices fall, stock market volatility increases, the price of natural gas increases, real economic activity falls and expected inflation rises. All these responses are significant and have economically relevant magnitudes. An increase of uncertainty of one standard deviation, for example, depresses the equity market valuation by one percent and raises gas prices by two percent. The opposite responses of activity and expected inflation suggest that the effect of uncertainty resembles the effect of an adverse supply shock. A subsample analysis reveals that the effect of uncertainty is much

larger in the first subsample, i.e. before August 2022, compared to the second subsample. Overall, our results suggest that geopolitical risks as reflected in the uncertainty expressed in tweets about the war in Ukraine, have strong and significant adverse effects on the German economy.

This paper is contributing to the literature on the macroeconomic effects of economic uncertainty, e.g. Bachmann et al. (2013), Baker et al. (2016), Jurado et al. (2015), Altig et al. (2020) and Ludvigson et al. (2021). Most of the papers use identified VAR models and find negative effects of uncertainty shocks on output. Bloom (2014) and Castelnuovo (2023) survey the vast literature on uncertainty shocks. In line with the notion of Bloom (2014), we use the concepts of uncertainty and risk synonymously. The work by Baker et al. (2016) is particularly relevant as these authors establish the construction of uncertainty indicators from text data using pre-specified dictionaries. In contrast to their paper, we use high-frequency Twitter data and apply a machine learning approach.³ We follow Bachmann et al. (2013), Baker et al. (2016), Altig et al. (2020) and Jurado et al. (2015) and use a recursive ordering of the endogenous variables to identify an exogenous uncertainty shock. Importantly, in our model, the recursive ordering holds by the construction of the variables. Hence, our paper avoids the critique of Kilian et al. (2025) of recursively identified uncertainty shocks. Recent literature sheds light on the economic effects of fluctuations in geopolitical risk. The pioneering contribution of Caldara and Iacoviello (2022) is the construction of the Geopolitical Risk Index. The authors search a century of newspaper articles for keywords associated with geopolitical threats, wars, terrorism and other tensions. A recursively identified VAR shows that shocks to geopolitical risk cause a persistent fall in investment, employment and equity prices. Caldara et al. (2022) update the Geopolitical Risk Index to include the war in Ukraine. Based on an estimated VAR model, they show that an increase in risk causes a strong fall in global economic activity and an increase in global inflation. The opposite responses of activity and inflation suggests that the geopolitical risk shock resembles a supply-side shock, which is consistent with our findings. Bondarenko et al. (2024) build a monthly index of geopolitical risk from articles in local newspapers. They find that risk shocks derived from Russian-speaking newspapers have strong effects on the Russian economy.

An alternative to the construction of uncertainty indices based on text data is an identification of uncertainty shocks based on the high-frequency responses of asset prices around important geopolitical events. Piffer and Podstawski (2018) estimate a VAR model and identify an uncertainty shock from the variation in the price of gold

³Naboka-Krell (2024) constructs economic policy indices based on newspaper articles in Germany, Russia and Ukraine.

in a tight window around geopolitically important events. Ha et al. (2022) estimate the consequences of geopolitical risk shocks on the Korean economy by instrumenting risks by the asset price responses around important events.

Another relevant but still small literature addresses the specific consequences of the war in Ukraine. Federle et al. (2026) analyze the stock market response to the Russian attack and document a proximity penalty, i.e. a stronger negative response of stock markets in countries that are geographically closer to Ukraine. The authors show that this excess response reflects closer trade integration. Neuenkirch et al. (2025) study how financial markets respond to Western support for Ukraine. In the first weeks after the invasion, stock prices fall as a response to announcements of support. Thereafter, hawkish announcements in support of Ukraine increase stock prices. Tong (2024) uses a heteroscedasticity-based identification of financial shocks related to the war in Ukraine and estimates their effects on 86 open economies.⁴

The paper is organized as follows. Section two explains the data collection. Section three outlines the construction of our uncertainty indicators. Section four introduces the VAR model and the identification of uncertainty shocks, while section five discusses the empirical results. Section six concludes. An appendix contains additional material.

2 Data

2.1 Data Collection

Our analysis is based on data from Twitter. Twitter data does not provide a representative sample of the German population. In fact, from the US we know that Twitter users tend to be younger and better educated than the average American adult (Wojcik and Hughes, 2019). They also have a higher income and are more likely to be Democrats. There is no comparable study available for Germany, but it is very likely that the same tendencies are present in Germany as well. Though the number of Twitter users is still relatively low compared to traditional media, the strong presence of journalists, politicians, academics and other opinion leaders make it a very influential platform.⁵

With these limitations in mind, we focus on German-speaking tweets containing

⁴Other papers conducting event studies on the effects of the war on equity markets are Ahmed et al. (2023), Bounou and Yatié (2022), Huang and Lu (2022) and Izzeldin et al. (2023).

⁵As of January 2022, there were 7.75 million registered Twitter users in Germany, see <https://www.statista.com/statistics/242606/number-of-active-twitter-users-in-selected-countries/>.

selected keywords related to the war, i.e. *Zelenskyy*, *Putin*, *Ukraine*, *Kyiv*.⁶ Adding additional search terms such as *Crimea* or *Northstream* would not result in many more tweets. Keep in mind that tweets that contain *Ukraine* and *Crimea* jointly, for example, are already in our data set. Hence, these additional search strings would matter only for the small number of tweets that use *Crimea* without also using *Ukraine*, *Putin* or the other search terms. At the same time, we acknowledge that the vocabulary becomes more specific as the war continues. After a year of war, Twitter users might use much more specific terms related to the Ukrainian geography or specific weapon systems without also using a term such as *Ukraine*. Nevertheless, we believe our subset of tweets is informative about the sentiment on Twitter.

We collect tweets and user information from January 1, 2021 until February 28, 2023 using the Twitter API. The data contains user information and includes the account name, number of followers and followees, total number of tweets, user location, and whether the account has been verified by Twitter. Note that Twitter users have to actively choose to add their location to their accounts and tweets. As a result, only a small percentage of tweets include geotags. With that in mind, we do not limit our analysis solely to Germany as a geographical location. As a result, we obtain over nine million tweets sent by a total of 468,596 distinct Twitter users.


To focus on the textual component only, we then apply the following preprocessing steps for each tweet. First, we remove Twitter handles such as profile names starting with @, hashtags, hyperlinks, RT (retweet) tags, digits, special characters and line breaks. Second, to ensure that the tweets collected are unique, we remove any retweets, duplicates and empty entries. As a result, we end up with a database of 8,457,134 unique tweets. Third, given that the Twitter API provides tweet data in UTC time, we convert all tweet dates to the Berlin/Europe time zone. From the 8.4 million tweets, 4.7 million are from users who provide information about their location. 2.9 million tweets are from users based in Germany, 0.2 million from users based in Austria and 0.2 million from users based in Switzerland. Figure (1) shows the total number of unique tweets according to the respective keywords.

2.2 Measuring uncertainty

To measure the perceived uncertainty expressed in tweets, we first make use of state-of-the-art machine learning methods. Over the past few years, the use of transformer-based models has gained in popularity. The most prominent example in

⁶To ensure that we also consider the Ukrainian and Russian spellings of these keywords, our analysis also accounts for the corresponding transliteration of specific variations. As a result, our analysis includes the following terms: Selenski, Selenskyj, Zelenskyj, Kiev, Kiew, Kyiv, Kyjiw, Putin and Ukraine.

Figure 1: Number of tweets



Notes: The figure shows the total number of unique tweets according to the respective keywords, i.e. All keywords ($n=8,457,134$), Ukraine ($n=4,948,399$), Putin ($n=3,441,673$), Kyiv ($n=346,305$) and Zelenskyy ($n=364,049$). Note that the category *Kyiv* (*Zelenskyy*) also includes alternative transliterations, i.e. *Kiew*, *Kyjiw*, and *Kiev* (*Selenski*, *Selenskyj* and *Zelenskyj*).

the field of Natural Language Processing is the approach introduced by Devlin et al. (2018), namely Bidirectional Encoder Representations from Transformers (BERT). Using a transformer architecture and the concept of transfer learning, these models are usually trained on large text corpora, which enables them to achieve deeper learning and understanding of general language patterns and semantic relationships.⁷ The latter is achieved by using the masked language model approach, namely the masking out of specific words in a text. By training the model to predict the original word using the surrounding text context, these models learn to understand general language patterns and human language structure.

Since we are primarily interested in the classification of tweets in the "uncertainty" category, we make use of zero-shot classification with German BERT (GBERT), a pre-trained transformer model for the German language.⁸ The main idea is to utilize pre-trained models for a specific task and apply them to a different target task or application than their original training purpose. As in our case, the main advantage of zero-shot classification with GBERT lies in the extensive acquisition

⁷Malte and Ratadiya (2019) provide an overview of the evolution and advantages of transfer learning techniques.

⁸For a detailed model description see https://huggingface.co/Sahajtomar/German_Zeroshot. The model performs with 85.5 and 83.6 accuracy depending on the training data set.

of general language during the pre-training phase. For the pre-training data set for GBERT, Chan et al. (2020) use over 160 GB of German text data and cover a wide variety of domains such as, among others, Wikipedia articles, parliament speeches, books and German court decisions. By incorporating different text data during the pre-training phase, this approach enables a deeper and more nuanced understanding of human language.⁹ Thus, it allows the model to perform different classification tasks without the need for domain-specific training for each individual class, making it highly efficient and flexible in the adaptation to new domains. Using generalized labels, this approach allows us to assign probabilities to categories that were not explicitly included in the training data. In our case, the model assigns a probability to "uncertainty", our pre-defined category.

Next, we utilize the "uncertainty" probabilities as a measure of uncertainty inherent in a given tweet. To better understand how our approach measures the uncertainty expressed in a tweet, we shed light on two examples from leading German politicians. The first tweet is from German Chancellor Olaf Scholz, who tweeted the following on May 5, 2022:

Olaf Scholz (@Bundestkanzler) on May 05 2022

Tweet: Es ist wieder Krieg in Europa – unser ganzer Kontinent lebt gerade in einer Ausnahmesituation. Es gibt kein Drehbuch für das, was vor uns liegt. Meine Aufgabe ist es, unser Land sicher durch diese Zeit zu steuern. Und das tue ich.

Translation: There is war again in Europe - our entire continent is currently living in an exceptional situation. There is no script for what lies ahead of us. My task is to safely navigate our country through this time. And that's what I'm doing.

This is an example of a tweet for which the model assigns a very high uncertainty value, with the probability of the label "uncertainty" being 0.92. Interestingly, this value is high even though the word "sicher" ("safely") appears in the tweet. The phrase "Es gibt kein Drehbuch für das, was vor uns liegt." ("There is no script for what lies ahead of us.") effectively conveys uncertainty without using uncertainty-specific vocabulary. As a result, dictionary-based methods, e.g. used by Baker et al. (2016) and others, might struggle to detect the uncertainty expressed in this sentence.

⁹However, it is noteworthy to mention that the results of pre-trained language models should be interpreted with caution. As the training data set may inadvertently inherit preexisting cultural biases or contain false and misleading information, it cannot be ruled out that these biases might potentially be incorporated during the pre-training phase and consequently impact the model results.

However, in our approach, the model takes the context of a word into account, thus detecting uncertainty even when uncertainty-specific words are absent. The model assigns an even higher uncertainty value of 0.94 to the mentioned sentence by itself. The second tweet is from Federal Foreign Minister Annalena Baerbock, who tweeted:

Annalena Baerbock (@ABaerbock) on May 10 2022

Tweet: Ich bin unfassbar froh, hier im freien #Kiew zu sein. Der dafür notwendige Mut der Ukrainer*innen ist ergreifend. Meine Botschaft ist klar: Die #Ukraine kann sich auf unsere Unterstützung verlassen - nicht nur militärisch, nicht nur heute.

Translation: I am incredibly happy to be here in free #Kyiv. The courage of the Ukrainian people is moving. My message is clear: Ukraine can rely on our support - not only military, not only today.

The probability of the label "uncertainty" for this tweet is a very low value of 0.01. This low level of uncertainty can be largely attributed to the last sentence: "Die #Ukraine kann sich auf unsere Unterstützung verlassen - nicht nur militärisch, nicht nur heute." ("Ukraine can rely on our support - not only military, not only today.").

3 A Twitter-based uncertainty indicator

3.1 Index construction


Using the probabilities obtained from the zero-shot classification, we construct our index based on the assumption that a higher probability value for the "uncertainty" label corresponds to a high level of uncertainty inherent in a given tweet. To create this index, we calculate daily averages of uncertainty probabilities across all tweets. As illustrated in Figure (1), there is a significant change in the number of tweets over the sample period. The huge difference between the number of tweets at the beginning of this period and the onset of the war could pose a challenge to the construction of the index. To account for the varying number of tweets, we assign weights to each day and then multiply the average uncertainty values by these weights. The weights w_t are constructed as follows

$$w_t = \frac{\log(\#tweets_t + 1)}{\sum_{t=1}^T \log(\#tweets_t + 1)}, \quad (1)$$

where $\#tweets_t$ is the number of tweets used to construct the average uncertainty of day t over the sample $t = 1, \dots, T$. This method of calculation results in what

we refer to as the Ukraine Uncertainty Index.¹⁰ In our construction, we assume that each tweet contributes equally to aggregate uncertainty. However, to account for the potential influence of a tweet’s reach, we also compute an alternative index that weights tweets based on their number of retweets. Since this index is extremely similar to our baseline index, we relegate it to the online appendix. The appendix also reports the results based on an alternative uncertainty index, which is constructed without assigning weights.

Figure 2: The Ukraine Uncertainty Index



Notes: The figure shows the standardized daily Ukraine Uncertainty Index derived from German tweets.

Figure (2) depicts the standardized uncertainty index, alongside the uncertainty associated with specific events.¹¹

The index fluctuates at a relatively low level from January 2021 to November 2021. As early as April 2021, information surfaced about Russia amplifying its military presence near the Ukrainian border and in occupied Crimea. Correspondingly, the uncertainty index exhibits brief initial spikes around this period, even though the general level of uncertainty remains low until November 2021. On November 13, 2021, Ukrainian President Volodymyr Zelenskyy reported that approximately 100,000 Russian troops had amassed at the Ukrainian border. The situation continued to escalate thereafter, and the uncertainty index surged correspondingly.

¹⁰The uncertainty index is available at www.ukraine-uncertainty.de.

¹¹A comprehensive chronology of the events during the war can be found at <https://edition.cnn.com/interactive/2023/02/europe/russia-ukraine-war-timeline/index.html> or <https://commonslibrary.parliament.uk/research-briefings/cbp-9476/>.

In the German Twitter-sphere, a notable event was the visit of German Chancellor Olaf Scholz to Russian President Vladimir Putin on February 15, 2022. This was when the escalating conflict and troop deployment along the Ukrainian border were discussed directly. Following the visit, there was increased talk that the visit had contributed to a de-escalation of tensions. The index also shows a decrease in uncertainty on this day and the day following the visit. However, as subsequent events unfolded, it became clear that the situation had not improved. On the contrary, Russian President Putin declared the commencement of a "special military operation" on February 24, 2022, effectively signaling the outbreak of war. This date corresponds to the highest recorded value of uncertainty for the entire sample. After this peak, the index declines but stays at an elevated level. The beginning of April 2022 sees another noticeable surge in uncertainty, coinciding with the disclosure of the Bucha massacre.

Counteroffensives from the Ukrainian army characterized the following summer. As they sought to reclaim territories seized by Russia, the uncertainty index decreased, reaching a low on August 9, 2022, when the Ukrainian military targeted the Crimean airport. From this point forward, uncertainty starts increasing again. October 10, 2022, records a significant level of uncertainty, a day marked by severe attacks on Ukraine's water and electricity infrastructure, which led to blackouts in Kyiv and seven other regions. November 2022 experiences a considerable fluctuation in day-to-day uncertainty. On November 11, the index records its lowest value since the start of the war, a day when the Ukrainian army successfully recaptured the city of Kherson and other regions. However, a few days later, on November 15, the uncertainty index significantly increases following a missile strike on Polish territory near the Ukrainian border.

3.2 Comparing our index

We now compare our Twitter-based uncertainty index with four alternatives. These include the Geopolitical Risk Index of Caldara and Iacoviello (2022) for Germany, the Economic Policy Uncertainty (EPU) Index of Baker et al. (2016) for Germany and the GPR Index Germany (Bondarenko et al., 2024). All three indicators are newspaper-based measure of uncertainty or risk. Our fourth benchmark is the ifo Dispersion Measure (Grimme, 2017), which measures the dispersion of firms' responses in the monthly ifo business survey.¹² As these four indices are available only at a monthly frequency, we aggregate our index to the same frequency by selecting

¹²Bachmann et al. (2022) shows that firm-level uncertainty remained moderate at the beginning of the war in Ukraine. In fact, uncertainty was much lower than during pandemic.

the last observation for each month.

Figure (3) shows that all indices are strongly correlated. Our index showcases the highest correlation with the GPR Index Germany and the Geopolitical Risk Index for Germany, with correlation coefficients of 0.84 and 0.76, respectively.

When comparing our index with both the Geopolitical Risk Index for Germany and the GPR Index Germany, it is apparent that the increase in geopolitical risk from January to March 2022 is larger in these indices than for our index. Despite this, the uncertainty indicated by our Twitter-based index remains elevated, in contrast to the decreasing trends in both the Geopolitical Risk and GPR Index Germany. It is important to stress that our index does not rely on the occurrence of specific words in a given tweet, in contrast to the dictionary-based indices.

Our index also correlates with the EPU Germany and the ifo Dispersion Measure, as indicated by correlation coefficients of 0.73 and 0.70. Both the EPU Germany and the ifo Dispersion Measure exhibit similar dynamics, yet our index signals an earlier increase in uncertainty - observable from December 2021 and reaching its peak in February 2022. By comparison, the EPU Germany and the ifo Dispersion Measure do not exhibit a rise until January, reaching their respective peaks in March 2022.

In conclusion, our index is particularly suited for our analysis due to its specific focus on the uncertainty surrounding the Ukraine war. Furthermore, the use of high-frequency Twitter data allows our index to detect changes early on, maintaining this advantage even when the data is aggregated to a monthly frequency.


4 Estimating the economic effects on the German economy

In this section, we estimate the consequences of an increase in uncertainty about the war in Ukraine on the German economy. Our preferred tool is a structural VAR model, in which we identify an exogenous increase in uncertainty and trace its effect on real and financial variables. The key challenge is identification. Therefore, we will first discuss our approach to identification before introducing the model.

4.1 Identification and data

The previous sections provide us with daily indicators of Uncertainty about the war in Ukraine as perceived by German Twitter users. We estimate a VAR model for a vector of endogenous variables that consists of the alternative uncertainty indicators as well as economic and financial data such as a measure of economic activity, stock

Figure 3: Comparing our index with alternative indices



Notes: The figure shows the monthly Ukraine Uncertainty Index derived from German tweets, the Geopolitical Risk Index for Germany, the EPU index for Germany, the ifo Dispersion Measure and the GPR Index Germany. Each index is standardized.

prices and stock market volatility. As a matter of fact, changes in uncertainty are not exogenous with respect to the state of the German economy. Real and financial variables respond to an increase in uncertainty, but at the same time uncertainty is sensitive to the state of the economy. This mutual, contemporaneous interaction requires careful identification of exogenous variation in uncertainty.

Our identification scheme draws on the fact that Twitter is active around the clock while the financial markets are active during trading hours only. This key difference allows us to isolate variations in uncertainty that are exogenous with respect to financial variables and most forms of economic activity. Consider an increase in uncertainty over-night, i.e. before the stock market opens. This increase is clearly exogenous with respect to the stock market opening in the morning. In contrast, consider an increase in uncertainty at noon. This could equally well drive stock prices or could itself be the result of events on the stock market. Our identifying assumption is that Twitter activity at night is exogenous with respect to financial variables and economic activity realized during the following day.

To implement this identification scheme, we split our uncertainty indicator into two parts: the first is uncertainty as reflected by Twitter activity between 05:30pm on day $t - 1$ and 08:59:59am on day t . We refer to this as unc_t^{night} . The second part is unc_t^{day} and reflects uncertainty on twitter between 09:00am and 05:29:59pm on day t . On a given day t , unc_t^{night} should affect financial variables and economic activity,

while the reverse effect can safely be excluded. Hence, we study the response of the endogenous variables to an increase in unc_t^{night} .

The vector of endogenous variables is

$$y'_t = \begin{bmatrix} unc_t^{night} & unc_t^{day} & stockp_t & stockv_t & gasp_t & activity_t & infl_t^e \end{bmatrix}. \quad (2)$$

We include the log ($\times 100$) of the German stock market index DAX, $stockp_t$, the log ($\times 100$) of the volatility index (VDAX) of the German stock market, $stockv_t$, and the log ($\times 100$) of the gas price (Dutch TTF), $gasp_t$. These variables are easily available from standard databases. In an alternative specification, we replace stock prices with German 10-year government bond yields.

We also want to include a measure of real economic activity in our VAR model. We need an activity measure at a daily frequency, which is naturally noisy. To address this, we use the weekly activity index provided by Deutsche Bundesbank, see Eraslan and Götz (2020), which is the common component of several high-frequency indicators such as electricity usage, Google search volume for "unemployment", pedestrian frequency, air pollution, credit card transactions and many more.¹³ We interpolate the weekly index (seasonally adjusted and calendar adjusted) to daily frequency using the Chow-Lin procedure with the daily truck toll mileage index as a reference series. The resulting daily index of activity is $activity_t$. In the online appendix, we depict this series as well as all other macroeconomic and financial time series used in the estimation. As an alternative to the interpolated activity indicator, we use the daily Perceived Economic Situation indicator provided by trendEcon.¹⁴ This indicator summarizes the Google search intensities on the following keywords: "Wirtschaftskrise" (economic crisis), "Kurzarbeit" (short-term work), "arbeitslos" (unemployed) and "Insolvenz" (insolvency).

The final variable, $infl_t^e$, is expected inflation measured as the break-even inflation rate, i.e. the difference between yields on inflation-indexed and conventional German government bonds. As a matter of fact, this difference is a measure of expected inflation only if we abstract from risk premia and liquidity premia, respectively. In the baseline model, we use five-year break-even inflation rates, but will also show results for ten-year expectations. The data is retrieved from Refinitiv Datastream. As we separated uncertainty into night and day, we can now safely impose a recursive ordering onto the endogenous variables. This implies that we restrict the contemporaneous interaction of the variables such that an increase in unc_t^{night} can

¹³The data is available at <https://www.bundesbank.de/en/statistics/economic-activity-and-price-s/weekly-activity-index/weekly-activity-index-for-the-german-economy-833976>.

¹⁴See <https://www.trendecon.org/>.

contemporaneously impact all the remaining variables in the system, while unc_t^{night} itself responds with a lag of one day to all other variables. Consistently, we will estimate the responses of the endogenous variables to an increase in unc_t^{night} . We expect equity prices to fall, stock market volatility and gas prices to increase and economic activity to contract as a response to an increase in uncertainty. Uncertainty during the trading hours is also expected to increase following higher uncertainty the night before.

In the online appendix, we report results from extensions of the model that include additional variables ordered first. These extensions address the problem of omitted variables. We show that the results remain unchanged if we allow an index of geopolitical risk or the Citigroup Economic Surprise Index to drive all variables contemporaneously.

We estimate the VAR model over the sample period from January 3, 2022, to February 28, 2023. All weekends are excluded. The uncertainty indicators are standardized to have a standard deviation of one over the sample period. This allows us to compare the responses of the endogenous variables to different types of uncertainty.

4.2 VAR model

Based on the 7×1 vector of endogenous variables, y_t , we estimate a VAR model with p lags

$$y_t = A_1 y_{t-1} + \dots + A_p y_{t-p} + \varepsilon_t, \quad (3)$$

where $A_1, \dots, A_p$ are 7×7 coefficient matrices. The model also includes a vector of intercept terms, which we ignore in the exposition. The vector of error terms, ε_t , follows a multivariate normal distribution, $\varepsilon_t \sim N(0, \Sigma)$, where Σ is the variance-covariance matrix with $E(\varepsilon_t \varepsilon_t') = \Sigma$. The residuals are mutually uncorrelated at all leads and lags.

We estimate the VAR model using Bayesian methods. Assuming a Normal-Wishart prior, we interpret the VAR parameters as random variables drawn from a probability distribution.¹⁵

As outlined in the previous section, the identification of the structural uncertainty shock is based on the exogenous variation in Twitter-based uncertainty at night when the stock market is closed. This amounts to a recursive identification scheme, such that we can write the model in its structural form

$$D_0 y_t = D_1 y_{t-1} + \dots + D_p y_{t-p} + \eta_t, \quad (4)$$

¹⁵All estimations are carried out using the BEAR toolbox for MATLAB, see <https://www.ecb.europa.eu/pub/research/working-papers/html/bear-toolbox.en.html>.

where $\eta \sim N(0, \Gamma)$ is the vector of structural shocks and the D matrices are defined appropriately. With $D = D_0^{-1}$, the reduced-form error terms and the structural shocks are linked by $\varepsilon_t = D\eta_t$. We assume that D is lower triangular, thus imposing restrictions on the contemporaneous relationship between the variables.

5 The effects of uncertainty

We now present the estimated impulse response functions, i.e. the responses of the endogenous variables to an exogenous increase in uncertainty outside the trading hours. The size of the uncertainty shock is one standard deviation. Each figure also shows 68% probability bands.

5.1 Baseline results


The baseline results for the aggregate uncertainty index are shown in Figure (4). A higher uncertainty at night immediately raises uncertainty during day. Both measures of uncertainty remain significantly above their mean for four weeks. The increase in uncertainty has strong and significant effects on financial markets, economic activity and expected inflation. Stock prices fall by one percent and stock market volatility increases by four percent. Both responses are highly persistent. Higher uncertainty about the war in Ukraine also increases the price of natural gas. A week after the shock, gas prices are about one to two percent higher. The daily measure of economic activity falls upon an increase in uncertainty. This drop in activity is consistent with the fall in the stock market valuation. However, the fall in economic activity is relatively short-lived. Finally, the uncertainty shock causes a persistent increase in expected inflation. Two weeks after the shock, five-year expected inflation is 0.05 percentage points higher.¹⁶

The responses of economic activity and break-even inflation have opposite signs. This stands in contrast to the influential work of Leduc and Liu (2016), show that the effects of an uncertainty shock on the U.S. economy resemble those of a negative aggregate demand shock, i.e. a fall in both output and prices. In contrast, our results suggest that uncertainty about the war in Ukraine resembles an adverse aggregate supply shock.

The results for the model with ten-year bond yields are shown in Figure (5). In times of uncertainty, German government bonds serve as a safe haven for investors. The safe-haven status would suggest that yields fall when uncertainty increases. However,

¹⁶The higher expected inflation is consistent with the findings of Afunts et al. (2024).

Figure 4: The response to an uncertainty shock (baseline model)




Notes: The figure shows the responses of the endogenous variables to an uncertainty shock at night one standard deviation in size. The shaded areas reflect 68% probability bands and the horizontal axis denotes days after the shock.

we find that the uncertainty shock causes an increase in yields by about three basis points after three weeks. This increase is consistent with the higher expected inflation rate and also with markets' anticipation of higher financing needs of the government, which apparently dominate the safe-haven effect. All other impulse responses remain unchanged.

Figure (6) shows the effect of the uncertainty shock in a model with the economic perception indicator as a measure of real activity. The perceived situation deteriorates after an exogenous increase in uncertainty. In contrast to the baseline specification, this response is not statistically significant. All other results remain unchanged.¹⁷

¹⁷All results also remain unchanged if we include numerical information on the military events in Ukraine. For that purpose, we use the total number of daily fatalities available from the Ukraine Conflict Monitor at <https://acleddata.com/ukraine-conflict-monitor>.


Figure 5: The response to an uncertainty shock (with bond yields)



Notes: The figure shows the responses of the endogenous variables to an uncertainty shock at night one standard deviation in size. The shaded areas reflect 68% probability bands and the horizontal axis denotes days after the shock.

The appendix to this paper contains additional results. We show that the results remain unchanged if we replace five-year with ten-year expected inflation. We also provide a comparison of the impulse responses for models with alternative lag orders. The appendix contains responses to additional variables. We find that online job postings fall, though this response remains insignificant, that the number of tweets increases and that page views on Wikipedia for "Ukraine" and "Kiev" increase strongly after an uncertainty shock. Finally, we show the results for a VAR model estimated using least-squares rather than Bayesian methods.

Figure 6: The response to an uncertainty shock (with perceived economic situation)



Notes: The figure shows the responses of the endogenous variables to an uncertainty shock at night one standard deviation in size. The shaded areas reflect 68% probability bands and the horizontal axis denotes days after the shock.

5.2 Alternative identifications

While the identification of uncertainty shocks in our baseline model rests on the distinction between Twitter posts before and during trading hours, we now pursue two alternative identifications. First, we adopt a straightforward recursive ordering with daily uncertainty ordered first. Hence, we no longer separate tweets at night from tweets at day. The results are reported in Figure (7) and are essentially unchanged compared to the baseline model.

Second, we follow Romer and Romer (2004), Coibion (2012) and Caldara and Herbst (2019), among others, and adopt a hybrid approach. Specifically, we select several dates with important news about Ukraine and calculate the change in the uncertainty index on these dates. These news are all unrelated to the business cycle in Germany.


Hence, changes in uncertainty on these days can be considered exogenous. We no longer differentiate between day and night but use all tweets. Finally, we put the cumulative sum of the changes of uncertainty on these dates as the first variable in the vector of endogenous variables. Hence, the exogenous news is allowed to contemporaneously affect the remaining variables in the system. The VAR structure is needed to model the dynamic effect of the uncertainty shock. The vector of endogenous variables now is

$$\mathbf{y}'_t = [\Delta^{cum}unc_t \quad unc_t \quad stockp_t \quad stockv_t \quad gasp_t \quad activity_t \quad infl_t^e]. \quad (5)$$

where $\Delta^{cum}unc_t$ is a cumulative change of the main uncertainty indicator on the following days: February 11 2022 (White House warns Russian invasion of Ukraine may be imminent); February 15, 2022 (Chancellor Scholz in Moscow); February 21, 2022 (Russia sends troops into Eastern Ukraine); February 24, 2022 (Russian invasion begins); March 24, 2022 (fighting near Zaporizhzhia nuclear power plant); April 4, 2022 (news about Bucha massacre emerge); April, 28 (Bundestag decides to send heavy weapons); July 22, 2022 (Black Sea grain deal); August 19, 2022 (Gazprom announces pipeline maintenance); September 21, 2022 (Russia declares partial mobilization); November 15, 2022 (misguided Ukrainian missile hits Poland); January 25, 2023 (German government decides to send Leopard tanks to Ukraine); February 20, 2023 (President Biden in Kyiv).


Figure (8) shows the estimated responses to an increase in $\Delta^{cum}unc_t$. The results are very similar to the baseline results. We find that higher uncertainty strongly depresses the stock index, increases stock market volatility, raises natural gas prices, and causes a drop in economic activity and a higher expected rate of inflation. Hence, the effects of uncertainty are not dependent on the baseline identification scheme based on the distinction between night and day. Note that the dimensions of the responses are not comparable across identification schemes. The alternative identification is based on the most important events only, such that the magnitude of the responses is larger than in the baseline model. The figure also shows the response based on this alternative identification but without the change in uncertainty on April 28, 2022; August 19, 2022 and January 25, 2022. The first and the third of these events reflects an endogenous policy response, while the second is tantamount to an adverse supply shock. Leaving these three events out has a minor effect on the results.

Figure 7: The response to an uncertainty shock (alternative identification I)



Notes: The figure shows the responses of the endogenous variables to an uncertainty shock one standard deviation in size. The shaded areas reflect 68% probability bands and the horizontal axis denotes days after the shock.

Figure 8: The response to an uncertainty shock (alternative identification II)




Notes: The figure shows the responses of the endogenous variables to an uncertainty shock one standard deviation in size. The shaded areas reflect 68% probability bands and the horizontal axis denotes days after the shock. The green dashed line is the response for the alternative identification that leaves out the change in uncertainty on April 28, 2022; August 12, 2022 and January 25, 2022.

5.3 The effect of uncertainty over time

Our sample period ends one year after the Russian invasion. We now investigate whether the effect of uncertainty changes over time. Since the day of the invasion, the German government introduced policies to mitigate the economic consequences of the war, such as the cap on average costs of natural gas for households by in October 2022. The government also initiated transfers to households particularly exposed to the energy shock. In addition, households and firms might have adjusted to the new geopolitical situation or pay less attention to it. We see that reflected in the number of tweets in Figure (1), which continuously fall after the peak in February 2022. Do these effects imply that the economy becomes less sensitive to an uncertainty

shock? In order to assess this question, we split the sample into two parts. The first subsample covers January 2022 to August 2022, while the second subsample covers September 2022 to February 2023. We re-estimate the baseline model over each subsample and present the resulting impulse response functions in Figure (9). The results clearly suggest that an uncertainty shock has a larger effect in the first subsample compared to the second. The increase in volatility, gas prices and expected inflation are much larger in the first subsample. In fact, gas prices and inflation expectations are no longer responding to uncertainty in the second subsample. An unexpected increase in uncertainty causes a drop in real activity in the first subsample, but not in the second. Put differently, most of the responses over the full sample discussed before reflect the economic adjustment in the early months of the war in Ukraine.

Figure 9: The response to an uncertainty shock (subsamples)



Notes: The figure shows the responses of the endogenous variables to an uncertainty shock at night one standard deviation in size. The model is estimated for January 2022 to August 2022 (first subsample, blue) and September 2022 to February 2023 (second subsample, red). The shaded areas reflect 68% probability bands and the horizontal axis denotes days after the shock.

6 Conclusions

The uncertainty about the war in Ukraine, in particular about its potential escalation, its duration and its long-term consequences, is a determinant of the European business cycle. This paper estimates the effect of uncertainty on Germany. We quantify the evolution of uncertainty using Twitter data. Based on more than eight million tweets, we apply a machine learning approach to construct an index of uncertainty. In the second step, we estimate VAR models on daily data and identify an exogenous shock to uncertainty.

We find that an uncertainty shock causes a strong and persistent drop in equity

valuations, a strong increase in stock market volatility, an increase in natural gas prices and inflation expectations as well as a significant fall in real economic activity. These results are robust with respect to the measurement of uncertainty and the identification of uncertainty shocks. We also find that the consequences of an uncertainty shock are particularly severe in the first six months of the war and weaker thereafter.

Our conclusions are threefold. First, activity on Twitter, and potentially also on other social media channels, is a useful source of information in order to learn about the evolution of uncertainty. In future work, we will further exploit the cross-section of Twitter users, i.e. how different the perceived degree of uncertainty is across accounts, in addition to the time-series dimension of average uncertainty used in this paper. Second, uncertainty about the war in Ukraine does not only cause military and political upheaval but also leads to macroeconomic adjustments that resemble those after adverse supply shocks. This implies that an eventual resolution of uncertainty about the war will be expansionary for real economic activity and, at the same time, disinflationary. Third, our results support the broader notion that geopolitical risks in general are an important source of economic fluctuations.

References

- Afunts, G., Cato, M. and Schmidt, T. (2024). Inflation expectations in the wake of the war in ukraine, *Journal of Behavioral and Experimental Economics* **113**: 102303.
- Ahmed, S., Hasan, M. M. and Kamal, M. R. (2023). Russia–ukraine crisis: The effects on the european stock market, *European Financial Management* **29**(4): 1078–1118.
- Altig, D., Baker, S., Barrero, J. M., Bloom, N., Bunn, P., Chen, S., Davis, S. J., Leather, J., Meyer, B., Mihaylov, E., Mizen, P. and Parker, N. (2020). Economic Uncertainty Before and During the COVID-19 Pandemic, *Journal of Public Economics* **191**: 104274.
- Bachmann, R., Carstensen, K., Menkhoff, M. and Schneider, M. (2022). Umsatzunsicherheit und Umsatzerwartungen deutscher Firmen zur Zeit des Ukraine-Kriegs: Gas und Gasausfall spielen eine eher geringe Rolle, *ifo Schnelldienst* **75**(6): 33–36.
- Bachmann, R., Elstner, S. and Sims, E. S. (2013). Uncertainty and economic activity: evidence from business survey data, *American Economic Journal: Macroeconomics* **5**(2): 217–249.
- Baker, S. R., Bloom, N. and Davis, S. J. (2016). Measuring Economic Policy Uncertainty, *The Quarterly Journal of Economics* **131**(4): 1593–1636.
- Baker, S. R., Bloom, N., Davis, S. J. and Renault, T. (2021). Twitter-derived measures of economic uncertainty, *unpublished*.
- Bloom, N. (2014). Fluctuations in Uncertainty, *Journal of Economic Perspectives* **28**: 153–176.
- Bobasu, A. and De Santis, R. (2022). The impact of the Russian invasion of Ukraine on euro area activity via the uncertainty channel, *ECB Economic Bulletin* **4/2022**.
- Bondarenko, Y., Lewis, V., Rottner, M. and Schöler, Y. (2024). Geopolitical risk perceptions, *Journal of International Economics* **152**: 104005.
- Boungou, W. and Yatié, A. (2022). The impact of the Ukraine-Russia war on world stock market returns, *Economics Letters* **215**: 110516.
- Caldara, D., Conlisk, S., Iacoviello, M. and Penn, M. (2022). The Effect of the War in Ukraine on Global Activity and Inflation, *Feds notes, may 27, 2022*, Board of Governors of the Federal Reserve System.
- Caldara, D. and Herbst, E. (2019). Monetary Policy, Real Activity, and Credit Spreads: Evidence from Bayesian Proxy SVARs, *American Economic Journal: Macroeconomics* **11**(1): 157–192.
- Caldara, D. and Iacoviello, M. (2022). Measuring geopolitical risk, *American Economic Review* **112**: 1194–1225.
- Castelnuovo, E. (2023). Uncertainty before and during covid-19: A survey, *Journal of Economic Surveys* **37**(3): 821–864.

- Chan, B., Schweter, S. and Möller, T. (2020). German’s next language model, *arXiv preprint arXiv:2010.10906*.
- Coibion, O. (2012). Are the Effects of Monetary Policy Shocks Big or Small?, *American Economic Journal: Macroeconomics* **4(2)**: 1–32.
- Devlin, J., Chang, M., Lee, K. and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding, *arXiv preprint arXiv:1810.04805*.
- Eraslan, S. and Götz, T. (2020). An unconventional weekly economic activity index for Germany, *Technical paper no. 02/2020*, Deutsche Bundesbank.
- Federle, J., Meier, A., Müller, G. J. and Sehn, V. (2026). Proximity to war: The stock market response to the russian invasion of ukraine, *Journal of Money, Credit and Banking* **58(3)**: 681–703.
- Garicano, L., Rohner, D. and Weder di Mauro, B. (eds) (2022). *Global Economic Consequences of the War in Ukraine – Sanctions, Supply Chains and Sustainability*, Centre for Economic Policy Research, London.
- Geissler, D., Bär, D., Pröllochs, N. and Feuerriegel, S. (2023). Russian propaganda on social media during the 2022 invasion of ukraine, *EPJ Data Science* **12(1)**: 35.
- Grimme, C. (2017). Measurement of Corporate Uncertainty in Germany - the ifo Dispersion Measure, *ifo Schnelldienst* **70**: 19–25.
- Ha, J., Lee, S. and So, I. (2022). The impact of uncertainty shocks: evidence from geopolitical swings on the Korean peninsula, *Oxford Bulletin of Economics and Statistics* **84**: 21–56.
- Huang, L. and Lu, F. (2022). The cost of Russian sanctions on the global equity markets, *unpublished*, University of Hong Kong.
- Izzeldin, M., Muradoglu, Y. G., Pappas, Y., Petropoulou, A. and Sivaprasad, S. (2023). The impact of the Russian-Ukrainian war on global financial markets, *International Review of Financial Analysis* **87**: 102598.
- Jurado, K., Ludvigson, S. C. and Ng, S. (2015). Measuring uncertainty, *American Economic Review* **105(3)**: 1177–1216.
- Kilian, L., Plante, M. D. and Richter, A. W. (2025). Macroeconomic responses to uncertainty shocks: The perils of recursive orderings, *Journal of Applied Econometrics* **40(4)**: 395–410.
- Leduc, S. and Liu, Z. (2016). Uncertainty shocks are aggregate demand shocks, *Journal of Monetary Economics* **82**: 20–35.
- Ludvigson, S. C., Ma, S. and Ng, S. (2021). Uncertainty and business cycles: exogenous impulse or endogenous response?, *American Economic Journal: Macroeconomics* **13(4)**: 369–140.

- Malte, A. and Ratadiya, P. (2019). Evolution of transfer learning in natural language processing, *arXiv preprint arXiv:1910.07370*.
- Naboka-Krell, V. (2024). Construction and analysis of uncertainty indices based on multilingual text representations, *Economics Letters* **237**: 111653.
- Neuenkirch, M., Repko, M. and Weber, E. (2025). Hawks and doves: Financial market perception of western support for ukraine, *European Journal of Political Economy* **89**: 102744.
- Piffer, M. and Podstawski, M. (2018). Identifying uncertainty shocks using the price of gold, *The Economic Journal* **128**: 3266–3284.
- Romer, C. D. and Romer, D. H. (2004). A New Measure of Monetary Shocks: Derivation and Implications, *American Economic Review* **94**: 1055–1084.
- Tong, E. (2024). Repercussions of the Russia-Ukraine war, *International Review of Economics and Finance* **89**: 366–390.
- Wojcik, S. and Hughes, A. (2019). Sizing up twitter users, *Pew Research Study*. <https://www.pewresearch.org/internet/2019/04/24/sizing-up-twitter-users/>.


Appendices

Appendix A The role of retweets

In our construction, we assume that each tweet contributes equally to aggregate uncertainty. However, to account for the potential influence of a tweet’s reach, we compute an alternative index that weights tweets based on their number of retweets. Following the methodology of Baker et al. (2021), each tweet is weighted by $(1 + \log(1 + \#retweets_i))$, where $\#retweets_i$ denotes the number of retweets of tweet i . Next, we adjust for the daily tweet count by applying weighting scheme explained in the main paper.

Comparing this retweet-weighted index with our baseline index as illustrated in Figure (A.1), we observe only minor differences.¹⁸ Given this minimal deviation, we derive all results from our baseline index.

Figure A.1: The Ukraine Uncertainty Index (weighted by retweets)




Notes: The figure shows the standardized daily Ukraine Uncertainty Index (Retweets) derived from German tweets.

¹⁸Similar results are obtained when using likes or the number of followers of the tweet author as alternative weighting metrics instead of retweets.

Appendix B Uncertainty based on unweighted tweets

The uncertainty index used for the baseline results assigns weights on each day reflecting the relative number of tweets including our search strings. We believe this weighting is necessary in order to avoid erratic fluctuations in uncertainty due to a low number of tweets. Nevertheless, Figure (B.2) shows the estimated impulse responses once we use the average uncertainty across all tweets on a given day without assigning weights. The weighting scheme has no material effect on our results.

Figure B.2: The response to an uncertainty shock (unweighted tweets)



Notes: The figure shows the responses of the endogenous variables to an uncertainty shock at night one standard deviation in size. The shaded areas reflect 68% probability bands and the horizontal axis denotes days after the shock.

Appendix C Additional empirical results

Figure (C.3) depicts this series as well as all other macroeconomic and financial time series used in the estimation.

The series of the structural uncertainty shock estimated in the baseline model is shown in Figure (C.4). We observe particularly large realizations of the uncertainty shock around the beginning of the Russian invasion in late February 2022 and in the night between November 15 and 16 2022, when the misguided Ukrainian missile hits a village in Poland.

It could be argued that most of the economic effects is due to the big and persistent uncertainty shock in February 2022 when Russia started the invasion. In order to address this concern, we estimate the baseline model over a shorter sample period that starts in February 28, 2022, i.e. four days after the invasion. Figure (C.5) shows the estimated responses. Most responses remain unchanged when compared to the baseline model. The only exception is the response of gas prices: gas prices no longer respond significantly once we exclude the run-up to the invasion and the immediate aftermath.

In another alternative specification, we replace five-year expected inflation with ten-year expected inflation. The results are shown in Figure (C.6). The results are very similar to the baseline findings.

Figure (C.7) shows the estimated impulse response functions for the baseline model with eight lags and compares them to alternative models with ten, six and four lags, respectively. For shorter lag lengths, the stock price response becomes smaller, while the response of real economic activity becomes stronger. Most importantly, all of our results remain qualitatively unchanged.

The estimated VAR model incorporates a small set of variables only, thus leaving out important information. Suppose a variable omitted from the VAR affects all variables contemporaneously. In this case, our identification would still be valid as, by construction, shocks to uncertainty at night are causal with respect to the variables ordered behind it. Nevertheless, the omitted variable could affect the shape and the significance of the impulse response. Suppose news about geopolitical risks breaks and affects both our uncertainty index and the economic variables on the same days. Geopolitical risk needs to be ordered first in the VAR model to reflect this. In an alternative model, we estimate such a specification. We use the daily version of the global newspaper-based Geopolitical Risk Index (in $\log s \times 100$) of Caldara and Iacoviello (2022) and put it into the VAR model as the first variable. Hence, geopolitical risk can drive all other variables contemporaneously. The index picks up all the news about the war and other geopolitically relevant incidents not

incorporated in our Twitter data.

Figure (C.8) shows the estimated impulse response functions. We still plot the response to an unexpected increase of uncertainty during the night. Importantly, all our results remain broadly unaffected. In addition, we see that an increase in uncertainty raises geopolitical risk by about 5%. Thus, this alternative specification confirms that our estimated effects of uncertainty are robust with respect to including additional information omitted from the baseline model.


Alternatively, we use the Citigroup Economic Surprise Index for the euro area as the first variable in the model. This index summarizes the surprise components of macroeconomic data releases. It is exogenous with respect to the other variables in the model and, hence, should be ordered first. All results remain unchanged, see Figure (C.9). This suggests that including omitted information about the state of the economy does not affect our main findings.

Figure (C.10) shows the responses of additional variables, i.e. the number of online job postings on indeed.com, the number of tweets in our sample and the views of the Wikipedia page on "Ukraine" and "Kiew". All four series are included in logs ($\times 100$).¹⁹ Both Wikipedia pages are in German. We estimate the baseline model and replace the seventh variable, i.e. expected inflation, with the three alternative variables discussed here one at a time. To save space, we show only the responses of these additional variables. The remaining variables do not change in an important way. Job postings fall as a response to an adverse uncertainty shock. However, the fall in postings is significant only on impact. An uncertainty shock causes a strong increase in the number of German tweets on the war by about 45%. We also find a large increase in the interest in Wikipedia pages on Ukraine. Both sites, on "Ukraine" and "Kiew", experience 20% more views following an uncertainty shock.

Figure (C.11) depicts the estimated impulse response from the baseline model estimated by least squares rather than Bayesian methods. We plot 90% confidence bands, which are more common in the frequentist VAR literature than the 68% bands from the Bayesian literature. All main results remain unchanged.


¹⁹The job postings are available at <https://www.hiringlab.org/de/>.

Figure C.3: Data series for the VAR model




Notes: The figure shows the data series used in the estimated VAR model. The weekly activity index is interpolated to daily frequency using the Chow-Lin procedure with the daily truck toll mileage index as a reference series. All series are explained in the text.

Figure C.4: Structural uncertainty shock




Notes: The figure shows the series of estimated structural uncertainty shocks (in standard deviations).

Figure C.5: The response to an uncertainty shock (sample begins in March 2022)




Notes: The figure shows the responses of the endogenous variables to an uncertainty shock at night one standard deviation in size. The shaded areas reflect 68% probability bands and the horizontal axis denotes days after the shock.

Figure C.6: The response to an uncertainty shock (with ten-year expected inflation)




Notes: The figure shows the responses of the endogenous variables to an uncertainty shock at night one standard deviation in size. The shaded areas reflect 68% probability bands and the horizontal axis denotes days after the shock.

Figure C.7: The response to an uncertainty shock (alternative lag orders)




Notes: The figure shows the responses of the endogenous variables to an uncertainty shock at night one standard deviation in size for the model with eight lags (black line) and alternative models with ten (green line), six (red line) and four (orange line) lags. The shaded areas reflect 68% probability bands and the horizontal axis denotes days after the shock.

Figure C.8: The response to an uncertainty shock (geopolitical risk ordered first)




Notes: The figure shows the responses of the endogenous variables to an uncertainty shock at night one standard deviation in size. The shaded areas reflect 68% probability bands and the horizontal axis denotes days after the shock.

Figure C.9: The response to an uncertainty shock (Citigroup Surprise Index ordered first)




Notes: The figure shows the responses of the endogenous variables to an uncertainty shock at night one standard deviation in size. The shaded areas reflect 68% probability bands and the horizontal axis denotes days after the shock.

Figure C.10: The response to an uncertainty shock (additional variables)



Notes: The figure shows the responses of the additional variables, each used as the alternative seventh variable replacing expected inflation, to an uncertainty shock at night one standard deviation in size. The shaded areas reflect 68% probability bands and the horizontal axis denotes days after the shock.

Figure C.11: The response to an uncertainty shock (OLS-VAR)



Notes: The figure shows the responses of the endogenous variables to an uncertainty shock at night one standard deviation in size. The shaded areas reflect 90% probability bands. The red dashed line is the response in the baseline model and the horizontal axis denotes days after the shock.

Essay II

Geopolitics in the Boardroom: How German Managers Respond to the War in Ukraine

Acknowledgement

We thank Lukas Hack and Davud Rostam-Afschar for sharing their data from the *German Business Panel* and for very helpful comments on an earlier draft. We also thank Michael Ehrmann for comments. David Finck kindly shared his MATLAB codes. Financial support from Fritz Thyssen Stiftung is gratefully acknowledged.

Geopolitics in the Boardroom: How German Managers Respond to the War in Ukraine

Moritz Grebe* Sinem Kandemir† Peter Tillmann‡

Abstract

We use unique daily data from the *German Business Panel*, a survey among German managers, to estimate the impact of uncertainty about the war in Ukraine on managerial plans and expectations. Each workday, the online survey invites a random set of firms to answer a standard set of questions. War-related uncertainty is extracted from a corpus of more than eight million German-language Twitter messages about the war in Ukraine. Higher uncertainty has strong adverse effects by prompting firms to plan higher prices and expect a drop in dividends, revenues, profits and investment. The persistent fall in expected R&D expenditure suggests that the shock has consequences for long-run productivity growth. Managers' satisfaction with economic policy increases after an uncertainty shock. The sensitivity of managers with respect to war-related uncertainty depends on the overall level of geopolitical risk.

Keywords: war, Twitter, geopolitical risk, machine learning, business cycle

JEL Classification: D8, E3, G1

*Justus Liebig University Giessen, Email: moritz.grebe@wirtschaft.uni-giessen.de

†Justus Liebig University Giessen, Email: sinem.kandemir@wirtschaft.uni-giessen.de

‡Justus Liebig University Giessen, Email: peter.tillmann@wirtschaft.uni-giessen.de

1 Introduction

The full invasion of Ukraine by Russian forces on February 24, 2022, caught many in the West off guard. Households, firms and policymakers had to adjust to the new level of geopolitical tensions. Due to its reliance on Russian gas, the German economy was particularly exposed to geopolitical uncertainty.

In this paper, we empirically study how German managers revised their plans and expectations in light of the uncertainty related to the war in Ukraine. Thus, we ask how the war resonates in German boardrooms. Specifically, we investigate how planned adjustments of prices, dividends, and R&D expenditure, as well as expectations about revenues, profits and investment respond to uncertainty shocks. We use data from the *German Business Panel* (Bischof et al., 2024), provided by Hack and Rostam-Afschar (2024). This online survey has the advantage of sampling firms randomly each workday, allowing us to construct a daily indicator from the responses. The large underlying sample also covers many small and medium-sized firms, which are the backbone of the German economy. Both the influential ifo survey and the relatively new Bundesbank Online Panel of Firms (BOP-F) offer time-stamped responses, which are published as monthly averages. However, the advantage of the data from the *German Business Panel* is that the cross-sectional sample sizes are stable within a month, and that composition effects are absent (Hack and Rostam-Afschar, 2024). Thus, our data allows us to track the views of managers on a daily frequency.¹ The daily frequency is necessary as uncertainty shocks materialize within days and can change plans and expectations rapidly. For the empirical analysis, we construct the weighted average of the daily average responses across firms, where the weights correspond to the number of surveys completed on each day.

Rising uncertainty about the war could affect managers' views and expectations through at least three different transmission channels.² First, higher uncertainty could cause markets and banks to raise risk premia on the external funding of firms. Second, higher uncertainty could lead managers to adopt a wait-and-see strategy and postpone investment decisions. Finally, higher uncertainty prompts households to save more. This type of precautionary saving is the flip-side of a drop in aggregate demand for consumption, which translates into lower demand for firms' products. Each of these three forces imply that higher uncertainty dampens real economic

¹See Sauer and Wohlrabe (2022) for early evidence on the consequences of the Ukraine war based on the ifo survey.

²Bobasu and De Santis (2022) discuss the impact of the Ukraine war on the euro area via the uncertainty channel.

activity. In this paper, we do not attempt to disentangle these transmission channels. Our measure of uncertainty about the war in Ukraine is the Ukraine Uncertainty Index from our companion paper (Grebe et al., 2024). This index provides a daily time series of the war-related uncertainty, which we derive from more than eight million German-speaking Twitter messages. Using methods from Natural Language Processing, we can assign to each tweet a probability that it expresses uncertainty. Aggregating over all tweets on a given day results in a time-varying uncertainty index. This index peaks on February 24, 2022, the day of the Russian invasion, and exhibits strong fluctuations thereafter. We estimate the impact of fluctuations in uncertainty on the views of German managers using a series of local projections (Jordà, 2005). These estimates provide impulse responses of the survey variables to changes in uncertainty.

Our results are threefold. First, we find that an increase in uncertainty about the war in Ukraine has strong adverse effects. Managers plan higher prices and expect a fall in dividends, profits, revenues, and investment spending. Hence, uncertainty has real economic consequences that are in line with the theoretical transmission channels outlined before. The expected survival rate of firms exhibits only a short-lived drop before returning to its mean. Therefore, managers do not perceive the uncertainty as a fatal shock for the existence of their businesses. Second, we find that firms also reduce their expected R&D spending after an increase in uncertainty. This suggests that the uncertainty has not just temporary business cycle effects but could also have long-run structural effects, as reduced R&D spending translates into weaker productivity growth in the long run. Third, we find that the satisfaction of managers with the economic policy of the government increases after a shock to uncertainty. This suggests that the German coalition government managed the immediate aftermath of the invasion quite well, at least in the eyes of corporate Germany.

There are two limitations that we should mention. First, we look at uncertainty related to the Ukraine war only. To be reflected in our corpus, a tweet needs to contain one of the pre-specified keywords. This also means that tweets about the indirect effects of the war, such as German energy policy or the international sanctions regime, are included only if they contain at least one of the keywords. A tweet about energy policy, for instance, without mentioning "Putin", "Ukraine" etc. is not included. Second, our survey data does not allow us to distinguish between sectors or regions. We also cannot control for factors such as size, financial leverage, integration into global supply chains, or energy intensity. All of these characteristics could affect a firm's sensitivity to the war in Ukraine.

Our paper is related to the burgeoning literature on the macroeconomic effects of uncertainty shocks (Baker et al., 2016; Jurado et al., 2015; Ludvigson et al., 2021). This literature was initiated by Bloom (2009).³ Though each paper draws on different ways to measure uncertainty and identify exogenous variation in uncertainty, they share a focus on the consequences of macroeconomic (policy) uncertainty.⁴ Bachmann et al. (2013) study the response of individual firms to aggregate uncertainty based on micro data from the ifo survey.

Another strand of the literature studies the consequences of geopolitical risk. The seminal paper by Caldara and Iacoviello (2022) introduces the Geopolitical Risk Index. Specifically, the authors search a large set of U.S. newspapers for terms related to geopolitical threats and acts. This index tracks the historical episodes of elevated geopolitical risks very well. In a VAR analysis, the authors show that an increase in geopolitical risk causes a decline in U.S. private fixed investment, stock prices and hours worked.⁵ Caldara et al. (2026) extend this research agenda and estimate the effects of geopolitical risk on inflation dynamics. In our companion paper, Grebe et al. (2024), we add to this literature by introducing the Ukraine Uncertainty index based on German-language Twitter posts. We show that higher war-related uncertainty depresses economic activity and stock prices, while raising inflation expectations and stock market volatility.⁶ Bondarenko et al. (2024) study how geopolitical risk is covered in local newspapers rather than U.S. media. They show that perceptions differ markedly. The effect of geopolitical shocks on the Russian economy depend on how the news is covered by newspapers. In a related strand of the literature, Hassan et al. (2019) and Hassan et al. (2024) construct measures of how political risk is perceived by firms from transcripts of quarterly earnings calls. They estimate the adverse effects of perceived risk on investment and hiring decisions. However, these measures reflect political risk in general, including domestic risk, and not just geopolitical risk.

The rest of this paper is structured as follows. In Section 2, we introduce the Ukraine Uncertainty Index. Section 3 explain the data from the *German Business Panel*. The empirical analysis is conducted in Section 4. Section 5 concludes. An online appendix contains additional material.

³For surveys on the vast literature that identifies the economic consequences of uncertainty shocks, see Bloom (2014), Cascaldi-Garcia et al. (2023) and Castelnovo (2023).

⁴Altig et al. (2020) apply the Baker et al. (2016) approach to measure economic policy uncertainty to the uncertainty to the Covid-19 pandemic.


⁵Wang et al. (2024) also show the strong negative impact of geopolitical risk and firm-level investment.

⁶Federle et al. (2026) study the responses of stock markets to the Russian invasion of Ukraine. The closer a country is to Ukraine, the stronger the drop in its share prices index.

2 The Ukraine Uncertainty Index

As a measure of uncertainty about the war in Ukraine, we use the *Ukraine Uncertainty Index*, which we present in more detail in Grebe et al. (2024).⁷ This index is based on a large dataset of 8,457,134 unique German tweets. These tweets were obtained over the period from January 1, 2021, to February 28, 2023, via Twitter’s API using specific keywords such as ‘Ukraine’, ‘Zelenskyy’, ‘Putin’, among others.

Figure 1: The Ukraine Uncertainty Index



Notes: Panel (a) shows the standardized daily Ukraine Uncertainty Index derived from German tweets, Panel (b) shows the standardized residual of a regression of the index on one lag of itself and day-of-the-week dummies.

We then perform a zero-shot classification using a German BERT (Bidirectional Encoder Representations from Transformers) model, which assigns a probability to the label ‘uncertainty’ for each tweet. We interpret this probability as a measure of the uncertainty conveyed by that tweet. The great advantage of this model is that it understands the linguistic context and can therefore identify expressions of uncertainty more effectively than dictionary methods, especially when classical uncertainty terms are absent. Next, we calculate the average uncertainty value for each day. Since the number of tweets per day varies widely in our sample - from a

⁷The index is available at www.ukraine-uncertainty.de.

few hundred before the war started to around 125,000 on February 24, 2022, the start of the war - we must adjust for this variability when constructing the index. Therefore, we multiply the daily averages by weights that account for the number of tweets each day.

Panel (a) of Figure (1) illustrates the resulting *Ukraine Uncertainty Index* from January 2022 onward. As shown, the index rises sharply at the beginning of the sample and reaches its maximum on the day the war began. The index falls over the summer of 2022 reaching a low when the Ukrainian military targeted the Crimean airport, but increases again during the fall. Uncertainty declines when the Ukrainian army recaptured the city of Kherson and other regions on November 11, though it jumps upwards on November 15, the day a misguided missile hit Polish territory near the Ukrainian border.

3 Daily views of German managers

The *German Business Panel* is a long-term online survey of German firms conducted by the supraregional project *TRR 266 Accounting for Transparency* of the German Research Foundation.⁸ A key property of this survey is particularly valuable for our purpose: on each workday a randomly selected set of firms is invited to answer the survey's core questions listed below. The sample from which the daily set of firms is randomly drawn is representative for the entire population of firms. Hence, the survey provides a reliable picture of the views of German managers at a high frequency. This is a major advantage over the Bundesbank Online Panel of Firms or the traditional survey conducted by the ifo Institute. According to Bischof et al. (2024), around 90% of the respondents belong to the top management of the firm, i.e. they are either the firm's owner or the CEO. Thus, the survey sheds light on backbone of the German economy. We received this daily dataset from Hack and Rostam-Afschar (2024).

We use the following ten variables from the survey, which are organized into two groups.⁹ Group I includes firms' responses to questions about planned managerial decisions. Firms are asked about their plans to change the following variables:

1. sales price
2. fixed costs, i.e. personnel or rents
3. dividends, i.e., payouts to owners

⁸For more information, see <https://gbpanel.org/>.

⁹In the appendix, we also provide the exact (German) questionnaire and its English translation.

4. R&D investment

In each case, respondents can choose between the options 1 (increase), -1 (decrease) and 0 (no change). The data we have available reports the average of the responses across firms on a given day. Hence, we do not have the individual firm responses available.

Group II includes the responses to questions about firms' expectations and assessments. Firms are asked about the following variables:

1. Expected firm survival rate within industry of their own firm (in percent).
2. Satisfaction with (contemporaneous) economic policy (11-point scale between 0 = very dissatisfied to 11 = very satisfied).
3. Expected percent change in revenue of their own firm in the current calendar year relative to the last calendar year (in percent).
4. Expected percent change in profit of their own firm in the current calendar year relative to the last calendar year (in percent).
5. Expected percent change in the number of employees of their own firm in the current calendar year relative to the last calendar year (in percent).
6. Expected percent change in investment of their own firm in the current calendar year relative to the last calendar year (in percent).

Again, the data we have available reports the average of the responses to each question across the firms that respond on a given day.

It should be stressed that several questions about expectations refer to the current calendar year. The idea is to capture information relevant for investors: the distance between targets typically announced in annual reports and the expectation of reaching these target. For example, a firm could announce in its annual report for 2021 a target of 3% profit growth for 2022 (relative to 2021). If the response to the question is 1%, then the target is missed by two percentage points.¹⁰ For surveys completed at the end of the calendar year, most realizations of the current calendar year are already known. Although the question is framed in terms of expectations, it becomes more and more backward-looking as the calendar year progresses. This should be kept in mind when interpreting the estimated impulse response functions. In particular, this property might explain why expected employment tends to increase after an adverse shock to uncertainty. The four questions on managers' plans address plans for the

¹⁰We thank Davud Rostam-Afschar for this point.

coming 12 months, not for the calendar year. Hence, the answers to these questions should truly reflect forward-looking expectations of managers.

In the online appendix, we show the number of responses to each survey question over the days of the week in our sample period from January 2022 to February 2023. Each question is answered by about 30 to 50 firms on an average working day. On the weekend, these numbers fall to single-digit responses. Therefore, we will use a weighted moving average of the daily responses below. The weight attached to each day is the relative number of responses submitted. Our baseline model relies on 28-day weighted moving averages. Hence, each average response in these 28 days is weighted by the number of responses relative to the total number of responses throughout the 28 days. We will also provide alternative estimates based on 14-day weighted moving averages constructed analogously.

The online appendix also provides plots of the evolution of each variable over time.

4 Empirical analysis


In this section, we first construct a series of uncertainty shocks. In a second step, we put this shock series into straightforward local projection models and estimate the impact of uncertainty on managerial plans and expectations.

4.1 Construction of shock series

Managers should change their plans and expectations when new information becomes available. Their decisions should not be driven by anticipated events. While the full invasion of Ukraine certainly came unexpectedly to most managers, a large part of the evolution of the Ukraine Uncertainty Index is potentially predictable. Therefore, we construct an uncertainty shock by removing the forecastable parts from the Ukraine Uncertainty Index. Specifically, we regress the index on one lag of itself and a set of seven day-of-the-week dummies, which account for the fact that Twitter users may not be equally active on the weekends. The standardized residual of this regression will serve as a driving variable of the empirical model estimated below. We refer to this residual as the uncertainty shock.

Panel (b) of Figure (1) shows our shock series. On February 24, 2022 the shock is more than six standard deviations higher than its mean. We also see considerable fluctuations over the remaining sample period with several realizations larger than two standard deviations. Figure (2) plots the distribution of the shock series. Again, the realization on the day of the full invasion is clearly visible as an outlier. In the online appendix, we show that the results remain robust even if we remove outliers.

Figure 2: Distribution of uncertainty shock



Notes: The figure shows the distribution of the uncertainty shock used in the estimation. We highlight the shock realization on February 24, 2022, the day of the full invasion.

4.2 Evidence from local projections

We now estimate the dynamic effects of an increase in average uncertainty on managers' plans, expectations, and opinions. The regression model in line of standard local projections (Jordà, 2005; Jordà and Taylor, 2025) is

$$y_{t+h} = c + \beta_h unc_t + \gamma_h X_t + \varepsilon_{t+h}. \quad (1)$$

The dependent variable is dated $t + h$, i.e. h days after the realization of the shock with $h = 0, \dots, 49$, where unc_t is the uncertainty shock at day t derived in the previous subsection. The estimated coefficient β_h reflects the impact of an increase in uncertainty on each of the dependent variables after h periods. Variations of the dependent variable explained by the uncertainty shock or the control variables are reflected in the residual ε_t . The vector X_t contains seven lags of the dependent variable and seven lags of firms' plans with respect to prices and dividends. We also include seven lags of the (log) Dutch TTF natural gas future to control for the war-related increase in energy costs. Generally, the results are very robust with respect to the specification of the vector of control variables.

Figure (3) presents the estimated coefficients for the four variables which reflect firms'

plans. We show the β_h coefficient with 68% and 90% confidence bands obtained from robust standard errors. All baseline results are for the 28-day weighted moving average of the responses as the dependent variables. The results are shown in blue. We also report the point estimates obtained for the 14-day weighted moving average of the responses as an orange (dashed) line.

An increase in uncertainty of one standard deviation leads firms to revise their planned price-setting. Initially, the response is positive, yet small. However, three weeks after the shock, the average firm becomes strongly inclined to raise prices. This response is very persistent and remains significantly positive over the projection horizon. Planned increases in prices are, of course, consistent with the subsequent surge in inflation in 2022.¹¹ We also find a significant drop in firms' planned dividends, fixed costs and R&D spending. One potential consequence of a drop in R&D spending could be a fall in productivity growth in the future. In this sense, a temporary increase in uncertainty could have long-run structural effects on the German economy. The responses of firms' expectations, see Figure (4) are consistent with these findings. Firms' expected revenues and profits sharply deteriorate. German managers' also reduce their expectations of investment expenditures. Taken together, these responses suggest that the uncertainty shock has strong adverse effects on German firms, which are consistent with the aggregate effects found in Grebe et al. (2024). The responses based on the 14-day weighted moving average of the data series are broadly similar to the baseline responses, though a bit more volatile.

While most responses indicate a relatively persistent drop in real economic activity, the expected survival rate dips only briefly and then quickly recovers. It even increases in the long-run. This could imply that after an initial transition, firms become more robust, leading to an increased expected survival rate. A less benign interpretation is that government support, such as transfers to cope with higher energy costs, keeps firms afloat that would otherwise exit the market.

Managers also expect more hires in the future. Note that this result could also be due to the backward looking nature of the question, which refers to current calendar year employment. If an adverse shock occurs at the end of the year, following increases in employment, the response to the shock could be positive.


Strikingly, higher war-related uncertainty causes an increase in the satisfaction with German policies. In times of geopolitical and economic uncertainties, managers feel well-governed. This is remarkable in light of the downward trend in satisfaction with policy. In the data series, satisfaction falls strongly after Russia's invasion of

¹¹The findings of Dräger et al. (2025) are also consistent with that. These authors use the Russian invasion of Ukraine as a natural experiment to conduct a survey among German economics professors. A geopolitical shock causes a significant increase in inflation expectations.

Ukraine. This underlines the importance of distinguishing between the unconditional trend in the data and the response of satisfaction with policy conditional on an uncertainty shock. Again, a somewhat less benign interpretation is also consistent with the results: higher satisfaction with economic policy could mean that managers believe that the government will cushion the consequences of adverse shocks and prevent firm exits. The data does not allow us to disentangle or distinguish between different transmission mechanisms.

The fact that higher uncertainty causes an upward revision of price-setting plans and a decline in expected revenues supports the notion that the shock resembles an aggregate supply-side shock, with prices and quantities responding with opposite signs. This is remarkable since Leduc and Liu (2016) argue that uncertainty shocks have consequences similar to aggregate demand shocks.


Figure 3: Response of business plans to an uncertainty shock



Notes: The figure shows the response to an uncertainty shock one standard deviation in size. The solid blue (dashed orange) line is the response of the 28-day (14-day) moving averages of the survey variables. The shaded areas reflect the 68% and 90% confidence bands based on robust standard errors.

Our Ukraine Uncertainty Index measures the perceived uncertainty about the war in Ukraine. We now want to control for news about actual war-related events to make sure that we do not confuse an increase in uncertainty with a higher number

Figure 4: Response of expectations to an uncertainty shock



Notes: The figure shows the response to an uncertainty shock one standard deviation in size. The solid blue (dashed orange) line is the response of the 28-day (14-day) moving averages of the survey variables. The shaded areas reflect the 68% and 90% confidence bands based on robust standard errors.


of incidents in Ukraine. We measure actual war-related events using four alternative indicators and include each of them separately as an additional control variable in our model: the first is the number of war-related fatalities in Ukraine. The second is the number of fatalities relative to the number of war-related events in Ukraine. The third is the number of battles. These series are taken from the Ukraine Conflict Monitor database.¹² As a fourth indicator, we use the number of tweets on a given day. More news about actual events on the battlefield or political decisions in Ukraine and abroad should be reflected in a higher number of tweets.

In Figure (5), we show the responses of selected variables to an uncertainty shock when controlling for each of these war indicators in turn. In order to save space, we do not report the responses of all survey series. Controlling for the number of tweets seems to have the largest effect compared to the other three indicators. In each case, our main findings do not change.

In an alternative specification of the model, we use the daily Geopolitical Risk

¹²See <https://acleddata.com/ukraine-conflict-monitor/> and Raleigh et al. (2023).

Figure 5: Response to an uncertainty shock: controlling for war events



Notes: The figure shows the response to an uncertainty shock of one standard deviation in size. The model includes the number of fatalities (blue), the number of fatalities relative to the number of war-related events (green), the number of battles (pink) or the number of tweets (orange) as an additional control variable. The shaded areas reflect the 68% and 90% confidence bands based on robust standard errors.

Index (Caldara and Iacoviello, 2022) as the driving variable instead of the Ukraine Uncertainty Index.¹³ Specifically, we replace our uncertainty shock unc_t by the standardized residual of a regression of the daily (global) GPR index on one lag of itself and re-estimate all responses. The results are shown in Figures (6) and (7). Again, higher geopolitical risk causes firms to revise their price-setting plans upward and to cut their planned fixed costs, dividends and R&D expenditures. Firms also expect lower revenues, profits, and investments. These responses are qualitatively very similar to our baseline findings and tend to be even more clear-cut. Hence, our results are robust with respect to the measurement of geopolitical uncertainty.

4.3 State-dependent effects of uncertainty shocks

It is plausible that the impact of uncertainty on firms changes over time. One reason is that firms simply are initially very sensitive to the new circumstances, but

¹³This series is available here: <https://www.matteoiacoviello.com/gpr.htm>.

over time adjust to the changing geopolitical landscape. Likewise, economic policy might adjust and stabilize the consequences of war-related shock to some extent. In fact, the German government took several measures to address the consequences of sky-rocketing gas prices for households and firms, among them the gas price brake announced by a government appointed commission in October 2022. These factors suggest that a shock on one standard deviation might not always be equally effective in driving managerial views.¹⁴

We now take account of these factors and allow the uncertainty shock to have different consequences depending on the prevailing level of geopolitical risk. Our prior is that uncertainty shocks have a stronger impact if the level of geopolitical risk is already high. Following the work of Auerbach and Gorodnichenko (2013) and Tenreyro and Thwaites (2016), we can easily generalize the model to a state-dependent local projections approach, which nests the model presented before as a special case. We estimate

$$y_{t+h} = F(\widehat{GPR}_{t-1}) [\alpha_h^I + \beta_h^I unc_t + \gamma_h^I X_t] + (1 - F(\widehat{GPR}_{t-1})) [\alpha_h^{II} + \beta_h^{II} unc_t + \gamma_h^{II} X_t] + u_{t+h}, \quad (2)$$

with separate coefficients in states I and II. The response to uncertainty in state I(II) is β_h^I (β_h^{II}). Without a distinction of two states, the model would collapse to the linear model from the previous subsection.


The transition function $F(\widehat{GPR}_{t-1})$ is between zero and one and indicates the state of the system. Thus, the linear model is a weighted average of the two states, with the weights being the values of the transition function. This function maps the standardized geopolitical risk index, $\widehat{GPR}_t$, into the interval between zero and one

$$F(\widehat{GPR}_{t-1}) = \frac{\exp(\theta \widehat{GPR}_{t-1})}{1 + \exp(\theta \widehat{GPR}_{t-1})}, \quad (3)$$

where θ determines how abruptly the system shifts between the states. We choose $\theta = 3$, which is consistent with the large literature on state-dependent local projections. We use the standardized seven-day backward moving average of the geopolitical risk index as the state variable. The raw index is very volatile and not well suited to indicate different states of risk. Figure (9) shows the transition function $F(\widehat{GPR}_t)$, i.e., the probability of being in the high geopolitical risk state. The probability rises before the full invasion and equals one on February 24, 2022. It falls in the summer of 2022 but increases again in the fall and winter of 2022.

¹⁴Such a non-linear impact of geopolitical risk is also supported by Brignone et al. (2024).


Figure 6: Response of business plans to an uncertainty shock (GPR)



Notes: The figure shows the response to an uncertainty shock one standard deviation in size. The shaded areas reflect the 68% and 90% confidence bands based on robust standard errors.


The impulse responses conditional on the state of geopolitical risk are shown in Figure (8). To save space, we report these responses only for four key variables from the survey. A sudden increase in uncertainty prompts firms to revise their price-setting plans if geopolitical risk is high. If risk is low, however, the response of the price-setting plans remains insignificant. We find a similar pattern for the other variables: if geopolitical risk is already elevated, higher uncertainty causes a strong and highly significant fall in expected dividends, revenues, and investment. Under low geopolitical risk, these responses are not different from zero or, as in the case of expected revenues, even have the opposite signs. To summarize, we find that the adverse effects of war-related uncertainty stem from the weeks and months characterized by very high geopolitical risk. The same increase in uncertainty about the war in Ukraine has no significant effects if it occurs against the backdrop of a lower level of geopolitical risk.

Figure 7: Response of expectations to an uncertainty shock (GPR)




Notes: The figure shows the response to an uncertainty shock one standard deviation in size. The shaded areas reflect the 68% and 90% confidence bands based on robust standard errors.

Figure 9: Probability of high geopolitical risk state



Notes: The figure shows the probability of state I, the high geopolitical risk state. The vertical line marks full invasion on February 24, 2022.

Figure 8: State-dependent responses to an uncertainty shock



Notes: The figure shows the state-dependent response to an uncertainty shock one standard deviation in size conditional on high and low levels of geopolitical risk. The shaded areas reflect the 68% and 90% confidence bands based on robust standard errors.

5 Conclusions

This paper studied how German managers adjust their plans and expectations in response to shocks related to war-related uncertainty. High-frequency data from the *German Business Panel* allows us to track the daily responses to the uncertainty about the war in Ukraine, which we measure from more than eight million German-language Twitter messages related to the war. The combination of survey responses and the indicator of uncertainty, both measured at a high frequency, sheds light on the transmission of news into German boardrooms.

After a shock to uncertainty, managers revise their planned price adjustment upwards and cut expected revenues, dividends, profits, and investments. Thus, the shock has strong adverse effects that resemble a classical aggregate supply shock. We also find that the expected survival rate of firms falls only briefly and that satisfaction with economic policy increases.

The findings underline the notion that the unfolding war in Ukraine is a major driving force of the German business cycle. In addition, the persistent fall in R&D

expenditure suggests that the shock could have consequences for long-run productivity growth.

Due to limited data, we cannot examine sectoral differences across managers' responses or distinguish the responses of managers from small and large firms. This data is generally available in the *German Business Panel*, though the number of daily responses is too small to use for our purposes. A useful extension could be to study the sensitivity of firms to selected geopolitical shocks by making use of the full cross-sectional information available in the survey.

References

- Altig, D., Baker, S., Barrero, J. M., Bloom, N., Bunn, P., Chen, S., Davis, S. J., Leather, J., Meyer, B., Mihaylov, E., Mizen, P. and Parker, N. (2020). Economic Uncertainty Before and During the COVID-19 Pandemic, *Journal of Public Economics* **191**: 104274.
- Auerbach, A. J. and Gorodnichenko, Y. (2013). Fiscal multipliers in recession and expansion, in A. Alessina and F. Giavazzi (eds), *Fiscal Policy after the Financial Crisis*, Chicago University Press, Chicago.
- Bachmann, R., Elstner, S. and Sims, E. S. (2013). Uncertainty and economic activity: evidence from business survey data, *American Economic Journal: Macroeconomics* **5**(2): 217–249.
- Baker, S. R., Bloom, N. and Davis, S. J. (2016). Measuring Economic Policy Uncertainty, *The Quarterly Journal of Economics* **131**(4): 1593–1636.
- Bischof, J., Doerrenberg, P., Rostam-Afschar, D., Simons, D. and Voget, J. (2024). The German Business Panel: Firm-Level Data for Accounting and Taxation Research, *European Accounting Review* pp. 1–29.
- Bloom, N. (2009). The impact of uncertainty shocks, *Econometrica* **77**: 623–685.
- Bloom, N. (2014). Fluctuations in Uncertainty, *Journal of Economic Perspectives* **28**: 153–176.
- Bobasu, A. and De Santis, R. (2022). The impact of the Russian invasion of Ukraine on euro area activity via the uncertainty channel, *ECB Economic Bulletin* **4/2022**.
- Bondarenko, Y., Lewis, V., Rottner, M. and Schöler, Y. (2024). Geopolitical risk perceptions, *Journal of International Economics* **152**: 104005.
- Brignone, D., Gambetti, L. and Ricci, M. (2024). Geopolitical risk shocks: when the size matters, *Working Paper no. 2972*, European Central Bank.
- Caldara, D., Conlisk, S., Iacoviello, M. and Penn, M. (2026). Do geopolitical risks raise or lower inflation?, *Journal of International Economics* **159**: 104188.
- Caldara, D. and Iacoviello, M. (2022). Measuring geopolitical risk, *American Economic Review* **112**: 1194–1225.
- Cascaldi-Garcia, D., Sarisoy, C., Londono, J. M., Sun, B., Datta, D. B., Ferreira, T., Grishenko, O., Jahan-Parvar, M. R., Loria, F., Ma, S., Rodriguez, M., Zer, I. and Rogers, J. (2023). What is certain about uncertainty?, *Journal of Economic Literature* **61**: 624–654.
- Castelnuovo, E. (2023). Uncertainty before and during COVID-19: a survey, *Journal of Economic Surveys* **37**(3): 821–864.


- Dräger, L., Gründler, K. and Potrafke, N. (2025). Political shocks and inflation expectations: Evidence from the 2022 russian invasion of ukraine, *Journal of International Economics* **153**: 104029.
- Federle, J., Meier, A., Müller, G. J. and Sehn, V. (2026). Proximity to war: The stock market response to the russian invasion of ukraine, *Journal of Money, Credit and Banking* **58**(3): 681–703.
- Grebe, M., Kandemir, S. and Tillmann, P. (2024). Uncertainty about the war in Ukraine: Measurement and effects on the German economy, *Journal of Economic Behavior and Organization* **217**: 493–506.
- Grimme, C. (2017). Messung der Unternehmensunsicherheit in Deutschland – das ifo Streuungsmaß, *ifo Schnelldienst* **70**(15): 19–25.
- Hack, L. and Rostam-Afschar, D. (2024). Understanding firm dynamics with daily data, *unpublished*, University of Mannheim.
- Hassan, T., Hollander, S., van Lent, L. and Tahoun, A. (2019). Firm-level political risk: measurement and effects, *Quarterly Journal of Economics* **134**: 2135–2202.
- Hassan, T., Hollander, S., van Lent, L. and Tahoun, A. (2024). The global impact of Brexit uncertainty, *The Journal of Finance* **79**: 413–458.
- Jordà, O. (2005). Estimation and inference of impulse responses by local projections, *American Economic Review* **95**: 161–182.
- Jordà, Ò. and Taylor, A. M. (2025). Local projections, *Journal of Economic Literature* **63**(1): 59–110.
- Jurado, K., Ludvigson, S. C. and Ng, S. (2015). Measuring uncertainty, *American Economic Review* **105**(3): 1177–1216.
- Leduc, S. and Liu, Z. (2016). Uncertainty shocks are aggregate demand shocks, *Journal of Monetary Economics* **82**: 20–35.
- Ludvigson, S. C., Ma, S. and Ng, S. (2021). Uncertainty and business cycles: exogenous impulse or endogenous response?, *American Economic Journal: Macroeconomics* **13**(4): 369–140.
- Raleigh, C., Kishi, R. and Linke, A. (2023). Political instability patterns are obscured by conflict dataset scope conditions, sources, and coding choices, *Humanities and Social Sciences Communications* **10**(1): 1–17.
- Sauer, S. and Wohlrabe, K. (2022). Auswirkungen des Krieges in der Ukraine auf deutsche Unternehmen: Unsicherheit, Lieferengpässe und Preisanstiege, *ifo Schnelldienst* **75**: 15–18.
- Tenreyro, S. and Thwaites, G. (2016). Pushing on a string: U.S. monetary policy is less powerful in recessions, *American Economic Journal Macroeconomics* **8**: 43–74.
- Wang, X., Wu, Y. and Xu, W. (2024). Geopolitical risk and investment, *Journal of Money, Credit and Banking* **56**(8): 2023–2059.

Appendices

Appendix A Comparing the Ukraine Uncertainty Index

The ifo Survey, conducted monthly by the Munich-based ifo Institute, is a very influential source of information about firm-level sentiment. However, its monthly frequency makes it less suitable for the analysis conducted in this paper. Nevertheless, it is instructive to compare our indicator of uncertainty, the Ukraine Uncertainty Index, to the uncertainty reflected in the ifo Survey. Specifically, we use the ifo dispersion measure (Grimme, 2017), which summarizes information about the dispersion of firms' expectations at each point in time. Figure (A.1) shows that both measures strongly increase around the time of the Russian invasion. Overall, both series exhibit roughly the same evolution over time.

Figure A.1: The Ukraine Uncertainty Index against the ifo Uncertainty Index




Notes: The figure shows the standardized daily Ukraine Uncertainty Index derived from German tweets against the monthly ifo dispersion measure (in index points, right scale).

Appendix B Additional control variables

In the main text, we include four alternative control variables which reflect the time paths of war-related events. The number of tweets is a natural by-product of our data set. The three series on total fatalities, the fatalities per event and the number of battles are taken from the Ukraine Conflict Monitor database (Raleigh et al., 2023). Figure (B.2) shows the evolution of all four additional control variables.

Figure B.2: Additional control variables reflecting war-events




Notes: The graph shows the four alternative control variables included in the estimation. Each variable is a proxy for the war-related events. The dashed vertical line indicates the begin of the full invasion on February 24, 2022.

Appendix C Dependent variables

Figure (C.3) plots the evolution of the first group of variables reflecting managers' plans. The second group, reflecting expectations and the assessment of policy, is shown in Figure (C.4). Each series represents the 28-day weighted moving averages of the average of the firm responses, with the weights being the relative number of responses on each day. It is particularly remarkable that the drop in firms' planned dividends, the fall in expected employment, and the decline in expected profits

coincide with the invasion on February 24, 2022. Likewise, we observe an increase in firms' plans to raise prices after this date. Firms' satisfaction with policy sharply deteriorates after the full invasion but recovers in the fall of 2022, likely as a result of the German government's announcement of the gas price brake in October 2022. Figure (C.5) compares four selected series with the monthly index of business expectations from the ifo Survey. At the beginning of the sample period, our data tracks the ifo expectations quite well. However, after August 2022, managers became more optimistic in our data compared to the ifo survey.

Figure C.3: Dependent variables: managers' plans




Notes: The graph shows the 28-day weighted moving averages of the daily average responses across firms. The weights are the relative number of responses on each day. The dashed vertical line indicates the begin of the full invasion on February 24, 2022.

Appendix D Number of survey responses

Figures (D.6) and (D.7) depict the average number of responses in our sample for each day of the week. The number of responses to each survey question is fairly stable between Monday and Friday, with a small peak on Wednesday. Over the weekend, the number of responses falls drastically to only single-digit numbers. Note that we address this by weighting the 28-day (and 14-day) moving averages of the responses by the relative number of responses on a given day.

Figure C.4: Dependent variables: managers' expectations and assessments




Notes: The graph shows the 28-day weighted moving averages of the daily average responses across firms. The weights are the relative number of responses on each day. The dashed vertical line indicates the begin of the full invasion on February 24, 2022.

Finally, we assess whether the number of responses to the individual survey questions is affected by the uncertainty shock. Specifically, we investigate whether a higher level of uncertainty prompts fewer or more firms to respond to the German Business Panel. Figures (D.8) and (D.9) show the estimated impulse response functions. We find that the number of responses drops on the first day of the shock. For almost all survey questions, a shock of one standard deviation reduces the number of responses by one. However, this effect is very short-lived. There is no significant increase or decrease in the number of responses beyond the first day.

D.1 Alternative purification of shock

Our baseline uncertainty shock is the residual of a regression of the Ukraine Uncertainty Index on one lag of itself, a constant, and one dummy for each day of the week. An alternative shock is the residual of the regression of the Ukraine Uncertainty Index on six lags of itself and a constant. In this alternative specification, we include additional lags and drop the day-of-the-week dummies. Figures (D.10) and (D.11) plot the estimates β_n coefficients based on this alternative shock series. The results remain broadly similar to our baseline findings. A notable exception is the response of managers' expected survival rate, which falls more persistently than in the baseline

Figure C.5: Comparing selected variables to the ifo index




Notes: The graph shows the 28-day weighted moving averages of the daily average responses across firms. The weights are the relative number of responses on each day. The dashed vertical line indicates the begin of the full invasion on February 24, 2022. The orange line is the monthly index of business expectations (index points, right scale) from the ifo Survey.

model.

D.2 Adjusting the shock series for outliers

The series of uncertainty shocks used in the estimation exhibits several obvious outliers. The clearest case is the realization on February 24, 2022, the day of the full Russian invasion, where the shock is more than six standard deviations above its sample mean. We now assess the robustness of our estimated impulse response functions with respect to deleting outliers. We set all realizations of the uncertainty shock greater than four or three standard deviations (in absolute terms) to zero and re-estimate the baseline model. Figure (D.12) shows the resulting impulse responses of selected variables. All of our main results remain qualitatively unchanged.

Figure D.6: Number of survey responses (part I)



Notes: Average number of survey responses for each day of the week.

Appendix E Survey Questions

We provide the exact formulation of the survey questions from the *German Business Panel* and their English translation. Both are taken from Hack and Rostam-Afschar (2024).

E.1 Q1: Plan to Change Prices

Original: Welche Entscheidungen planen Sie in den nächsten 12 Monaten?

1. Absatzpreise: Reduzierung
2. Absatzpreise: Keine Änderung
3. Absatzpreise: Erhöhung

Translation: What decisions are you planning to make in the next 12 months?

1. Decrease sales prices
2. No change in sales prices
3. Increase sales prices

Figure D.7: Number of survey responses (part II)



Notes: Average number of survey responses for each day of the week.

E.2 Q2: Plan to Change Fixed Costs

Original: Welche Entscheidungen planen Sie in den nächsten 12 Monaten?

1. Fixkosten (z.B. festes Personal, Miete): Reduzierung
2. Fixkosten (z.B. festes Personal, Miete): Keine Änderung
3. Fixkosten (z.B. festes Personal, Miete): Erhöhung

Translation: What decisions are you planning to make in the next 12 months?


1. Decrease fixed cost (e.g., permanent personnel, rent)
2. No change in fixed cost (e.g., permanent personnel, rent)
3. Increase fixed cost (e.g., permanent personnel, rent)

E.3 Q3: Plan to Change Dividends

Original: Welche Entscheidungen planen Sie in den nächsten 12 Monaten?

1. Entnahmen von Unternehmer/Gesellschafter; Ausschüttungen an Anteilseigner:
Reduzierung
2. Entnahmen von Unternehmer/Gesellschafter; Ausschüttungen an Anteilseigner:
Keine Änderung

Figure D.8: Response of the number of survey responses to an uncertainty shock (I)



Notes: The figure shows the response to an uncertainty shock one standard deviation in size. The shaded areas reflect 68% and 90% confidence bands based on robust standard errors.

3. Entnahmen von Unternehmer/Gesellschafter; Ausschüttungen an Anteilseigner:
Erhöhung

Translation: What decisions are you planning to make in the next 12 months?

1. Decrease withdrawals of entrepreneurs/owners, or distributions paid to shareholders
2. No change in withdrawals of entrepreneurs/owners, or distributions paid to shareholder
3. Increase withdrawals of entrepreneurs/owners, or distributions paid to shareholder

E.4 Q4: Plan to Change R&D

Original: Welche Entscheidungen planen Sie in den nächsten 12 Monaten?

1. Forschung und Entwicklung: Reduzierung

Figure D.9: Response of the number of survey responses to an uncertainty shock (II)



Notes: The figure shows the response to an uncertainty shock one standard deviation in size. The shaded areas reflect the 68% and 90% confidence bands based on robust standard errors.

2. Forschung und Entwicklung: Keine Änderung
3. Forschung und Entwicklung: Erhöhung

Translation: What decisions are you planning to make in the next 12 months?

1. Decrease R&D investment
2. No change in R&D investment
3. Increase R&D investment

E.5 Q5: Expected Survival Rate


Original: Was schätzen Sie: Wie viel Prozent der Unternehmen Ihrer Branche werden ihr Geschäft in den nächsten 12 Monaten aufgeben?

Prozentzahl zwischen 0 und 100.

Translation: What do you estimate: What percentage of firms in your industry will go out of business in the next 12 months?

Percentage number between 0 and 100

Figure D.10: Response of plans to an uncertainty shock: alternative shock



Notes: The figure shows the response to an uncertainty shock one standard deviation in size. The shaded areas reflect the 68% and 90% confidence bands based on robust standard errors.

E.6 Q6: Satisfaction with Policy

Original: Wie zufrieden sind Sie mit der Wirtschaftspolitik in Deutschland?

Zufriedenheit in 11 diskreten Möglichkeiten von “sehr unzufrieden” (0) bis “sehr zufrieden” (10).

Translation: How satisfied are you with economic policy in Germany?


Satisfaction in 11-point Likert scale from “very dissatisfied” (0) to “very satisfied” (10).

E.7 Q7-10: Expected Change in Revenue, Profit, Employment and Investment

Original: Welche Veränderung (in %) der folgenden Kennzahlen erwarten Sie für das laufende Kalenderjahr für Ihr Unternehmen im Vergleich zum vorangegangenen Kalenderjahr?

- Jahresumsatz $\in [-100, 100]$

Figure D.11: Response of expectations to an uncertainty shock: alternative shock




Notes: The figure shows the response to an uncertainty shock one standard deviation in size. The shaded areas reflect the 68% and 90% confidence bands based on robust standard errors.

- Beschäftigte $\in [-100, 100]$
- Investitionen $\in [-100, 100]$
- Jahresgewinn $\in [-100, 100]$

Translation: What change (in %) in the following key figures do you expect for your firm in the current calendar year compared to the previous calendar year?

- Revenue $\in [-100, 100]$
- Number of employees $\in [-100, 100]$
- Investment $\in [-100, 100]$
- Profit $\in [-100, 100]$

Figure D.12: Response to an uncertainty shock: excluding outliers



Notes: The figure shows the response to an uncertainty shock of one standard deviation in size. The orange (grey) line represents the response when all shocks larger than four (three) standard deviations are set to zero. The shaded areas reflect the 68% and 90% confidence bands based on robust standard errors.

Essay III

Monetary Metanomics: The Effect of Monetary Policy on Virtual Real Estate Prices

This paper was presented at the following workshops and international conferences:

- i. 4th MORSE Conference (Maastricht)
Date: October 2024
Presenter: Salah Hassanin
- ii. ECONDAT spring meeting: Economics with nontraditional data and analytical tools (London)
Date: June 2025
Presenter: Salah Hassanin (Poster session)

Acknowledgement

We would like to thank the marketplaces OpenSea, X2Y2, Rarible, LooksRare, and Reservoir for providing free API keys and special API rights for research purposes. We are grateful to Peter Tillmann, David Finck, Daniel Grabowski, Lucas Hafemann and Malik Hebbat for many helpful comments and suggestions. Seminar participants at the University of Giessen, the 4th MORSE Conference (Maastricht, 2024) and ECONDAT Spring Meeting (London, 2025) provided valuable feedback. David Finck shared his fantastic local projections toolbox. Last but not least, we thank Yudhy Darma for accompanying us on our exploration of Decentraland.

Monetary Metanomics: The Effect of Monetary Policy on Virtual Real Estate Prices

Moritz Grebe *

Salah Hassanin †

Abstract

This paper examines the impact of monetary policy on virtual real estate prices within the Decentraland metaverse. Using transaction data from various blockchain marketplaces, we construct a daily Decentraland Price Index (DPI). Applying state-dependent local projections, we show that the effect of the Federal Reserve’s monetary policy depends on user activity. In high-activity states, a 25 basis point monetary policy shock leads to an immediate 20% DPI drop, with the effect peaking at 37% by the third day. Our results demonstrate the sensitivity of virtual real estate markets to monetary policy and highlight the economic parallels between virtual and physical real estate.

Keywords: metaverse, Decentraland, virtual real estate, hedonic price index, monetary policy, price dynamics, user activity, blockchain data

JEL Classification: E52, R31, G12

*Justus Liebig University Giessen, Email: moritz.grebe@wirtschaft.uni-giessen.de

†Justus Liebig University Giessen, Email: salah.hassanin@wirtschaft.uni-giessen.de

1 Introduction

The concept of a metaverse existed before Facebook’s rebranding to Meta popularized it, yet this event marked the beginning of a significant surge in interest and adoption. The metaverse is a vast interconnected network of 3D virtual worlds designed for synchronous, persistent experiences by countless users, characterized by real-time rendering, a sense of individual presence, and continuity of data such as identity, ownership, history, and transactions, while aiming for seamless interoperability across platforms (Ball, 2022). Decentraland, one of the most prominent platforms in this space, represents an early-stage realization of the metaverse paradigm.¹ It is described as "paving the way to an open, interconnected metaverse of diverse virtual worlds" (Ordano et al., 2024, p. 4). Powered by the Ethereum blockchain, Decentraland integrates the principles of decentralization and offers users a unique and evolving virtual environment (Ordano et al., 2024). In alignment with the mentioned metaverse features in Dwivedi et al. (2022) and Park and Kim (2022), users can interact through avatars in social activities, attend community-driven events, engage in gaming, pursue 3D-based learning, and create and monetize content. This environment emphasizes user governance and ownership, blurring the lines between physical and digital realms and introducing novel economic dynamics.

Although the initial hype has subsided, the metaverse’s significant future potential remains a subject of rigorous academic scrutiny. Specifically, very little is known about the economics of metaverses (Goldberg et al., 2024), termed metanomics, and even less about its relationship with the real-world economy.

In this paper, we therefore examine the effect of monetary policy on prices in the metaverse. More specifically, we look at the prices of virtual real estate in Decentraland. These prices have an important characteristic that distinguishes them from many other virtual assets and makes them comparable to their real-world

¹Ball (2022) includes the following metaverse key attributes: massive scale and interoperability, real-time rendering, 3D virtual worlds, synchronicity and persistence, effectively unlimited users, individual sense of presence, and continuity of data. Decentraland aligns almost with all those key attributes, except for individual sense of presence and interoperability. In terms of individual sense of presence, Decentraland provides a visually and audibly engaging experience using 3D environments and spatial audio. However, it relies on local software clients (standard screens), standard browsers, and basic sound design. It lacks deeper sensory engagement, such as virtual reality or haptic feedback, which are crucial for achieving the individual sense of presence. Regarding interoperability, Decentraland’s blockchain foundation enables asset ownership and transferability. While this aligns partially with Ball’s vision of a massively interconnected metaverse, the platform’s practical interoperability remains limited to blockchain-compatible systems, falling short of the seamless asset and identity mobility he describes. Nonetheless, while it is not the full embodiment of Matthew Ball’s idealized metaverse, it fits enough key attributes to be considered part of the broader metaverse ecosystem. Therefore, it is valid and reasonable to regard Decentraland as a metaverse within the current landscape.

counterparts: they can not only be bought and sold, but also rented to users. If we understand the price of a property as the present value of future rental payments, then monetary policy could influence prices via the discount rate. An unexpected rise in interest rates would then lead to a fall in prices, as would be expected in the real world.

To test this, we first need information about Decentraland’s virtual land, such as prices and location. This virtual world comprises 90,601 land parcels, approximately 40,000 of which are privately owned. Owners of private parcels can modify their content, merge them with other parcels, rent them out, or offer the parcel for sale. The open architecture of the underlying blockchain offers a transparent dataset of virtual land information. Every economic activity is recorded on the blockchain and consequently can be tracked.

We collect relevant on-chain transaction data for parcel sales and transfers in Decentraland from its native marketplace, various other blockchain marketplaces (OpenSea, X2Y2, Rarible, and LooksRare), and cross-reference it with the data aggregation service Reservoir to ensure data integrity. This dataset includes detailed information on buyers, sellers, daily sales prices, transaction dates, and unique parcel identifiers. Additionally, we gather metadata for each parcel, such as coordinates, adjacency to roads and districts, and whether the coordinates are easy to remember. Combining all successful sales into our dataset, we end up with 24,237 successful sales transactions between March 20, 2018, and September 20, 2023. The average transaction price is \$7,207.46, with the highest reaching \$405,406.40.

We then use these transactions to calculate the daily Decentraland Price Index (DPI), based on the hedonic price index model methodology used for the German Real Estate Index (Amaral et al., 2023). This approach is particularly suited for this analysis as it allows us to decompose the price of virtual land into constituent attributes that affect its value, such as location within the virtual world, proximity to roads and plazas, ease of remembering parcel coordinates, and adjacency to prominent districts. This enables us to isolate the effect of each characteristic on land prices, ensuring that our analysis accurately reflects the value drivers in this unique market and estimate a quality-adjusted price index.

To estimate the effect of monetary policy, we update the time series of Fed monetary policy surprises from Bu et al. (2021) and employ local projections (Jordà, 2005). We focus on the Federal Reserve’s monetary policy as it has a significant impact on global financial markets due to the dominant role of the U.S. dollar in international trade and finance. Its decisions are closely monitored worldwide which makes it a relevant benchmark for analyzing monetary policy impacts, even in decentralized and globally

distributed platforms like Decentraland. We observe a negative effect on prices: after a monetary policy shock of 25 basis points, the DPI falls by approximately 20%. However, this substantial effect is only statistically significant on the fourth day following the shock.

User activity on Decentraland has seen notable fluctuations since its inception. Initially, the platform was relatively unknown, with only a few users interacting and trading land parcels. During the height of its popularity, both user numbers and transaction volumes surged significantly, only to decline once the hype subsided. We argue that the level of user activity, whether high or low, impacts market dynamics. One might expect that the more users engage with Decentraland, the more prices should react as in the real world. However, if user activity is low and the majority of users are high-risk investors accustomed to investing in virtual assets, the effect is less clear.

We use state-dependent local projections and identify states with high and low user activity using the approach of Born et al. (2020). In the state with high user activity prices fall by 15% on the day of the shock. The largest drop occurs on day three, with around 37%. After day 4 the effect gets insignificant. In the state with low activity, we see a positive effect on the first and second day after the shock.

We also update the shock series of Ciminelli et al. (2022) and test whether Decentraland prices react to central bank information shocks which could occur when the central bank releases unknown information about the economic outlook during its meetings.

In summary, we can show that the prices of virtual land in Decentraland actually respond to the Fed's monetary policy. When activity and popularity in the virtual world are high, prices respond to monetary policy and information shocks in a manner that closely resembles real-world market reactions. However, important differences emerge: in the metaverse, prices adjust much more rapidly—on impact following a pure monetary policy shock or within a few days after an information shock. Moreover, these effects fade relatively quickly, usually within a week. Lastly, responses can be quite pronounced, with even stronger reactions to information shocks.

Our paper contributes to different strands of the literature. First, our paper adds to the literature by advancing the understanding of economic principles governing virtual real estate markets, particularly within metaverse platforms like Decentraland. Recent research on virtual real estate in the metaverse has yielded several important insights relevant to pricing and economic dynamics within these platforms. Studies indicate that virtual land prices often display inefficiencies and high volatility. Yet,

they have shown a consistent upward trend, driven by speculative behavior and the unique characteristics of NFTs (Dowling, 2022a,b). There is also evidence of limited but significant co-movement between virtual land prices and cryptocurrency markets, particularly with Ether, highlighting a potential link between these digital asset classes (Dowling, 2022b; Nakavachara and Saengchote, 2022). The value of virtual land is further influenced by location, similar to physical real estate, where proximity to key virtual landmarks significantly affects prices (Goldberg et al., 2024; Güven and Ercan, 2022). Additionally, the medium of exchange used, such as cryptocurrency versus fiat currency, can impact the willingness to pay for virtual land, suggesting that monetary policies could influence these emerging digital markets (Nakavachara and Saengchote, 2022; Yench, 2023). These findings underscore the need for a nuanced understanding of virtual real estate, particularly as it relates to broader economic and monetary policy considerations in digital economies.

Second, we contribute to the literature investigating the relationship between monetary policy and asset prices. House prices, in particular, are seen as key elements of the business cycle. Especially since the house price boom and the subsequent financial crisis, understanding the interaction between monetary policy and house prices has become increasingly important (Del Negro and Otrok, 2007; Iacoviello, 2005). Jordà et al. (2015) find a crucial role of loose monetary policy in house price booms and argue, just like Kiley (2021), that house price dynamics can be predictive of crises.

Overall, it seems that monetary policy has a large effect on house prices. For instance, Benati (2021) find a negative effect of a positive monetary policy shock on real house prices that is three to five times larger than the effect on real GDP. Furthermore, the effects seem persistent, and prices do not react on impact but usually only after about a year (Goodhart and Hofmann, 2008; Williams, 2015; Hülsewig and Rottmann, 2021; Xie, 2024). Xie (2024) attributes this "sluggishness" to the information channel, arguing that households are not informed about current changes in monetary policy and only become aware of them much later through news about changes in house prices.

The reaction of house prices, however, is heterogeneous across the countries considered, and seems to depend on the mortgage markets. Calza et al. (2013) and Benati (2021) find that house prices react more strongly when mortgage markets are more developed and flexible. Paul (2020) finds that the effects are stronger when house prices are high and less strong when they are low.

This paper is structured to first provide an overview of Decentraland and its virtual real estate market, followed by the methodology used to construct the Decentraland

Price Index (DPI) from on-chain transaction data. It then analyzes the effects of Federal Reserve monetary policy shocks on virtual real estate prices using state-dependent local projections, with a focus on user activity. The paper concludes with a discussion of the findings’ implications for stakeholders in the metaverse and robustness checks to validate the results.

2 Decentraland

Decentraland, launched as beta and raised more than \$24 million in its initial coin offering (ICO) in 2017 (ICO Drops, 2017) and opened to the public in 2020 (Decentraland, 2020), is a pioneering metaverse platform that amalgamates the principles of decentralization and virtual reality, powered by the Ethereum (ETH) blockchain. Ball’s (2022) metaverse definition of the metaverse applies to Decentraland: Decentraland is a vast interconnected network of 3D virtual worlds designed for synchronous, persistent experiences by countless users. In alignment with the mentioned metaverse features in Dwivedi et al. (2022) and Park and Kim (2022), users can engage through their avatars in social activities, attend community-driven events, play games, educate themselves, explore the virtual world, and create, experience, and monetize content (e.g., avatar fashion and art) and applications (e.g., gambling) in a 3D immersive environment. By extending the social significance, the metaverse integrates with the real world (Park and Kim, 2022). Moreover, Decentraland expands human activities by transcending the natural limitations of space, time, and resources, a characteristic of the metaverse described by Momtaz (2022).

Inspired by Web 3.0, the metaverse Decentraland operates within a decentralized metaverse model. This model represents a paradigm shift towards a model where control and ownership are distributed among its users rather than being held by a single entity (Cantú et al., 2024). This virtual world is built on blockchain technology, utilizing fungible tokens (digital currency), non-fungible tokens (NFTs), and DAO, which empower users to govern without a centralized authority. These technologies collectively enable an open, interoperable, and user-governable digital universe (Huynh-The et al., 2023; Despotović et al., 2023; Goldberg and Schär, 2023). Consequently, users in Decentraland have complete control over the land they own, the assets they create, and their transactions, contributing to an ecosystem where they are major stakeholders and beneficiaries. Additionally, they support a personal sense of presence and maintain data continuity across identity, history, entitlements, objects, communications, and payments. This model not only overcomes the limitations of centralized systems, such as lack of trust and potential censorship

(Goldberg and Schär, 2023), but also fosters the creation of virtual environments that are equitable, secure, and user-driven. The decentralized model places the responsibility of privacy and security on the individual users, lacking a centralized body to address issues like abuse or enforce policies. It places the responsibility of privacy and security on individual users, as there is no centralized body to address issues like abuse or enforce policies, thereby laying the foundation for an open digital economy.

Fungible and non-fungible tokens play a crucial role in this virtual world by representing Decentraland’s currency, MANA, and a variety of virtual assets. MANA follows like ETH the Ethereum Request for Comments 20 (ERC-20) standard, allowing for standardized transactions and economic activities. ERC-20 tokens are identical and can be exchanged on a one-to-one basis. NFTs include unique land parcels (LAND), estates, wearable items for avatars, unique avatar names, and digital art and collectibles, all of which can be bought, sold, or traded on Decentraland’s marketplace or alternative marketplaces, such as OpenSea. In the Ethereum blockchain, NFTs follow the ERC-721 standard. They are distinctive assets, each with unique identification codes and metadata that set them apart. They cannot be replicated or divided like ERC-20 tokens. This uniqueness allows NFTs to represent ownership or proof of authenticity.

Alongside the decentralized metaverse The Sandbox (The Sandbox, 2020), Decentraland (Ordano et al., 2024) represents the forefront of integrating virtual economies with blockchain technology. The two most powerful platforms utilize their MANA and SAND tokens for transactions and governance. These platforms offer unique digital asset experiences, like virtual land (LAND), and empower user participation in platform evolution via DAOs. Both platforms have garnered attention from major brands (The Sandbox, 2024; Decentraland, 2023; Decentraland Fandom, 2024) and celebrities (Hazan et al., 2022), underscoring the real-world relevance of virtual economies. In these worlds, LAND scarcity mirrors the physical world. It ensures no risk of dilution by central authorities and safeguards against unilateral rule changes that could disrupt the ecosystem. Among all NFTs, LAND is not merely a digital asset but a foundation for diverse creative and social endeavors. Their open architecture offers a transparent dataset of – among other things – LAND prices, which makes these metaverses a perfect setting for observable experiments with real monetary stakes. Moreover, both platforms’ parcel size and shape uniformity facilitate comparability and consistency in the virtual real estate market. Compared to The Sandbox, Decentraland stands out for its symmetrical geographic layout, where the parcels are embedded in a strict grid structure. Moreover, as the

pioneering major blockchain-based virtual world, Decentraland potentially provides a richer repository of historical data, offering more profound insights into digital land economics. Furthermore, Decentraland’s focus on a wide array of virtual experiences, including art, events, and fashion - beyond gaming - (Ordano et al., 2024; The Sandbox, 2020) paired with lower transaction fees (2.5% Vs. 5%; Chanda, 2023), positions it as a particularly appealing environment for creators and investors. Given its pioneering status, a wealth of historical data, and a broad spectrum of virtual experiences coupled with economic incentives, Decentraland emerges as the primary focus of our research, promising unique insights into the dynamics of virtual real estate within decentralized metaverses.

From the economic lens, Decentraland embodies, like The Sandbox (Nakavachara and Saengchote, 2022), an open economy characterized by the absence of capital control and monetary sovereignty, and a fully flexible exchange rate, according to the Impossible Trinity (Mundell, 1963; Fleming, 1962).² This is facilitated by the unrestricted convertibility of blockchain tokens that ensures perfect capital mobility, the completely decentralized financial system (DeFi), and the easy virtual economy’s currency exchange by other currencies. The platform’s emphasis on DAO for governance and property rights (LAND ownership) introduces a novel dimension to virtual real estate. This invites an analysis of how decentralized environments might reflect or diverge from traditional real estate markets under varying monetary policies and economic conditions.

The rest of the chapter describes the collected Decentraland data and is accompanied by descriptive statistics and facts. Based on this data, the construction of a hedonic real estate price index and the monetary policy analysis will follow.


2.1 LAND Data

The parcels on the square-shaped Decentraland map are described two-dimensionally by x- and y-coordinates. Decentraland’s basis consists of a central plaza with eight surrounding plazas, all equidistant in a circular pattern (see Figure (1)). The symmetrically structured metaverse consists of 90,601 parcels ($x, y \in [-150; 150]$), comprising 43,689 private parcels, 33,886 parcels within 39 DAO-approved community-built districts, 9,438 road parcels, and 3,588 parcels divided over nine plazas. All parcel sizes are standardized and are 256 square meters (16m x 16m). The owner of a parcel can modify its content (e.g., create a scene or build a house) and merge it

²In this context, ‘open economy’ refers primarily to Decentraland’s integration with the real-world economy. While openness to other metaverses is not excluded, the focus is on its interaction with global economic systems.

with other neighboring parcels to form an estate (and vice versa). In addition, the owner can transfer or rent the parcel to other users, give other users parcel editing rights, name the parcel, and give the parcel a public description.

Figure 1: Decentraland Map



Notes: Decentraland at a glance

We collect on-chain Decentraland parcel transaction data from various marketplaces, covering the period from March 2018 to August 2023, and cross-checked it with Reservoir’s data aggregation service for integrity.³ Afterward, the datasets of the various marketplaces were filtered. For example, some marketplaces only show

³The lack of on-chain Decentraland LAND parcel transaction data before March 2018 is because the Terraform Event, which took place from December 2017 to January 2018, involved off-chain bidding, with the actual on-chain registration and transfer of LAND ownership occurring later. The Decentraland LAND smart contracts were deployed in March 2018, marking the beginning of on-chain transactions. Therefore, on-chain data only became available after the contracts were live and LAND parcels were officially registered on the blockchain. The Decentraland Marketplace was released allowing users to buy and sell parcels (LAND) as well as explore the map.

successful transactions, while others also show failed transactions.⁴ As the various marketplaces use, for example, different terminology for sellers (e.g., *sender*), buyers (e.g., *recipient*), closing time (e.g., *end time*) of a transaction, the datasets are harmonized to achieve consistency. Then, we concatenated and filtered the dataset as we only considered sales for the methodology application. Since transaction payments are accepted in various cryptocurrencies (e.g., Bitcoin and Ethereum) beyond the Decentraland marketplace (MANA), we have standardized all transactions by converting them to daily USD sales prices for consistency and later analysis. Out-of-market trust-based sales are not and cannot be recorded.⁵ They only result in a transfer. However, a transfer can also be a gift or part of an exchange transaction. The final transaction dataset contains the daily USD sales prices, the closing dates (Central Standard Time), and the unique token ID with which the LAND NFTs can be identified. This allows us to add data on the parcels, such as the x-y-coordinates, direct access to roads, distance to a plaza, and others to describe and understand the value of a given parcel.

2.2 LAND Ownership

The decentralization of LAND ownership is a key factor in determining the inclusivity and sustainability of Decentraland’s virtual ecosystem. The pattern of LAND ownership not only reveals the degree of decentralization but also highlights the shifts in control and accessibility within the marketplace. We calculate the Gini coefficient G over time to examine the distribution of land ownership in Decentraland’s virtual real estate market, specifically assessing the decentralization development of LAND ownership. A G close to 0 (1) would indicate that LAND parcels are distributed relatively (un)equally among participants.

The Gini coefficient shows a long-term downward trend with two sharp drops (see Figure (2)). At the publication of Decentraland and public access to parcel trade, the


⁴Failure may be due to fraud attempts, suspicious activity, or insufficient liquidity. A common reason for a failed transaction is gas. A fee (gas) is due for carrying out a transaction on the Ethereum blockchain, which depends on the network demand and transaction complexity. Accordingly, gas varies from the submission of a bid/purchase to the completion of the transaction. Thus, the buyer’s cash balance may be insufficient to carry out the transaction if the gas fee increases unexpectedly during the transaction process. The gas fees are payouts to Proof-of-Work miners (pre-Merge) / Proof-of-Stake validators (post-Merge) on the ETH protocol who support and secure the ETH network. Therefore, gas has an activity-regulating function and maintains the network security and functionality in the ETH blockchain.

⁵Out-of-market trust-based sales refer to sales that occur outside established marketplaces, relying on trust between the parties involved rather than on the security and protocols of the marketplace. They depend on personal agreements, trust, or reputation of the parties to ensure that the terms of the sale are honored. Therefore, the transaction parties save on marketplace fees but bear the risk of fraud.

coefficient falls from around 94% to 83%, and at the end of 2021, when trading activity has increased enormously. This indicates a redistribution of parcels of land across more landowners. Large landowners own fewer parcels. This finding is underpinned by other indicators. In the beginning, the top 5% (top 1%) owned almost 90% (82%) of the parcels, while in 2023, 66% (53%) of the parcels are still owned by them after the drastic drop at the time of publication, and the short-term sharp increase in trading activity (see Figure (3)). The Gini coefficient rose briefly during the period of high market activity, but then returned to the same level and fell even further. This could be due to the active land purchases of the top 5% (1%) largest landowners who sold their land after the peak phase. Over the entire period, the average number of parcels per landowner has fallen from around 38 to 4 (-89.47%), and the number of different LAND owners has risen from 2,059 to 6,595 (+220.30%) (see Figure (4)). The number of properties in circulation has remained roughly the same after the drastic increase from around 12,800 to 65,000 from the launch of Decentraland until the end of the data set time horizon. At the end of the observed time window, the Gini coefficient remains at 70.71%. Compared to other virtual worlds, such as The Sandbox (77% in November 2022; Blockchain Research Lab, 2022), the LAND parcels are less unevenly distributed. Compared to the real world, G lies between the USA (75%) and Germany (68%) (UBS, 2024).⁶ Overall, there has been a continuous long-term progression in decentralizing Decentraland LAND ownership which reflects a shift toward a more equitable distribution of virtual land. This democratizing and maturing market promotes greater accessibility, participation, and the potential for a more stable and sustainable virtual real estate ecosystem.


⁶It is important to note that the Gini coefficients for the USA and Germany in UBS (2024) reflect the distribution of wealth assets, which encompass a broader range of items beyond land alone. While this comparison is not perfectly aligned, it serves as a useful point of reference for understanding the relative inequality in virtual and real-world contexts, given the lack of more directly comparable data. In literature, one directly comparable but outdated source is Deininger and Olinto (2000): Germany Land Gini was at 55.39% and USA 73.10%.

Figure 2: Gini Coefficient




Notes: The calculated Gini coefficient over the entire time horizon based on the collected LAND data. A Gini coefficient G close to 0 (1) would indicate that LAND parcels are distributed relatively (un)equally among participants.

Figure 3: Monthly LAND Ownership among the largest LAND Owners



Notes: The LAND distribution across the top 1% and 5% largest LAND owners.

Figure 4: Monthly Number of LAND Owners & Average LAND Holdings



Notes: Comparison of the number of LAND owners (blue) with the average number of LAND held (green) over the time horizon based on the data collected.

2.3 Internal Factors influencing LAND Prices

In the context of Decentraland, several key factors influence land prices. We explore the spatial distribution of average LAND parcel prices in Decentraland at a descriptive level and compare our findings about the factors affecting land valuations to existing literature on the metaverse and the real world. Analyzing the heatmap (Figure (5)), several patterns emerge:

- *Centrality:* Parcels located near the center of Decentraland (around 0, 0) command higher prices. This central region includes Genesis Plaza, a major hub known for high user traffic, significant landmarks and the starting point for users in Decentraland. The peripheral regions of Decentraland generally exhibit lower average sale prices. These areas tend to have fewer attractions or be less developed, leading to reduced demand. Central locations are more desirable due to their accessibility and visibility.
- *Nearby Landmarks:* Proximity to notable landmarks and attractions, such as Fashion Street, District X (formerly known as Red Light District), Vegas City, Decentraland University, and the Conference Center, enhances parcel value. These areas are vibrant with activities, drawing consistent user engagement and visitors.
- *Community Projects:* The presence of community-driven projects like the

AETHERIAN Project and themed areas such as Dragon Kingdom and Dragon City contribute to higher land values due to their active community involvement and themed events.

- *Roads and Plazas*: Parcels adjacent to major roads and plazas often show higher values. Good connectivity enhances accessibility, making these locations more attractive for development and user visits. Roads and plazas serve as critical infrastructure that improves navigation within Decentraland.


Overall, the spatial distribution of property prices in Decentraland indicates that location, proximity to key amenities, connectivity, and accessibility are influential factors. This insight aligns with previous literature on both virtual worlds (Goldberg et al., 2024) and the real world (Heikkila et al., 1989; Tyrväinen and Miettinen, 2000). For example, Heikkila et al. (1989) demonstrate that residential land values in polycentric urban areas like Los Angeles are significantly influenced by multiple subcenters rather than just the central business districts. Similarly, Tyrväinen and Miettinen (2000) find that proximity to forested areas and having forest views significantly increase property values. However, in Decentraland, the ability to walk or teleport between locations reduces the traditional significance of physical proximity to transport hubs (Bowes and Ihlanfeldt, 2001). Instead, the location value is more closely related to expected visitor density and the potential for random encounters (Goldberg et al., 2024). Consequently, calculations of land prices in Decentraland should focus on the virtual environment’s attributes and strategic positioning relative to popular and high-traffic areas, rather than conventional proximity-distance relationships. Therefore, areas hosting regular events, community gatherings, or collaborative projects may boost (nearby) parcel values. Active community involvement and engagement drive up demand and prices as they attract more visitors and participants. This shift in locational preferences reflects the unique dynamics of virtual worlds, where connectivity and accessibility take on new meanings. Finally, Decentraland’s map infrastructure is polycentric. It is characterized by multiple centers of activity, decentralized ownership, and community-driven development, making it a vibrant and diverse virtual world.

2.4 LAND Market Activity


Market dynamics in the digital land sector reveal significant fluctuations driven by varying levels of activity, influencing both demand and pricing trends over time. The increasing activity and thus the potential demand for parcels, while the number of available parcels remains constant (see Figure (6)), leads to greater spikes in LAND

Figure 5: Decentraland Heatmap of average LAND parcel prices

(a) Decentraland background




(b) POI background



Notes: The heatmaps describe the spatial distribution of average LAND parcel prices in Decentraland. These are the average prices over the entire time horizon of the data set to detect possible price-influencing factors. Darker areas indicate a higher average price. Heatmap (a) uses the bird's eye view of Decentraland and (b) the points of interest (POI) map of the key locations we have considered in this paper.


prices in times of high activity (see Figure (7) & (8)). At times of low activity, such as 05/2019-01/2020, the average trading price of a LAND parcel is approx. \$1022.82, the trading volume is approx. \$0.12 million (see Figure (9)), and the successful transaction count is 159. Compared to the time frame (03/2021-08/2022), where activity peaks, several differences exist in the magnitude of several indicators. At the peak, the average trading price of a property is \$17928.78 (+1652.88%), the trading volume is \$32.92m (+27333.33%), and the number of transactions is 1910 (+1101.26%). The number of parcel owners during this period increased by 70.23% (from 3641 to 6198). This was followed by a decline in market activity to the previous level. Meanwhile, the decentralization indicators continued to show a steady increase, albeit very slow, such as in the LAND ownership of the top 5% (1%) largest LAND owners. In conclusion, the strong fluctuations in trading prices and volumes during peak periods underscore the speculative nature of this market, while the overall trend towards greater decentralization suggests a gradual stabilization as the user base continues to grow and diversify.

Figure 6: Monthly Circulating LAND




Notes: This figure shows the monthly circulating tradable quantity of LAND over the entire time horizon.

Figure 7: Average LAND Price




Notes: This figure shows the monthly average trading price of LAND over the entire time horizon.

Figure 8: Number of Active Wallets



Notes: This figure shows the number of monthly active wallets over the entire time horizon.

Figure 9: LAND Market Trading Volume and Number of Sales



Notes: This figure shows the comparison of trading volume (in mln. USD; blue) with the number of sales (green) over the entire time horizon.

3 The Decentraland Price Index

Given our database of successful sales transactions, we aim to construct a Decentraland price index. Simply averaging prices on a daily basis would not be an informative approach because parcels are not only sold infrequently, but also vary in their characteristics and consequently in their values. For example, on days with few transactions, the sale of a high-value parcel could skew the daily average and not accurately reflect the true price level. Therefore, average prices must be quality-adjusted. Following the method used to construct the German Real Estate Index (Amaral et al., 2023), we use a time-dummy hedonic regression.⁷

We construct the Decentraland price index (DPI) as follows:

$$\ln p_{i,n} = c + \sum_{t=1}^T \gamma^t D_i^t + \sum_{k=1}^K \delta^k z_n^k + \varepsilon_{i,n} \quad (1)$$

For each transaction i , where a parcel n was sold, we regress the log of the price p on a constant c , a set of time dummies D_i^t , and parcel-specific characteristics z_n^k , denoting the k^{th} characteristic of parcel n . The time dummy D_i^t is set to one if the transaction occurred at time t and zero otherwise. In a daily price index, $t = T$ represents the last day, and $t = 0$ the first day in the sample, with the first dummy

⁷For a more detailed description of the method, see European Commission. Eurostat. (2013).

serving as the omitted category. This allows the coefficients γ^t to be interpreted relative to the first day. $\varepsilon_{i,n}$ is the error term.

Fortunately, it is unnecessary to include as many characteristics (z_n) as would be required for real-world land or real estate data. As previously discussed, each parcel is uniformly sized. However, we know from previous insights (Chapter 2.3) that the location of a parcel is an important determinant of its value. We include three different location-related characteristics (z_n) for each parcel: the distance to the nine plazas and adjacency to a district as well as to a street. Additionally, we employ a dummy variable to indicate whether the parcel's coordinates are numerically memorable, defined as having identical absolute x and y coordinates.

In order to obtain as much information as possible about the effects of location, while accounting for the fact that many of the variables are correlated with each other, we estimate Equation (1) using four different specifications.

For all specifications, we use dummy variables that are equal to one if a parcel is adjacent to a given district and zero otherwise. We also always use a dummy that tells us whether the parcel is adjacent to a street and the aforementioned dummy for easy-to-remember coordinates. The specifications differ in their consideration of distances to plazas. The distance to a plaza is calculated as the Euclidean distance from the center of the parcel to the center of the plaza.

For specification (I), the log distance to the nearest plaza is used, regardless of the specific plaza. Table (1) indicates that an increase in distance by 1% is associated with a reduction in price by 0.24%: The further away the parcel is from a plaza, the lower the price. However, in this specification, we do not differentiate between the plazas. Plaza 1, in particular, has a distinctive position: it is not only the center of the map, but also serves as the point of origin for all avatars upon logging into Decentraland. Consequently, we utilize the log distance to Plaza 1 and the smallest log distance to the other plazas for specification (II). The coefficient associated with Plaza 1 is more negative than those associated with the other plazas.

Table 1: The effect of LAND characteristics

	(I)	(II)	(III)	(IV)
Distance nearest plaza	-0.240 [0.007***]			-0.992 [0.039***]
Distance plaza 1		-0.184 [0.009***]		
Distance nearest plaza (2-9)		-0.113 [0.010***]		
Nearest plaza (Plaza 1 omitted)				
Plaza 2			-0.179 [0.017***]	-2.510 [0.166***]
Plaza 3			-0.124 [0.020***]	-2.971 [0.207***]
Plaza 4			-0.270 [0.015***]	-2.737 [0.164***]
Plaza 5			-0.326 [0.014***]	-2.942 [0.156***]
Plaza 6			-0.336 [0.014***]	-2.942 [0.156***]
Plaza 7			-0.321 [0.014***]	-3.229 [0.149***]
Plaza 8			-0.187 [0.015***]	-3.111 [0.157***]
Plaza 9			-0.168 [0.017***]	-3.136 [0.170***]
Distance nearest plaza $\times$ Nearest plaza ($\times$ Plaza 1 omitted)				
Plaza 2				0.705 [0.047***]
Plaza 3				0.827 [0.059***]
Plaza 4				0.738 [0.046***]
Plaza 5				0.840 [0.043***]
Plaza 6				0.772 [0.044***]
Plaza 7				0.859 [0.042***]
Plaza 8				0.854 [0.045***]
Plaza 9				0.868 [0.048***]
AETHERIAN Project	0.490 [0.066***]	0.527 [0.065***]	0.316 [0.068***]	0.396 [0.070***]
Vegas City	0.386 [0.044***]	0.404 [0.043***]	0.168 [0.045***]	0.253 [0.048***]
Dragon City	0.056 [0.019***]	0.076 [0.019***]	0.028 [0.021]	0.078 [0.021***]
Dragon Kingdom	0.158 [0.051***]	0.195 [0.051***]	0.079 [0.124***]	0.183 [0.051***]
Red Light District	0.131 [0.050***]	0.158 [0.049***]	0.133 [0.051***]	0.107 [0.050***]
University and Conference Center	0.226 [0.032***]	0.250 [0.032***]	0.095 [0.034***]	0.143 [0.034***]
Fashion Street	0.300 [0.031***]	0.344 [0.031***]	0.237 [0.033***]	0.259 [0.032***]
Road	0.408 [0.008***]	0.400 [0.008***]	0.435 [0.008***]	0.428 [0.008***]
Easy to remember	0.134 [0.078*]	0.065 [0.077]	0.105 [0.078]	0.149 [0.076**]
# obs.	24,232	24,232	24,232	24,232

Notes: The table shows the coefficients of the parcel characteristics (z_n) resulting from the estimation of Equation (1). The estimates also include a constant and a time dummy for every day in the sample. ***, ** and * denote significance at the 1%, 5% and 10% level.

In specification (III), dummy variables are used to indicate which of the plazas are the closest to each parcel. The coefficients are estimated relative to Plaza 1 and are all negative. This implies that if a parcel has Plaza 1 as its closest plaza, the price there is significantly higher than if the parcel has one of the other plazas as its closest plaza. If plaza 6 is the nearest plaza, the effect is most pronounced in comparison to plaza 1.

The final specification (IV) incorporates interactions between the plaza dummies

and the distance to the nearest plaza.

Across all specifications, it can be seen that direct adjacency to a district has a significantly positive effect on the price, as does adjacency to a road. Furthermore, coordinates that are easy to remember only have a significantly positive effect on the price for specifications (I) and (IV).

To sum up, our regression results identify several key location factors that significantly influence parcel values in Decentraland. These findings support the observations discussed in Chapter 2.3: central parcels near Genesis Plaza, particularly around coordinates (0, 0), command higher prices, while peripheral areas tend to have lower values. Similarly, proximity to prominent landmarks such as Fashion Street, District X, and Decentraland University, as well as community-driven projects like the AETHERIAN Project, enhances parcel value. Finally, parcels adjacent to major roads and plazas show higher values.

Controlling for the mentioned characteristics leads to the quality adjusted price index which can be calculated by:

$$P^{0t} = \exp(\hat{\gamma}^t), \quad (2)$$

where P^{0t} is the price on day t relative to the base period ($t = 0$). $\hat{\gamma}^t$ is the estimated γ coefficient for day t from Equation (1). Given that the resulting index exhibits relatively minor differences between the four specifications, we use the estimated beta coefficients from specification (I) for our baseline index.⁸

⁸Figure A.1 in the appendix shows the indices resulting from all four specifications.

Figure 10: The Decentraland Price Index



Notes: This figure presents the Decentraland Price Index, derived from successful sales on the Decentraland platform, using the methods discussed in Section 3.

Figure (10) shows the resulting Decentraland Price Index, which begins on March 20, 2018. All values are interpreted relative to March 19, 2018, the official launch date of the Decentraland marketplace. By this point, the initial coin offering on August 17, 2017 and the first land auction on December 15, 2017 had already occurred.⁹ During the initial phase following the marketplace’s opening, when only off-chain transactions were possible, the price index ranged between 0.23 and 4.8. Due to the low number of successful daily transactions, the index fluctuated significantly, with some days recording no transactions at all. For our subsequent empirical analysis, we use the index starting from February 20, 2020, the public launch date, when transactions became on-chain and accessible to all users. A notable spike in the index occurs on this day and the following day, reaching a value of 4.3, indicating that prices were 4.3 times higher than on the first day of our sample. Prices continued to rise steadily, reaching a new high of just under 23 on May 3, 2021 as Decentraland gained popularity, with companies and institutions establishing their first virtual branches. For example, Sotheby’s opened a virtual gallery on June 4, 2021. The hype period began in October 2021, with prices rising sharply thereafter. The rebranding of Facebook to Meta, announced on October 28, 2021, likely contributed to this surge, sparking increased interest in metaverse platforms such as Decentraland. Our index peaked on November 1, 2022, when prices were almost 57 times higher than at the start of the sample period. However, within six months, prices fell drastically

⁹For further details on land auctions and pricing, see Goldberg et al. (2024).

to around 5, and in the final third of our sample, values fluctuated significantly around this level.

In the next chapter, we will examine how the Decentraland Price Index (DPI) responds to real-world monetary policy.

4 The effect of monetary policy on Metaverse prices

Monetary policy impacts real-world housing or real estate prices through several direct and indirect channels. For instance, an increase in interest rates could raise the cost of credit or mortgages, which, *ceteris paribus*, would lower the demand for housing and subsequently lead to a decrease in prices. Moreover, if households anticipate this response and observe the rise in interest rates, their expectations may directly impact house prices via the expectations channel.

But why should we expect virtual real estate prices to react to real-world monetary policy? As mentioned above, virtual land prices exhibit characteristics that make them comparable to real land prices: parcels can be bought, sold, and rented to other users. If we assume that the price of a parcel can be approximated by the present value of future rental payments, monetary policy can effect prices via the discount rate.

The key question is determining the relevant discount rate for users of the metaverse. If users invest and borrow exclusively within Decentraland and assume that the metaverse economy is entirely closed, then real-world monetary policy should have no effect on the discount rate and, consequently, on prices in Decentraland. However, we believe that this relationship is also influenced by the level of activity and popularity of the metaverse: the more users participate in the marketplace, the higher the likelihood that they are not exclusively active within the metaverse. Consequently, the transmission of real-world monetary policy becomes more relevant.

In such a market state, a positive monetary policy shock would likely increase the discount rate and lead to a decrease in prices - mirroring the expected reaction in real-world markets. We will test these conjectures later in the paper.

Before proceeding, however, it is essential to define what constitutes real-world monetary policy in our context. Decentraland operates globally, and we lack information about the nationalities of individual users or which specific monetary policies are most relevant to them. Nonetheless, the Fed's monetary policy is a main driver of global financial conditions (Rey, 2015; Miranda-Agrippino and Rey, 2020) and often considered a key factor behind the Global Financial Cycle of risky assets (Karau, 2023).

Therefore, in the next section, we discuss the calculation of Fed monetary policy shocks and subsequently estimate their effects on metaverse prices.

4.1 Monetary policy shocks

To make causal claims about the effects of monetary policy, it is essential to identify exogenous or unexpected changes in the policy rate, often referred to as shocks or surprises. A common approach to constructing these shocks involves examining changes in relevant indicators on Federal Reserve meeting days. Some of these high-frequency shocks focus on changes in asset prices within narrow time windows, typically around 30 minutes, surrounding the meetings (as in Gertler and Karadi (2015), Nakamura and Steinsson (2018) or Bauer and Swanson (2023)). This small time interval is assumed to capture only the information released during the meeting, thereby isolating its effect on asset prices. Moreover, asset prices should react only if something unexpected occurs during the meeting.

We update the shock series by Bu et al. (2021), which is originally available until 2020. This shock series offers several advantages over the high-frequency shocks mentioned above. Specifically, it examines changes in zero-coupon rates of various maturities on meeting days of the FOMC, eliminating the need for sometimes noisy and difficult-to-access tick-data. The estimation process is straightforward, and the required data are readily accessible.¹⁰ Further, according to Brennan et al. (2025), unlike some high-frequency shocks, these shocks often produce results that are consistent with New Keynesian models.

In the case of a pure monetary policy shock, interest rates and stock prices are expected to move in opposite directions. However, it has been observed that stock prices sometimes move with interest rates. One explanation for this could be the presence of a central bank information shock, which occurs when the central bank releases previously unknown information about the economic outlook during its meetings. According to Bu et al. (2021), their shock series captures a pure monetary policy shock and explicitly excludes central bank information shocks.

Building on this, Ciminelli et al. (2022) extend the procedure of Bu et al. (2021) by estimating information shocks through a simple regression of the change in the policy indicator on the monetary policy shock:


$$\Delta i_t = \theta_t + \gamma \hat{e}_t + \xi_t, \quad (3)$$

where the residuals ξ_t are the information shock.

¹⁰For the calculation of the shocks see B.

Figure (11) shows the estimated shock series alongside the original shocks. As observed, our updated shocks are highly correlated with the originals, although they should not be exactly identical due to a larger sample size and slightly different meeting days. In contrast to Bu et al. (2021), who only include regular meeting days, we also consider unscheduled meeting days.

Figure 11: Pure monetary policy (mp) and information (info) shocks



Notes: The figure shows the updated monetary policy and information shocks, derived using the methods of Bu et al. (2021) and Ciminelli et al. (2022), alongside the original shocks. For better comparability, unscheduled meetings are excluded from the figure.

4.2 Linear model

After calculating the shock series, we use local projections (Jordà, 2005) to estimate the effect of the Fed’s monetary policy on Decentraland’s prices:

$$DPI_{t+h} = \alpha_h + \beta_h s_t + \gamma_h X_t' + \eta_{t+h} \quad (4)$$

DPI is the log of the 3-day moving average of our Decentraland price index multiplied by 100. s_t is either our pure monetary policy shock or the information shock. The daily series takes the value of the corresponding shock on meeting days and is zero otherwise. We further include control variables (X_t). We use the log of the VIX index of implied stock market volatility, the Daily News Sentiment Index, as described in Shapiro et al. (2022) and interest rates measured by the yield on 10-year U.S. government bonds as control variables for real-world economic conditions. To control for developments in the crypto market, we use the first principal component of selected cryptocurrencies.¹¹ We further estimate the model with seven lags of the dependent variable *DPI*, the exogenous control variables and the shock, respectively. Our sample spans from February 20, 2020, the day of the public release, until September 20, 2023.


The variables of interest are the β_h -coefficients measuring the response of the Decentraland prices h days after the shock occurs. If monetary policy has an effect on prices in the metaverse, we would expect prices to fall after a positive pure monetary policy shock and to rise after a positive information shock.

Figure (12) plots the estimated β_h -coefficients. In case of a positive monetary policy shock of 25 basis points prices fall by about 20%. This pronounced effect, however, is only significant at the 90% significance level on the fourth day after the shock.

If we look at the effects of an information shock, we can also observe the expected sign here: prices rise by up to 38%. Again, this effect is only significant on one day, six days after the shock occurs.

¹¹For more information, please refer to Appendix C.

Figure 12: The response to a monetary policy and central bank information shock



Notes: This figure shows the daily responses of Decentraland prices to a 25-basis-point monetary policy (a) and information (b) shock. The shaded areas represent the 68% and 90% probability bands.

4.3 Different states of user activity

The number of users in Decentraland fluctuates considerably over our sample period. Following the official launch of Decentraland in February 2020, user numbers remained relatively low before experiencing a significant increase during the hype period, only to decline again thereafter.

The effect of monetary policy varies depending on the state of user activity. When activity is high, and a large number of users are actively buying and selling land parcels, the effect of monetary policy should closely mirror that observed in real-world markets. Specifically, we expect prices to fall following a monetary policy shock and to rise in response to an information shock when user activity is high. However, when user activity is low, the effects are less predictable. To test this, we estimate state-dependent local projections:

$$\text{DPI}_{t+h} = F(a_{t-1})[\alpha_h^{\text{high}} + \beta_h^{\text{high}} s_t + \gamma_h^{\text{high}} \mathbf{X}_t'] + (1 - F(a_{t-1}))[\alpha_h^{\text{low}} + \beta_h^{\text{low}} s_t + \gamma_h^{\text{low}} \mathbf{X}_t'] + \eta_{t+h} \quad (5)$$

In this smooth transition model, the effect of the shock on the dependent variable is a weighted average of the effects in high and low activity states. The weight of each regime for each observation is determined by $0 \leq F(a_{t-1}) \leq 1$. We follow Born et al. (2020) and use the empirical cumulative distribution function of the state variable. We measure activity in Decentraland by the number of successful sales recorded in our transaction database and use the lagged series of this measure as our state variable a_{t-1} .


The transition function evolves as:

$$F(a_{t-1}) = \frac{1}{N} \sum_{j=1}^N \mathbb{1}_{a_j \leq a_{t-1}}, \quad (6)$$

where N is the number of observations in our sample, and $\mathbb{1}$ is an indicator function that equals one if $a_j \leq a_{t-1}$ and zero otherwise. $F(a_{t-1}) = 0$ if the state variable is at its minimum and $F(a_{t-1}) = 1$ when it is at its maximum. The advantage of this method is that it does not require further assumptions about the parameters of the transition function, unlike the approach used by Auerbach and Gorodnichenko (2012).

Figure (13) plots the values of the transition function which can be interpreted as the probability of being in the high activity state. Immediately after the official launch of Decentraland, we observe high probabilities of being in the high activity state. Following this, the probabilities decline but exhibit substantial fluctuations. Notably, the probabilities increase again after Meta's rebranding in February 2021. Thereafter, the state variable reaches its peak during the hype period between November 2021 and January 2022, with the probabilities of being in a high state approaching one. Subsequently, the weights decrease again, showing considerable variation between very high and very low probabilities.

Figure 13: Transition function $F(a_{t-1})$




Notes: This figure shows the values of the transition function $F(a_{t-1})$ based on the method of Born et al. (2020). State variable is the number of successful sales in Decentraland.

We estimate Equation (5) using the described transition function and the same

variables as in the linear model. Figure (14) presents the corresponding impulse responses. We see much clearer effects of the shocks compared to the linear model. In the case of a monetary policy shock and high activity, prices react even on impact and fall by up to 40% on the third day after the shock, while they rise in the case of low activity. Overall, the effects are not very persistent. In the high state, the effect disappears on the sixth day, in the low state on the fourth day.

Figure 14: The response to a mp and in shock under different states of user activity



Notes: This figure shows the daily responses of Decentraland prices to a 25-basis-point monetary policy (a) and information (b) shock under two different states of user activity. The state variable is defined as the daily number of successful sales. The shaded areas represent the 68% and 90% probability bands.

Even with an information shock, prices react more strongly than with the linear model when activity in decentraland is high. They rise to a maximum of 70% on the sixth day, while no significant effects can be detected with low activity.

In summary, Decentraland prices respond to pure monetary and information shocks in high user activity states with the expected signs, similar to their real-world counterparts. However, there are notable differences: First, while real-world prices tend to react with a delay, the DPI in Decentraland responds immediately or within just a few days. Second, these effects are very short-lived; no significant impact can be detected after a week at the latest. Third, the magnitude of the effects is substantial.

We use the results from our state-dependent model as baseline for the following robustness analysis.

4.4 Robustness

We use the number of successful sales as the state variable for our previously discussed baseline results. This variable derives from our self-collected transaction dataset.

We now examine whether our results remain robust if we change the state variable to one not self-collected or constructed by us. We therefore use the daily number of active wallet addresses in Decentraland, as reported by Messari.¹² Figure 15 illustrates the corresponding impulse responses. As is apparent, a positive monetary policy shock leads to an immediate and significant drop in prices when user activity is high. The effect persists for approximately 4 to 5 days and yields values very similar to our baseline results. If user activity is low, we again observe a positive effect on prices. Also the effects of the information shock mirrors our baseline results closely: Prices experience a rise of about 50% which is only significant to the 90% level at the sixth day after the shock when user activity is high. We don't see an effect in the low user activity state.

As Brennan et al. (2025) show, monetary policy shocks can sometimes differ greatly depending on the methods and data chosen. Since there are only a comparatively small number of meeting days in our sample, the choice of monetary policy shock could have a significant influence on our results. We therefore estimate Equation (5) using the updated monetary policy and information shocks of Jarociński and Karadi (2020). They use high-frequency data around FOMC meetings and estimate the shocks using a bayesian VAR and sign restrictions. Figure 16 shows that our main results hold even with these shocks. With high user activity, we see a negative effect of a monetary policy shock with very similar values to our baseline estimate. The effect of information shocks is positive in the high state, but the estimated coefficients are not significant. We see no effects of monetary and information shocks on prices when activity is low.

Finally, we test whether our results depend on the chosen method of state-dependent local projections. We now use the approach of Auerbach and Gorodnichenko (2012) and Tenreyro and Thwaites (2016). This differs from the method by Born et al. (2020) in the calculation of the transition function $F(a_{t-1})$, which is now expressed by a logistic function:

$$F(a_{t-1}) = \frac{\exp\left(\theta \frac{a_{t-1}-c}{\sigma_a}\right)}{1 + \exp\left(\theta \frac{a_{t-1}-c}{\sigma_a}\right)} \quad (7)$$


θ measures the smoothness of the transition between the two states, with higher values indicating a smoother transition. As common in the literature, we set $\theta=3$. We set c , such that for 30% of the time Decentraland is in the high activity state.

Figure 17 shows the resulting impulse responses, which closely mirror our baseline results. Even the values for both states and shocks are nearly identical to those observed previously.

¹²www.messari.io


In summary, our findings remain robust with respect to alternative state variables, shocks, and transition functions.

Figure 15: The response to a monetary policy and central bank information shock: Alternative state variable




Notes: This figure shows the daily responses of Decentraland prices to a 25-basis-point information shock under two different states of user activity. The state variable is defined the daily number of active addresses. The shaded areas represent the 68% and 90% probability bands.

Figure 16: The response to a monetary policy and central bank information shock: Alternative shock series (Jarociński and Karadi, 2020)



Notes: This figure shows the daily responses of Decentraland prices to a 25-basis-point monetary policy (a) and information (b) shock under two different states of user activity. The state variable is defined as the daily number of successful sales. The shaded areas represent the 68% and 90% probability bands.

Figure 17: The response to a monetary policy and central bank information shock: Alternative transition function (Auerbach and Gorodnichenko, 2012)



Notes: This figure shows the daily responses of Decentraland prices to a 25-basis-point monetary policy (a) and information (b) shock under two different states of user activity. The state variable is defined as the daily number of successful sales. The shaded areas represent the 68% and 90% probability bands.

5 Conclusions

This paper examines the effects of real-world monetary policy on virtual land prices. We collect relevant on-chain transaction data from various marketplaces and construct a price index to capture the price dynamics of virtual real estate on the platform. Our findings demonstrate that virtual land prices do respond to Federal Reserve monetary policy. Using state-dependent local projections, we observe that when activity and popularity in the virtual world are high, prices respond to monetary policy and information shocks in the expected direction, mirroring real-world reactions.

However, important differences are evident: Prices in the metaverse react much faster, either immediately in the case of a pure monetary policy shock or within a few days following an information shock. Additionally, the effects are short-lived, disappearing within a week. Finally, the magnitude of the effects is substantial, with prices responding even more strongly to information shocks.

Understanding the economic dynamics of virtual real estate offers valuable insights for investors and developers aiming to capitalize on metaverse opportunities. Our findings equip stakeholders with the ability to make informed decisions in anticipation of monetary policy trends, while also highlighting the importance of managing risks associated with policy-induced price volatility to ensure the resilience of virtual asset investments against macroeconomic fluctuations. By providing empirical evidence on how monetary policy impacts virtual real estate prices, we contribute to the broader

discourse on the increasing interconnection between digital and physical economic realms.

Furthermore, by demonstrating that economic mechanisms and reactions also apply in the metaverse, our findings suggest that the metaverse can serve as a valuable setting for further economic research. For instance, field experiments that would be challenging to conduct in the real world become feasible in the metaverse due to a crucial property: virtually all behavior is observable.

References

- Amaral, F., Dohmen, M., Schularick, M. and Zdrzalek, J. (2023). German real estate index (greix), *ECONtribute Discussion Papers Series* **231**.
- Auerbach, A. J. and Gorodnichenko, Y. (2012). Measuring the output responses to fiscal policy, *American Economic Journal: Economic Policy* **4**(2): 1–27.
- Ball, M. (2022). The metaverse: And how it will revolutionize everything, *Liveright Publishing*.
- Bauer, M. D. and Swanson, E. T. (2023). A reassessment of monetary policy surprises and high-frequency identification., *NBER Macroeconomics Annual* **37**: 87–155.
- Benati, L. (2021). Leaning against house prices: A structural VAR investigation, *Journal of Monetary Economics* **118**: 399–412.
- Blockchain Research Lab (2022). An analysis of LAND in The Sandbox: Q4 2022 | Research Note.
URL: <https://www.blockchainresearchlab.org/2022/11/22/an-analysis-of-land-in-the-sandbox-q4-2022-brl-research-note/>
- Born, B., Müller, G. J. and Pfeifer, J. (2020). Does austerity pay off?, *The Review of Economics and Statistics* **102**(2): 323–338.
- Bowes, D. R. and Ihlanfeldt, K. R. (2001). Identifying the impacts of rail transit stations on residential property values, *Journal of Urban Economics* **50**(1): 1–25.
- Brennan, C. M., Jacobson, M. M., Matthes, C. and Walker, T. B. (2025). Monetary policy shocks: Data or methods?, *The BE Journal of Macroeconomics* **25**(2): 595–659.
- Bu, C., Rogers, J. and Wu, W. (2021). A unified measure of Fed monetary policy shocks, *Journal of Monetary Economics* **118**: 331–349.
- Calza, A., Monacelli, T. and Stracca, L. (2013). Housing finance and monetary policy, *Journal of the European Economic Association* **11**: 101–122.
- Cantú, C., Franco, C. and Frost, J. (2024). The economic implications of services in the metaverse, *Global perspectives in the metaverse: law, economics, and finance*, Springer, pp. 83–118.
- Chanda, P. (2023). Decentraland vs Sandbox: A Detailed Comparison (2024).
URL: <https://ambcrypto.com/blog/decentraland-vs-sandbox-a-detailed-comparison-2024/>
- Ciminelli, G., Rogers, J. and Wu, W. (2022). The effects of U.S. monetary policy on international mutual fund investment, *Journal of International Money and Finance* **127**: 102676.
- Decentraland (2020). The gates to Decentraland have opened!
URL: <https://decentraland.org/blog/announcements/decentraland-launch>

- Decentraland (2023). Brands and the Metaverse: Embracing the future of customer engagement.
URL: <https://decentraland.org/blog/announcements/brands-and-the-metaverse-embracing-the-future-of-customer-engagement>
- Decentraland Fandom (2024). Brands | Decentraland Wiki | Fandom.
URL: <https://decentraland.fandom.com/wiki/Brands>
- Deininger, K. and Olinto, P. (2000). Asset distribution, inequality, and growth, *Policy Research Working Paper Series 2375*, The World Bank.
- Del Negro, M. and Otrok, C. (2007). 99 Luftballons: Monetary policy and the house price boom across U.S. states, *Journal of Monetary Economics* **54**(7): 1962–1985.
- Despotović, V., Bjelica, D., Bjelica, A. and Barać, D. (2023). Decentralized autonomous organizations: State and perspectives, *E-business technologies conference proceedings* **3**(1): 238–243.
- Dowling, M. (2022a). Fertile LAND: Pricing non-fungible tokens, *Finance Research Letters* **44**: 102096.
- Dowling, M. (2022b). Is non-fungible token pricing driven by cryptocurrencies?, *Finance Research Letters* **44**: 102097.
- Dwivedi, Y. K., Hughes, L., Baabdullah, A. M., Ribeiro-Navarrete, S., Giannakis, M., Al-Debei, M. M., Dennehy, D., Metri, B., Buhalis, D., Cheung, C. M., Conboy, K., Doyle, R., Dubey, R., Dutot, V., Felix, R., Goyal, D., Gustafsson, A., Hinsch, C., Jebabli, I., Janssen, M., Kim, Y.-G., Kim, J., Koos, S., Kreps, D., Kshetri, N., Kumar, V., Ooi, K.-B., Papagiannidis, S., Pappas, I. O., Polyviou, A., Park, S.-M., Pandey, N., Queiroz, M. M., Raman, R., Rauschnabel, P. A., Shirish, A., Sigala, M., Spanaki, K., Wei-Han Tan, G., Tiwari, M. K., Viglia, G. and Wamba, S. F. (2022). Metaverse beyond the hype: Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy, *International Journal of Information Management* **66**: 102542.
- European Commission. Eurostat. (2013). *Handbook on residential property prices indices (RPPIs)*., Publications Office, LU.
- Fleming, J. M. (1962). Domestic Financial Policies under Fixed and under Floating Exchange Rates (Politiques financières intérieures avec un système de taux de change fixe et avec un système de taux de change fluctuant) (Política financiera interna bajo sistemas de tipos de cambio fijos o de tipos de cambio fluctuantes), *Staff Papers - International Monetary Fund* **9**(3): 369.
- Gertler, M. and Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity, *American Economic Journal: Macroeconomics* **7**(1): 44–76.
- Goldberg, M., Kugler, P. and Schär, F. (2024). Land valuation in the metaverse: Location matters, *Journal of Economic Geography* **24**(5): 729–758.

- Goldberg, M. and Schär, F. (2023). Metaverse governance: An empirical analysis of voting within Decentralized Autonomous Organizations, *Journal of Business Research* **160**: 113764.
- Goodhart, C. and Hofmann, B. (2008). House prices, money, credit, and the macroeconomy, *Oxford Review of Economic Policy* **24**(1): 180–205.
- Güven, İ. and Ercan, T. (2022). Determining Factors of Virtual Land Value: The Case of Decentraland, *ITESDES: International Technology Sciences and Design Symposium*, pp. 518–537.
- Hazan, E., Kelly, G., Khan, H., Spillecke, D. and Yee, L. (2022). Marketing in the metaverse: An opportunity for innovation and experimentation, *The McKinsey Quarterly*.
- Heikkilä, E., Gordon, P., Kim, J. I., Peiser, R. B., Richardson, H. W. and Dale-Johnson, D. (1989). What Happened to the CBD-distance gradient?: Land values in a polycentric city, *Environment and Planning A: Economy and Space* **21**(2): 221–232.
- Huynh-The, T., Gadekallu, T. R., Wang, W., Yenduri, G., Ranaweera, P., Pham, Q.-V., da Costa, D. B. and Liyanage, M. (2023). Blockchain for the metaverse: A review, *Future Generation Computer Systems* **143**: 401–419.
- Hülsewig, O. and Rottmann, H. (2021). Euro area house prices and unconventional monetary policy surprises, *Economics Letters* **205**: 109962.
- Iacoviello, M. (2005). House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle, *American Economic Review* **95**(3): 739–764.
- ICO Drops (2017). Decentraland (MANA) - All information about Decentraland ICO (Token Sale). Section: Ended ICO.
- 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.
- Jordà, Ò., Schularick, M. and Taylor, A. M. (2015). Betting the house, *Journal of International Economics* **96**: S2–S18.
- Karau, S. (2023). Monetary policy and Bitcoin, *Journal of International Money and Finance* **137**: 102880.
- Kiley, M. T. (2021). What macroeconomic conditions lead financial crises?, *Journal of International Money and Finance* **111**: 102316.
- Miranda-Agrippino, S. and Rey, H. (2020). U.S. monetary policy and the global financial cycle, *The Review of Economic Studies* **87**(6): 2754–2776.


- Momtaz, P. P. (2022). Some very simple economics of Web3 and the Metaverse, *FinTech* **1**(3): 225–234.
- Mundell, R. A. (1963). Capital mobility and stabilization policy under fixed and flexible exchange rates, *Canadian Journal of Economics and Political Science* **29**(4): 475–485.
- Nakamura, E. and Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: The information effect, *The Quarterly Journal of Economics* **133**(3): 1283–1330.
- Nakavachara, V. and Saengchote, K. (2022). Does unit of account affect willingness to pay? Evidence from metaverse LAND transactions, *Finance Research Letters* **49**: 103089.
- Ordano, E., Jardi, Y., Meilich, A. and Araoz, M. (2024). Decentraland white paper, *Decentraland* p. 14. Section: player.
- Park, S.-M. and Kim, Y.-G. (2022). A metaverse: taxonomy, components, applications, and open challenges, *IEEE Access* **10**: 4209–4251.
- Paul, P. (2020). The time-varying effect of monetary policy on asset prices, *The Review of Economics and Statistics* **102**(4): 690–704.
- Rey, H. (2015). Dilemma not trilemma: The global financial cycle and monetary policy independence, *NBER Working Paper Series* (21162).
- Rigobon, R. and Sack, B. (2004). The impact of monetary policy on asset prices, *Journal of Monetary Economics* **51**(8): 1553–1575.
- Shapiro, A. H., Sudhof, M. and Wilson, D. J. (2022). Measuring news sentiment, *Journal of econometrics* **228**(2): 221–243.
- Tenreyro, S. and Thwaites, G. (2016). Pushing on a string: US monetary policy is less powerful in recessions, *American Economic Journal: Macroeconomics* **8**(4): 43–74.
- The Sandbox (2020). The Sandbox, *The Sandbox* p. 44.
- The Sandbox (2024). Become a Partner with The Sandbox — Join Over 400 Brands.
URL: <https://www.sandbox.game/en/partnerships/>
- Tyrväinen, L. and Miettinen, A. (2000). Property prices and urban forest amenities, *Journal of Environmental Economics and Management* **39**(2): 205–223.
- UBS (2024). Global Wealth Report 2024.
URL: <https://www.ubs.com/content/dam/assets/wm/static/noindex/wm-germany/2024/doodownload/Global-Wealth-Report-2024.pdf>
- Williams, J. C. (2015). Measuring monetary policy’s effect on house prices, *FRBSF Economic Letter* **28**: 1–6.

- Xie, S. (2024). Monetary policy shocks and local housing prices, *International Journal of Central Banking* **21**(3): 327–373.
- Yencha, C. (2023). Spatial heterogeneity and non-fungible token sales: Evidence from Decentraland LAND sales, *Finance Research Letters* **58**: 103628.

Appendices

Appendix A Additional figures

Figure A.1: Different specifications of the Decentraland Price Index



Notes: The figure shows the different specifications of the Decentraland Price Index derived from Equation (1), varying by the choice of parcel-specific control variables (specifications I to IV, Table (1)).

Appendix B Constructing monetary policy shocks

The method of Bu et al. (2021) involves a two step Fama-MacBeth procedure. First, they assess the response of zero-coupon yields with maturities ranging from 1 to 30 years to a policy indicator. They estimate:

$$\Delta R_{m,t} = a_m + \kappa_m \Delta i_t + \epsilon_{m,t}, \quad \text{for } m = 1 \dots 30, \quad (\text{B.1})$$

where $\Delta R_{m,t}$ is the change in the zero-coupon yield at maturity m on day t , and Δi_t is the change in the policy indicator. Following Bu et al. (2021), we use the two-year Treasury yield. The day t can be either a meeting day or a day one week before a meeting. This follows from the heteroskedasticity-based identification of Rigobon and Sack (2004) used to estimate the equation above. The key assumption is that the monetary policy shock is heteroskedastic, i.e. the variance of the shock is higher on meeting days than on non-meeting days (e.g. seven days before the meeting), while other shocks are homoskedastic. Equation (B.1) can then be estimated using

Δi_t and $-\Delta i_t$ as instruments for meeting and non-meeting days, respectively. To identify meeting days, we use the same dates as Bauer and Swanson (2023) and update them to the end of 2023. Our dataset spans from 1994 to the end of 2023. The first step yields the estimated $\hat{\kappa}$ coefficients for each maturity. In the second step, these coefficients are used to run a cross-sectional regression for each meeting day t :

$$\Delta R_{m,t} = \delta_m + e_t \hat{\kappa}_m + v_{m,t}, \quad \text{for } t = 1 \dots T, \quad (\text{B.2})$$

where the resulting coefficients e_t are the monetary policy shocks.

Appendix C Principal component analysis

To adequately capture the price movement of the cryptocurrency market, we consider various cryptocurrencies that have existed at least since the launch of the Decentraland MANA Initial Coin Offering and whose characteristics serve to diversify the dataset (see Table (C.1)). In our case, the criteria are market representation, sector representation, and correlation and redundancy.¹³ We then describe the application of PCA.

While all these are cryptocurrencies, their uses vary from payment systems and digital currencies (e.g., BTC) to platforms enabling smart contracts, decentralized applications (dApps), or specific use cases like supply chain management (VET). Each has a unique role in the broader blockchain ecosystem. Here is a brief overview of each:

1. Application-Specific Tokens: These tokens are tailored for specific applications or services within a blockchain.
 - Chainlink (LINK): A Utility token for decentralized oracles with a highly application-specific focus. LINK bridges blockchain ecosystems with real-world data. Moreover, it is the leading token in the blockchain oracle sector. LINK enables smart contracts to interact with external data. Finally, it

¹³In the context of cryptocurrency diversification, redundancy refers to the overlap in functionality, behavior, or market movements among the cryptocurrencies in the dataset. In PCA, redundancy can reduce the effectiveness of dimensional reduction because highly correlated data points contribute less unique information. While some redundancy is inevitable and even desirable for capturing the underlying structure, excessive redundancy might overshadow unique relationships or insights from less-correlated assets. Balancing the dataset with cryptocurrencies that cover diverse functions, sectors, and behaviors ensures the analysis provides a more comprehensive view of the cryptocurrency market.

captures the market dynamics of blockchain ecosystems focusing on data interoperability.

- Decentraland MANA (MANA): Used for transactions in the Decentraland metaverse and enables virtual real estate and asset trading. See the paper content for more information.
- Stellar (XLM): A payment-focused cryptocurrency designed for cross-border transactions on the Stellar network. It represents financial inclusion and targets unbanked populations. While it shares some similarities with XRP, Stellar’s open-source and community-driven approach sets it apart. XLM diversifies the dataset within the payment sector.

Table C.1: Overview of the cryptocurrencies used in the PCA

#	Cryptocurrency	Abbreviation	PoW	PoS	Diverse
1	Binance Coin	BNB			DPoS
2	Bitcoin	BTC	X		
3	Bitcoin Cash	BCH	X		
4	Cardano	ADA		X	
5	Chainlink	LINK		X	
6	DASH	DASH	X		2nd tier masternode network
7	Decentraland Mana	MANA		X	
8	Dogecoin	DOGE	X		
9	EOS	EOS			DPoS
10	Ethereum	ETH		X	
11	Litecoin	LTC	X		
12	NEO	NEO		X	
13	Polygon	MATIC		X	
14	Stellar	XLM			SCP
15	Tezos	XTZ		X	
16	VeChain	VET			PoA
17	XRP	XRP			RPCA

Notes: The table presents an overview of the cryptocurrencies selected for PCA to record cryptocurrency market movements. Proof-of-Work (PoW) and Proof-of-Stake (PoS) represent different consensus mechanisms, with which the different underlying blockchains regulate the creation of new units of the respective currency and state of the blockchain.

2. Cryptocurrencies: These coins primarily act as digital currencies or stores of value.

- Bitcoin (BTC): The first decentralized cryptocurrency, primarily acting as a digital currency and store of value. It is underpinned by blockchain technology. BTC serves as the benchmark for the entire market. Its movements often dictate or correlate with broader trends, making it essential for inclusion in any cryptocurrency analysis.

- Bitcoin Cash (BCH): A fork of Bitcoin designed to enable faster and more cost-effective transactions. It represents an alternative perspective within payment-oriented cryptocurrencies due to its emphasis on peer-to-peer cash use. BCH focuses on scalability for daily payments. Although it exhibits moderate correlation with Bitcoin and Litecoin, Bitcoin Cash adds diversity by exploring blockchain scalability through its larger block size.
 - DASH: Optimized for fast, low-cost payments, with additional privacy features and functions as digital cash. DASH represents the privacy-focused and payment-optimized cryptocurrencies within the dataset.
 - Dogecoin (DOGE): Initially a meme-based cryptocurrency, now widely used for low-cost tipping (e.g., on social media platforms for creating valuable or entertaining content) and transactions. DOGE captures the dynamics of community-driven and speculative market movements.
 - Litecoin (LTC): A peer-to-peer Bitcoin alternative cryptocurrency that offers faster transactions and accessibility in mining. LTC captures the dynamics of Bitcoin alternatives and their relationship to market movements.
3. Hybrid Token: Hybrid tokens are designed to enhance versatility and efficiency within blockchain ecosystems. They enable broad use cases while supporting seamless integration across platforms. Their adaptability makes them essential in ecosystems requiring diverse token functionalities (utility, currency, governance).
- XRP: A cryptocurrency designed for efficient cross-border payments and has been widely adopted in the banking sector. It operates on the Ripple Consensus Ledger (RCL) and uses a unique consensus algorithm with pre-mined tokens validated by select validators. This approach combines elements of decentralization and central control. Unlike traditional cryptocurrencies, XRP does not rely on proof-of-work or proof-of-stake mechanisms. It serves both as a currency and as a bridge asset for international transactions. As a result, this demonstrates its enterprise utility and importance in blockchain-based payments. As a result, XRP captures the dynamics of financial services and their impact on the crypto market.
4. Platform Tokens: These are coins designed to enable or facilitate operations within a blockchain ecosystem, often with specific functionalities beyond being

a currency.

- EOS: A smart contract platform emphasizing scalability and ease of use in blockchain applications. It adds diversity to the smart contract sector with its DPoS consensus model. Moreover, it supports dApps on the EOS blockchain. EOS is included for its relevance in the performance-oriented blockchain sector.
 - Ethereum (ETH): Powers transactions and smart contract execution on the Ethereum blockchain. Therefore, it enables dApps. ETH represents the largest and most influential smart contract platform in the market.
 - Cardano (ADA): A cryptocurrency and leading competitor to ETH in the smart contract space, powered by the Cardano blockchain. Moreover, ADA powers dApps (e.g., Minswap and JPG Store). It emphasizes scalability, sustainability, and academic rigor in blockchain development. ADA represents Layer-1 blockchain platforms with scientific approaches to scalability and sustainability.
 - Tezos (XTZ): A token focusing on governance and self-amending blockchain technology. It represents innovation in blockchain governance and development. Unlike ETH and ADA, XTZ incorporates a built-in mechanism allowing stakeholders (XTZ holders) to directly participate in governance decisions, such as proposing upgrades to the blockchain protocol and voting on these proposals to ensure consensus-driven decision-making (on-chain governance). Moreover, it can upgrade without requiring a hard fork (a split in the blockchain) which is called self-amendment. This reduces the risk of community splits and ensures smoother platform evolution. XTZ reflects self-amending blockchain platforms and their stability-oriented influence on the market.
 - NEO: A digital asset often referred to as the “Chinese Ethereum” with a regional focus on digital identity and regulatory compliance. It represents geographical diversity within smart contract platforms and adds insights into government partnerships. NEO reflects the market activity of blockchain ecosystems focused on digital economies.
5. Utility Tokens: A utility token is a digital asset that gives holders access to specific goods, services, or functionalities within a blockchain-based ecosystem or platform. Unlike cryptocurrencies like Bitcoin, which primarily function as a medium of exchange or a store of value, utility tokens are designed to serve a specific purpose or utility within their ecosystem.


- **Binance Coin (BNB):** An exchange token integral to the BNB ecosystem that reflects trends and movements in exchange-centric crypto assets. The ecosystem plays a central role in trading fee reductions and DeFi applications.
- **Polygon (MATIC):** A utility token based on a leading Layer-2 scaling solution representing efforts to enhance the scalability and efficiency of smart contract platforms. Moreover, the Layer-2 solution facilitates faster and cheaper transactions on Ethereum. MATIC captures the scaling solutions sector and its influence on broader market trends. While it correlates with Ethereum, its focus on reducing transaction costs and congestion differentiates it.
- **VeChain (VET):** A cryptocurrency specialized in supply chain management and enterprise blockchain applications. It uniquely represents integrating blockchain technology into logistics and real-world data tracking. With a highly distinct enterprise-level use case, VeChain exhibits minimal redundancy with other tokens in the dataset.

The dataset was designed to encompass a broad spectrum of blockchain and crypto use cases. By including assets from the above categories, we aimed to capture market trends (Bitcoin and major platform tokens provide insights into overall market movements), represent niche sectors (application-specific tokens and hybrid tokens reflect specialized areas, such as oracles, supply chains, and cross-border payments), highlight innovation (governance-focused tokens like XTZ and scaling solutions like MATIC emphasize technological advancements), and include geographical and functional diversity (assets like NEO and VeChain add regional and industry-specific perspectives). This selection ensures that the dataset reflects the crypto market’s underlying correlation structure and each category’s unique contributions.

Based on the cryptocurrencies presented, we apply principal component analysis to capture the underlying correlative structure and thus the joint market movement. The methodological approach involves transforming the raw time series price data into logarithmic returns to achieve stationarity and z-score standardizing them in the next step to prevent the PCA input correlation matrix domination by large magnitudes of single variables. The next step involves evaluating the appropriateness of the dataset as input for the PCA. Therefore, we apply the Kaiser-Meyer-Olkin (KMO) test (suitability of the dataset for identifying latent structures) and Bartlett’s test of sphericity (significant difference of the correlation matrix from the identity matrix). The correlation matrix input for the tests and the subsequent PCA analysis

can be seen in Figure (C.2). The tests indicate an excellent sampling adequacy ($KMO = 0.977$) and sufficient correlations to justify applying PCA ($p < 0.05$).

Figure C.2: Correlation Heatmap Matrix of Standardized Log Returns




Notes: The correlation heatmap matrix presents the standardised logarithmic returns of the various cryptocurrencies included in the PCA.

Next, we now apply PCA and evaluate the result to determine the optimal number of principal components to at least adequately capture the correlative structure of the cryptocurrency market movement. Afterwards, we evaluate the result to determine the optimal number of principal components to at least adequately capture the correlative structure of the cryptocurrency market movement. For this purpose, we use a scree plot and look at the loadings within the optimal number of principal components with the highest explained variance ratio. In Figure (C.3), the scree plot shows that the first principal component (PC1) explains approximately 66% of the

total variance, a significant proportion. This substantial proportion demonstrates its dominant role in capturing the variability in the data. The steep drop in explained variance from PC1 to subsequent components reflects that PC1 captures the most critical underlying structure, while the remaining components represent more specific or idiosyncratic factors with less explanatory power.

Figure C.3: Scree Plot



Notes: The scree plot visualizes the explained variance ratio for each principal component. It highlights their relative contribution to the total variance in the dataset. In this case, the first principal component (PC1) explains approximately 66% of the total variance and, therefore, significantly dominates the data representation.

The loadings for the first principal component (PC1) are all negative, with values ranging narrowly between -0.902 and -0.550 (see Table (C.2)). These values show that all cryptocurrencies contribute similarly to this component. Moreover, this suggests that PC1 represents a market-wide latent factor that likely captures the collective behavior or systemic trend of the cryptocurrency market. Together, these results suggest that the cryptocurrency market dynamics are largely driven by a single overarching trend, as captured by PC1, with additional components to finer distinctions or sector-specific patterns. As only PC1 is of interest here, we will not go into further detail about the other principal components.

Finally, we generated the time series resulting from PC1 by projecting the standard-

Table C.2: Loadings Matrix from PCA Analysis

Cryptocurrency	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8	PC9	PC10	PC11	PC12	PC13	PC14	PC15	PC16	PC17
Cardano	-0,847	-0,021	0,089	0,148	-0,044	-0,057	-0,249	-0,182	-0,107	0,278	0,085	-0,058	0,049	-0,202	-0,029	0,087	-0,032
Bitcoin Cash	-0,870	0,092	0,020	-0,267	-0,062	0,021	0,080	-0,122	0,058	-0,076	0,017	-0,191	0,006	0,109	-0,204	0,196	0,009
Binance Coin	-0,804	-0,197	-0,143	-0,038	-0,081	0,183	-0,105	0,418	-0,177	0,089	-0,083	-0,137	-0,024	0,022	-0,020	0,015	-0,012
Bitcoin	-0,852	0,039	-0,104	-0,202	-0,077	0,121	-0,284	-0,019	0,100	-0,132	0,075	0,180	0,042	0,069	0,130	0,074	-0,161
DASH	-0,831	0,025	0,007	-0,259	0,019	0,042	0,272	-0,001	-0,031	0,290	-0,073	0,252	-0,073	0,016	0,042	0,062	0,033
Dogecoin	-0,550	0,753	-0,222	0,256	-0,035	0,088	0,042	0,077	0,008	0,013	0,005	0,001	-0,001	-0,013	-0,006	0,002	0,011
EOS	-0,883	0,052	0,117	-0,135	-0,007	0,027	0,166	-0,055	0,015	-0,032	0,023	-0,259	0,039	-0,059	0,283	-0,021	0,019
Ethereum	-0,902	-0,026	-0,063	-0,105	-0,084	-0,013	-0,204	-0,004	0,104	-0,061	0,041	0,062	-0,024	-0,022	0,003	-0,056	0,314
Chainlink	-0,803	-0,077	-0,103	0,111	0,032	-0,403	-0,002	0,139	0,324	0,071	-0,168	-0,030	0,053	0,007	0,007	0,023	-0,049
Litecoin	-0,895	0,056	0,003	-0,210	-0,046	0,045	-0,015	-0,075	0,084	0,023	-0,015	-0,052	-0,156	-0,073	-0,117	-0,276	-0,113
Decentraland MANA	-0,689	-0,196	-0,302	0,148	0,554	0,219	0,024	-0,116	0,071	0,000	-0,002	-0,027	0,006	0,007	-0,010	0,002	0,011
Polygon	-0,692	-0,303	-0,148	0,352	-0,437	0,181	0,188	-0,120	0,085	-0,046	0,003	0,038	0,009	0,004	-0,008	-0,002	-0,011
NEO	-0,874	0,003	-0,026	-0,093	0,030	-0,073	0,095	-0,007	-0,197	-0,138	-0,099	0,093	0,344	-0,079	-0,069	-0,080	-0,005
VeChain	-0,838	-0,015	-0,093	0,113	0,016	-0,237	-0,047	-0,140	-0,293	-0,150	-0,185	0,018	-0,218	0,069	0,050	0,026	-0,006
Stellar	-0,826	0,033	0,375	0,203	0,056	0,033	-0,072	-0,031	-0,023	0,122	0,026	-0,006	0,086	0,302	0,001	-0,098	0,008
XRP	-0,765	0,015	0,495	0,152	0,125	0,139	0,030	0,127	0,085	-0,178	-0,097	0,087	-0,091	-0,154	-0,037	0,076	-0,007
Tezos	-0,833	-0,080	-0,022	0,063	0,076	-0,215	0,141	0,146	-0,070	-0,070	0,431	0,047	-0,059	-0,005	-0,025	-0,003	-0,022


Notes: The loadings matrix in the PCA results represents the contribution of each original variable (cryptocurrency) to each principal component. Each value in the matrix indicates the strength and direction of the relationship between a given cryptocurrency and a principal component. Higher absolute values suggest a stronger influence of the cryptocurrency on that principal component, with positive and negative values indicating the direction of the relationship.

ized log returns onto the principal component space:

$$\text{PC1}_{\text{Scores}_t} = X \cdot W_{\text{PC1}}$$

To calculate the PC1 time series, the standardized log returns matrix X , which represents the transformed and scaled data, is used as the input. The loadings for the first principal component W_{PC1} from the first column of the PCA components matrix are then applied to the standardized data. This involves performing a dot product multiplication between the standardized log returns and the first principal component's loadings to project the data onto the PC1 axis. The result $\text{PC1}_{\text{Scores}_t}$ is a time series that represents a latent factor that captures the overall market-wide trend influencing cryptocurrency returns (see Figure (C.4)). Each PC1 value is a weighted combination of the standardized log returns. They summarize the behavior of the market on a given day. High absolute PC1 values indicate significant deviations from the average trend, with positive values suggesting synchronized movements in one direction (e.g., general increases) and vice versa. Analyzing the temporal patterns of PC1 can reveal periods of high volatility, systematic trends, or responses to market-wide events. The PC1 time series can serve as a proxy for overall market sentiment and provide insights into the collective behavior of cryptocurrencies. By analyzing the PC1 trend, periods of systemic risk or synchronized market movement can be identified, supporting risk management strategies. Additionally, comparing the PC1 time series with external macroeconomic indicators can reveal the impact of broader economic factors on the market, aiding in a comprehensive understanding of how systemic events or external drivers influence aggregated cryptocurrency behavior.

Figure C.4: PC1 Score Development over Time



Notes: The PC1 score time series represents the temporal development of the first principal component, which captures the dominant market-wide trend among the cryptocurrencies. Each value in the series reflects the aggregate behavior of the cryptocurrencies on a given day, weighted by their contributions (loadings) to PC1. Positive values indicate days where the majority of cryptocurrencies moved in one direction (e.g., price increases), while negative values indicate collective movement in the opposite direction. The PC1 time series can be interpreted as a latent factor summarizing the overall sentiment or systemic trends in the cryptocurrency market over time.

Essay IV

Household Expectations and Dissent Among Policymakers

This paper has been published as:

Grebe, M. and Tillmann, P. (2025). Household expectations and dissent among policymakers, *European Journal of Political Economy* **86**: 102638.
<https://doi.org/10.1016/j.ejpoleco.2024.102638>

This paper was presented at the following workshops and international conferences:

- i. 4th Behavioral Macroeconomics Workshop: Heterogeneity and Expectations in Macroeconomics and Finance (online)
Date: June 2022
Presenter: Moritz Grebe
- ii. 53rd Annual Conference of the Money, Macro and Finance Society (Canterbury)
Date: September 2022
Presenter: Peter Tillmann
- iii. SUERF-Banca d'Italia-ECB-EIB conference: The use of surveys for monetary and economic policy (Rome)
Date: April 2023
Presenter: Moritz Grebe
- iv. Workshop on Challenges for Monetary Policy in Times of High Inflation (Hannover)
Date: September 2023
Presenter: Peter Tillmann
- v. VfS Annual Conference 2023 (Regensburg)
Date: September 2023
Presenter: Moritz Grebe

vi. 4th BOC-ECB-FRBNY Conference on Expectations Surveys (Ottawa)

Date: October 2023

Presenter: Moritz Grebe

Acknowledgement

We thank the editor of the EJPE and three anonymous reviewers for helpful comments. We also thank Olivier Armantier from the Federal Reserve Bank of New York for sharing his SAS code, the team from the Bundesbank Online Panel Households for their support and Klodiana Istrefi for discussing our paper. Lena Dräger, Michael Ehrmann, David Finck, Paul Hubert, Michael Lamla, Alistair Macaulay, Giang Nghiem, Ben Wang and seminar participants at Giessen, the Money, Macro and Finance Group (Canterbury, 2022), the 4th Behavioral Macroeconomics Workshop (Bamberg, 2022), the SUEF-Banca d'Italia-ECB-EIB conference "The use of surveys for monetary and economic policy", the workshop in Hannover on "Challenges for Monetary Policy in Times of High Inflation" and the 4th BOC-ECB-FRBNY Conference on Expectations Surveys (Ottawa, 2023) provided valuable feedback.

Household Expectations and Dissent Among Policymakers

Moritz Grebe*

Peter Tillmann[†]

Abstract

This paper studies the impact of dissent in the ECB's Governing Council on uncertainty surrounding households' inflation expectations. We conduct a randomized controlled trial using the Bundesbank Online Panel Households. Participants are provided with alternative information treatments concerning the vote in the Council, e.g. unanimity and dissent, and are asked to submit probabilistic inflation expectations. The results show that the vote is informative. Households revise their subjective inflation forecast after receiving information about the vote. Information about unanimity or dissent causes a wider individual distribution of future inflation for those households that were relatively certain before the treatment. For the remaining 60% of households, this information reduces uncertainty. Information about dissent increases uncertainty more than information about a unanimous vote, though the difference is not statistically significant. A unanimous vote unambiguously reduces inflation uncertainty for households with anchored inflation expectations.

Keywords: central bank communication, disagreement, inflation expectations, randomized controlled trial, survey

JEL classification: E52, E43, E32

*Justus Liebig University Giessen, Email: moritz.grebe@wirtschaft.uni-giessen.de

[†]Justus Liebig University Giessen, Email: peter.tillmann@wirtschaft.uni-giessen.de

1 Introduction

In recent years, central banks such as the Federal Reserve (Fed) or the European Central Bank (ECB) provide an array of information to the public. This involves press releases about monetary policy decisions, regular post-meeting press conferences, projections about future macroeconomic variables, speeches by senior central bankers and analytic reports about economic developments. Central banks acknowledged the adjustment of public expectations as a key transmission channel of monetary policy and try to manage these expectations. The provision of information by the central bank is seen as a stabilization and coordination device for private expectations. This aim became all the more relevant when many advanced economies hit the effective lower bound on nominal interest rates. When the conventional short-term interest rate was no longer available in order to implement further stimulus, central banks engaged in forward guidance to raise inflation expectations. Recently, central banks broadened their communication efforts and not just address investors and professional central bank watchers, but also households.¹

One important piece of information is the vote in the decision making body of the central bank. In case of the ECB, this is the Governing Council (GovC). While the Fed and the Bank of England (BoE) are very transparent about the vote in their committees and publish the formal vote together with the post-meeting press release, the ECB remains opaque.² The press release and the ECB president's Introductory Statement to the press conference remain silent on the vote. When asked by journalists, the ECB president only communicates a thin assessment of his or her reading of the GovC's majority using codewords such as "consensus" or "overwhelming support" to communicate the fact that some members dissented. The ECB does not reveal information about the direction of dissent or the name(s) of the dissenting member(s). At the same time, disagreement among members of the GovC is often headline news in the media.

In this paper, we study the effect of the vote in the GovC on the inflation expectations of households. For that purpose, we use a randomized controlled trial (RCT) design, which mimics the information provision during ECB press conferences. The RCT is implemented using the Bundesbank Online Panel Households (BOP-HH), a monthly household survey conducted by the Deutsche Bundesbank among more than 3,000 households. The design of the Bundesbank survey resembles the Survey of Consumer Expectations run by the Federal Reserve Bank of New York (Armantier et al., 2013,

¹See Haldane et al. (2021) and Ehrmann and Wabitsch (2022) for analyses of central bank communication with a broader audience as opposed to a small group of experts.

²See Rieder (2022) for a survey on central bank decision-making by committees.

Armantier et al., 2017).

The survey experiment is the following: survey participants receive information about an ECB decision to keep interest rates unchanged at 0%. In a first stage, each participant is asked to submit her minimum and maximum inflation forecast. In a second step, we provide three information treatments and request participants to submit their probabilistic inflation projections. Each participant assigns probabilities to alternative inflation bins. Besides a control group, which receives no additional information, one group receives the information that the ECB decision was unanimous. Another group receives the information that there was dissent in the GovC. A fourth group learns that the decision was unanimous despite different views in the council. We use the method introduced in Engelberg et al. (2009) and Manski (2018) to infer the second moments of each participant's individual distribution of future inflation. This set-up allows us to derive the *causal* effect of unanimity and dissent, respectively, on the second moments such as the interquartile range or the standard deviation of individual distributions about future inflation, which are reflecting individual uncertainty about the inflation outlook.

We obtain four findings. First, information about the vote among policymakers is informative. Households receiving information about the vote revise their first-stage inflation forecasts more strongly relative to the control group which receives no information about the vote. A potential mechanism could be that the vote is informative about future monetary policy, which is why respondents change their assessment of future inflation.

Second, the effect of the information about the vote on uncertainty depends on the degree of pre-treatment uncertainty. Receiving information about either unanimity or dissent increases inflation uncertainty for the bottom 40% of households in the distribution of pre-treatment uncertainty. Hence, household which were relatively certain before being informed about the vote face higher uncertainty after the information treatment. For the remaining 60% of households, inflation uncertainty falls upon receiving the information. The change in individual inflation uncertainty caused by the treatment among policymakers should have real economic consequences: as all survey participants face the same interest rate, differences in individual distributions of expected inflation translate into different subjective distributions of real interest rates.

Third, when receiving information about dissent, relatively certain households revise their inflation projections more strongly than households receiving information about a unanimous vote. Households receiving information about dissent have a wider distribution of future inflation as reflected in the interquartile range or

the standard deviation of their probabilistic forecast. However, the effect between receiving information about a unanimous vote and about dissent, respectively, is not statistically significant. This suggests that it is the revelation of the fact that there is a vote in the GovC that causes inflation uncertainty to increase, not the nature of the votes cast. Hence, stressing the unanimous nature of the policy decision does not generally lead to a lower inflation uncertainty.

Some of these findings are rather complex and not easy to interpret in light of specific economic mechanisms. This is why we shed light on the role of inflation anchoring, which gives rise to our fourth result: we infer participants' long run inflation forecast from another question in the same survey round. This allow us to distinguish between respondents with anchored and unanchored inflation expectations, i.e. with long-run forecasts aligned with the ECB's inflation target or not. We use this anchoring dummy to construct interaction terms. Interestingly, we find an unambiguous effect of unanimity reducing inflation uncertainty of participants with anchored expectations. Our results remain unchanged if we make different assumptions about the shape of individual probability distributions in the first stage of the survey. While in the baseline model we assume an isosceles triangle spanned between the minimum and maximum inflation expectation, we also allow for asymmetric distributions in our robustness section. Specifically, we impose the skewness from a question earlier in the survey, which is routinely asked by the Bundesbank, or the average skewness from the control group.

Our results shed light on the role of the vote in GovC for the ability of the ECB to manage household expectations. Communicating a unanimous decision does not reduce the uncertainty of households around the inflation path for all participants. In fact, relative to remaining opaque about the vote, releasing *any* information about the vote, i.e. either information about unanimity or dissent, can indeed reduce uncertainty for households with high prior uncertainty.

In our survey, we deal with the inflation expectations of households. Though households might not be as well informed about the details of monetary policy as experts, they are aware of rifts in the GovC. This is particularly true for households in Germany, where the disagreement between the president of the Bundesbank, who is a member of the GovC, and the ECB president was headline news for more than a decade.

The remainder of this paper is organized as follows: Section two surveys the related literature. Section three provides some background on the vote in central bank decision making bodies. Section four introduces the survey and section five presents the regression model. The results are discussed in section six with the robustness of

the findings analyzed in section seven. Section eight draws conclusions. An online appendix contains additional material.

2 Related literature

This paper contributes to multiple strands of the literature. First, we add to the large literature on information provision experiments in RCTs, e.g. Coibion et al. (2024), Coibion et al. (2022), Coibion et al. (2023), Hoffmann et al. (2022) and many others.³ In particular, researchers use RCT designs in order to understand the impact of policy communication on expectations. Coibion et al. (2022) use data from the Nielsen Homescan Panel and provide participants with eight different information treatments, each of which contains a specific piece of information about monetary policy. The authors establish that policy communication has a significant effect on expectations and also on households' willingness to spend. In contrast to this paper, the authors focus on the first moments of the probability distributions of inflation only. Our focus lies on the second moments of the distributions.

Coibion et al. (2023) estimate the causal effect of the Fed's forward guidance on inflation expectations. The paper finds that the pass-through of forward guidance to real interest rates is limited. Coibion et al. (2024) study the effect of uncertainty on household spending. The authors confront participants with exogenous variation in individual uncertainty and ask participants to submit their spending intentions. As a key result, they find that higher uncertainty reduces spending. Other papers confirm that consumers with a higher degree of idiosyncratic inflation uncertainty have a lower attitude towards spending on durable goods (Binder, 2017; Kostyshyna and Petersen, 2024; Georgarakos et al., 2024).

Second, this paper adds to the literature on central bank communication and inflation expectations of households and firms. Hubert (2015) argues that central bank forecasts do indeed affect private inflation forecasts. Anderes et al. (2021) argue that ECB communication during the press conference adds valuable information to market participants. In contrast, Lustenberger and Rossi (2020) cannot establish firm evidence of an effect of central bank communication on forecast accuracy and forecast dispersion of professional forecasters. Likewise, Jain and Sutherland (2020) show that central bank projections do not cause a smaller dispersion of private-sector forecasts.

Lamla and Vinogradov (2019) conduct a survey among UK households before and

³Coibion et al. (2020) stress the role of policy communication for inflation expectations and underline the importance of household expectations.

after decisions of the Bank of England. This design allows a clean identification of the effects of central bank communication. The authors show that announcements have small effects on the first moments of macroeconomic variables. However, the effects of central bank announcements differ in the cross-section of households. Upon receiving the information from the central bank, better informed survey participants make smaller forecast errors. Lewis et al. (2020) use a daily survey of U.S. households from Gallup. They show that households are attentive to monetary policy decisions. Their economic confidence falls after a positive monetary policy surprise. However, this effect mostly stems from interest rate decisions. Changes in forward guidance are much less effective. As in Lamla and Vinogradov (2019), the authors find large cross-sectional differences in the effectiveness of policy communication. Conrad et al. (2022) support the notion of cross-sectional variation in the responses to policy. They use survey data from Germany and condition the effect of central bank information on idiosyncratic characteristics. In particular, the authors find that individual lifetime experiences affect the adjustment of expectations in light of new information. Kryvtsov and Petersen (2021) conduct a laboratory experiment, which allows a clean identification. Central bank communication has a stabilizing effect on forecasts. Again the authors find evidence for cross-sectional heterogeneity in the responsiveness of participants to policy communication.

De Fiore et al. (2021) trace the responses of households to decisions of the Fed's Federal Open Market Committee (FOMC) using the Survey of Consumer Expectations. Comparing responses before and after the policy announcements, the authors find little evidence of an effect of policy decisions on inflation expectations. Rast (2021) designs a survey among German households that elicits expectations before and after ECB announcements. He finds that households respond to the release of information about conventional monetary policy, i.e. about interest rates, but remain insensitive to information about forward guidance and asset purchase programs, respectively. Brouwer and de Haan (2022) conduct an RCT among Dutch households. Households which receive information about the central bank's actions, in particular the ECB choice of policy instruments, expect inflation to be closer to the ECB's inflation target.

In contrast to this literature, which looks mostly on first moments of households' inflation expectations, we concentrate on second moments such as the interquartile range of the individual distributions of inflation projections. Using micro data from the ECB's Survey of Professional Forecasters, Fernandes (2025) is one of the few papers studying second moments. She shows that ECB communication about forecasts reduces individual forecast uncertainty.

Third, this paper is related to the literature that studies the effect of dissenting votes and cacophonous communication of policy makers on market participants. This literature applies event-study and time series model rather than RCTs. Ehrmann and Fratzscher (2013) shed light on the role of dispersed communication by committee members. Their evidence suggests that speaking with one voice reduces the prediction errors of the public. The effect of formal dissent in the FOMC is studied by Madeira and Madeira (2019). Their findings show that the response of the stock market differs between unanimous policy decisions and decisions with dissent. Equity prices increase if a policy decision of the FOMC is unanimous and decrease if dissenting votes are cast. Tillmann (2021) constructs an index of dissent for the ECB and shows that dissent weakens the transmission of monetary surprises on long-term interest rates. These papers infer the consequences of dissent from time series data. Our paper suggests a much cleaner identification of the effect of the vote in an RCT setting.

3 Dissent in monetary policy committees

In this paper, we concentrate on dissent in the ECB's GovC. On a meeting day of the GovC, the ECB publishes a press release at 13:45 CET, which contains the latest policy decision. At 14:30 CET, the ECB president opens the press conference. After reading the Introductory Statement, which summarizes the Council's assessment of the economic situation and the policy change, both the ECB president and the vice-president are available for questions from journalists. The press conference is very informative: Ehrmann and Fratzscher (2009) find that information revealed during the ECB press conferences has a larger impact on asset prices than information revealed through the policy announcement. Information from the Q&A part of the press conference is particularly informative during volatile periods.

In the Q&A part, journalists often ask whether the policy decision was unanimous. Normally, the president responds by giving his or her assessment of the support the policy enjoys in the GovC. The president does not, however, reveal the name of the dissenter(s), nor the direction of dissent. Thus, the public can learn about the vote from information given to the press. Often the president gives a straightforward answer such as on July 06, 2006:

Question: "Was the decision unanimous today?"

Trichet: "Yes, very much."

In many other instances, the president reflects on the decision being supported "overwhelmingly" (e.g. August 03, 2006) or was taken "by consensus" (e.g. June 05,

2008).⁴ These characterizations imply that some members voted against the policy proposal. In addition, the president sometimes uses a sharper formulation, e.g. on September 06, 2012, when the exchange with a journalist reads as follows:

Question: "My question regards the vote today. Was it unanimous and, if not, what does it mean?"

Draghi: "Well, it was not unanimous. There was one dissenting view. We do not disclose the details of our work. It is up to you to guess."

In our survey experiment introduced below, we provide participants with information treatments that mimic the information about unanimity and dissent given in the press conference. In some press conferences, the ECB president stresses the unanimous vote but at the same time acknowledges different opinions among GovC members, e.g. in the following quote from February 05, 2009:

Trichet: "We were unanimous in taking our decision, which does not mean that we all have the same view."

Consistently, we allow for another treatment in our RCT which we label "unanimity despite different views".

Tillmann (2021) studies the transcripts of the ECB press conferences and constructs a binary index of dissent from the president's answers to journalists. In the 123 meetings considered in the analysis, he finds 25 meetings with at least one dissenting member. The frequency of dissent in the GovC is lower than for the Federal Open Market Committee (FOMC), the policy-making body of the Fed. Using the data set compiled by Thornton and Wheelock (2014), Tillmann (2021) reports a share of at least one dissenting vote of 40% for the 202 meetings between February 1994 and September 2018. For the BoE's Monetary Policy Committee, Riboni and Ruge-Murcia (2014) calculate a frequency of dissent of 63%. For the Swedish Riksbank, the authors find a dissenting vote in 38% of all policy meetings. Ruge-Murcia and Riboni (2017) study the decision making procedure at the bank of Israel and find a probability of at least one dissenting member of 31% between 2011 and 2015.

Thus, compared to other central banks dissent in the GovC is relatively infrequent. This diagnosis reflects the nature of policy making at the ECB as being highly consensual. In fact, ECB president Draghi argued that "Given the peculiar nature of the ECB, one of my objectives is that we have as much consensus as possible."⁵ Consequently, Blinder (2007) characterizes the GovC as a "genuinely-collegial

⁴The quotes from the press conferences are taken from the transcripts, which are available on the ECB's website at <https://www.ecb.europa.eu/press/pressconf/html/index.en.html>.

⁵See <https://www.ecb.europa.eu/press/key/date/2012/html/sp120224.en.html>.

committee". The relatively low frequency of dissent also implies that dissent is particularly noteworthy. In contrast, Blinder (2007) argues that there is "no (or negligible) public disagreement" among members of the GovC. However, the fierce disagreement between the presidents of the Bundesbank and the ECB in the decade after the 2008 financial crises indicates that the nature of the committee might have changed over time.⁶

Gerlach-Kristen (2004), Gerlach-Kristen and Meade (2010), Horváth et al. (2012), Neuenkirch (2013) and Riboni and Ruge-Murcia (2014) provide evidence for different central banks which is consistent with the notion that the vote in the decision making committee helps the public forecast future policy. This implies that dissent is an important piece of information that observers take into account when forming expectations.

4 The survey

We study the impact of the vote in the GovC on inflation expectations in a Randomized Controlled Trial (RCT). For that purpose, we contributed two questions to wave 19 (July 2021) of the BOP-HH. This survey is conducted regularly since 2020 as an online panel of several thousand households in Germany. Each month between 2,500 and 5,000 households, which are representative for the German population, are asked about their expectations of inflation, property prices and interest as well as their planned expenditures. The survey also collects information on the respondents' socio-demographic characteristics. Researchers have the chance to submit one-off questions in order to augment the set of regular questions.⁷

4.1 Design of the RCT

The RCT is conducted in two stages. In the first stage, the pre-treatment stage, households are prompted to submit their minimum and maximum inflation expectation. In the treatment stage, participants are randomly assigned to four groups: a control group and three treatment groups. In the following, we explain the survey design in detail.

Pre-treatment stage. We provide participants with the following information:

⁶Baerg and Krainin (2022) present a model in which members of the monetary policy committee bargain over the communication of the committee.

⁷Details about the survey are available at <https://www.bundesbank.de/en/bundesbank/research/survey-on-consumer-expectations>. The survey is conducted in German. Here, we use the translation available on the Bundesbank website.

"Assume that the European Central Bank (ECB) is aiming for an annual inflation rate of 2% over the medium term. Please also assume that the inflation rate is 1% in 2021. The ECB Governing Council decides to keep the policy rate at 0%."

Participants could click on two information boxes in order to acquire additional information about the meaning of "Governing Council" and the "policy rate". At the time of the survey (July 2021), the latest available inflation rate was 1.9% (June 2021) and the interest rate was effectively at zero. We then ask the following question:

"In your opinion, how high will the inflation rate be at least over the next one to two years? And at most?"

Figure (C.1) shows a screen shot of the survey question. The answers provide us with the minimum and maximum expected inflation rate for each participant.

Treatment stage. We repeat the information about the inflation target, the current inflation rate and the policy decision from the pre-treatment stage. We then randomly assign each participant to one of four alternative groups, three of which correspond to a specific information treatment. After receiving the information treatment, participants are asked the following probabilistic question:

"In your opinion, how likely is it that the rate of inflation will be as follows over the next one to two years? Note: The aim of this question is to determine how likely you think it is that something specific will happen in the future. You can rate the likelihood on a scale from 0 to 100, with 0 meaning that an event is completely unlikely and 100 meaning that you are absolutely certain it will happen. Use values between the two extremes to moderate the strength of your opinion. Please note that your answers to the categories have to add up to 100."

Participants are offered ten bins ranging from a deflation rate of 12% or higher to an inflation rate of 12% or higher. The width of the outer bins is four percentage points, while the width of the inner bins is two percentage points. The number of bins and the varying width of bins is identical to the Survey of Consumer Expectations run by the Federal Reserve Bank of New York. From these answers, we obtain the individual probability distribution of future inflation as explained below. We write "one or two years" because this is the typical horizon over which monetary policy affects inflation. We randomly allocate participants into groups and provide the treatments (T) given in Table (1).⁸ The first group ($T = 1$) corresponds to our control group. We do not

⁸Screenshots of these survey questions are available in Figures (C.2) to (C.5).

Table 1: Groups

$T = 1$	control group
$T = 2$	unanimity (<i>una</i>)
$T = 3$	dissent (<i>dis</i>)
$T = 4$	unanimity despite different views (<i>unadis</i>)

provide any additional information to participants in this group. Survey participants in the first treatment group ($T = 2$) receive the following information:

"The ECB President informs the media that this was a unanimous decision."

As discussed before, this treatment corresponds to the information provided during several press conferences in the past. In the regression model below, we index this treatment by *una*.

The second treatment group ($T = 3$) receives the following information:

"The ECB President informs the media that this was a majority decision, i.e. there were dissenting votes."

This is the treatment we are most interested in. In the regression model below, we index this treatment by *dis*.

Finally, the last group ($T = 4$) receives the following information:

"The ECB President informs the media that this was a unanimous decision despite different opinions."

In the regression model below, we index this treatment by *unadis*. We include this specific treatment for two reasons. First, as mentioned before, we find ECB press conferences during which the president classifies a decision as unanimous and at the same time mentions different opinions of GovC members. Second, it could be argued that unanimity despite different views is particularly informative as this implies that the full range of arguments has been put forward during the meeting. As a consequence, households could interpret such a decision as particularly thoughtful and, hence, informative for inflation expectations.

Our treatments closely reflect the provision of information during the ECB press conferences. We confront participants with alternative treatments, each of which is hypothetical. We do not replicate one specific ECB meeting day, but rather the range of information sets of ECB observers on meeting days over the past two decades. The hypothetical statements of the ECB president we mention in the treatments have all been made before during other ECB press conferences. Thus, we provide survey

participants with hypothetical, but well-known scenarios. Because the treatment are familiar from previous press conference, we think the treatments are not too complex. Nevertheless, asking hypothetical questions might affect the external validity of the results.

In our robustness analysis, we also draw on survey question CM004, which is routinely asked by the Bundesbank. In this question, participants are asked:

In your opinion, how likely is it that the rate of inflation will be as follows over the next twelve months?

Participants are offered the identical set of bins that we use in our questions from the treatment stage and are asked to allocate 100 points across the bins. This question is very similar to our question from the treatment stage. There are two main differences, though: the first difference is that the horizon of the inflation forecast is specified as one year, while in our question the horizon is "the next one to two years". The second difference is that in question CM004, participants are not provided with information about the inflation target, the current inflation rate and the current policy rate.


2927 participants completed the survey. Survey participants also provide socio-demographic information about gender, age, household income, their employment status, the years of schooling and much more. We use this information as control variables in our regression model.

4.2 Inferring individual probability distributions

In order to estimate the effect of dissent on individual inflation uncertainty, we first need to fit distributions to the respondents' answers. We fit a triangular distribution to the answers from the pre-treatment stage. Figure (1) shows that the triangular distribution is defined by its lower limit a , upper limit b and mode c . We set a and b to the minimum and maximum expectation elicited from the survey. Because we do not know c , we assume the triangular distributions to be symmetric. This implies that the mean, median and mode of the triangular distribution equal $\frac{a+b}{2}$. Below, we show that the results are robust with respect to assuming a symmetric distribution at the pre-treatment stage.

In order to fit distributions to the histograms of the answers from the treatment stage, we follow the method used in the Survey of Consumer Expectations from the Federal Reserve Bank of New York. This approach in turn is based on Engelberg et al. (2009). The distribution of respondents who use two bins is assumed to be a symmetric triangular distribution. Either the upper limit a or the lower limit b is set

Figure 1: Probability density function of a symmetric triangular distribution



Notes: The lower (upper) limits are denoted by a (b). The mode is given by c .

to the outer limit of the bin in which the forecaster places more than 50% probability mass. The remaining parameters must be fitted by taking the height and center of the triangle into account. For example, if a respondent puts 10% probability in the interval $[0\%, 2\%]$ and 90% in the interval $[2\%, 4\%]$, the upper limit b of the distribution is fixed to 4. a is fitted to 1.42. Because the distribution is symmetric, c equals $\frac{a+b}{2} = 2.71$. Note that this procedure can only be applied to respondents using two adjacent bins.

If a forecaster places weights in more than two bins, a generalized β distribution is fitted. This distribution is not necessarily symmetric. The bins, moreover, do not have to be adjacent. For a more detailed description of the method, see Armantier et al. (2017). We deviate from this method by fitting symmetric triangles to histograms with only one bin instead of fitting a uniform distribution. This is in line with Engelberg et al. (2009) and consistent with our method used in the pre-treatment step. The parameters a and b of the triangular distribution are then set to the limits of the interval. Once we fitted these distributions to the individual responses, we can compute the mean, the standard deviation and the interquartile range, which we use as dependent variables in our regression model.

5 Empirical model

Coibion et al. (2023) suggest that regressing the change of a variable between the

pre-treatment and treatment stage on a dummy variable for each treatment may not always be feasible to investigate treatment effects. Since our treatments are very specific, one might expect little variation in average effects. Accordingly, we find no significant effects of the treatments on average effects when estimated as described (Table (2)). Therefore, we follow the literature (e.g. Coibion et al., 2024; Coibion et al., 2023) and test the updating of information in light of a noisy public signal. Suppose the expectation formation of forecaster i can be described by the Bayesian updating rule:

$$x_i^{post} = (1 - G) \times x_i^{pre} + G \times signal \quad (1)$$

Here, the expectation x_i^{post} is a weighted average of the prior expectations x_i^{pre} and a new signal. When applied to our survey, this signal corresponds to information from the treatments. If the treatment provides information that is informative relative to the prior expectation, the respondent should increase their Kalman gain G i.e. assign more weight to this new information and thereby less weight to prior information. This effect can be measured by comparing the weight on prior expectations between respondents from the control group, who do not receive any new information, and respondents from the treatment groups. We estimate:

$$x_i^{post} = c + \alpha x_i^{pre} + \sum_{T=2}^4 \beta_T I_i^{(T)} + \sum_{T=2}^4 \gamma_T \left(I_i^{(T)} \times x_i^{pre} \right) + \omega_i, \quad (2)$$

where x_i^{post} is the individual expectation in the treatment stage, x_i^{pre} denotes expectations from the pre-treatment stage, I_i^T is an indicator variable equal to one if respondent i receives treatment T and zero otherwise. x_i^{pre} and x_i^{post} , respectively, denote the interquartile range or the standard deviation of the individual distribution. The control group ($T = 1$) is our reference category such that we include three treatments. The error term is ω_i . The coefficients are denoted α , β and γ . We will also estimate a specification with an additional vector of respondent-specific control variables.

α can be interpreted as the weight the control group attaches to prior information. Because the control group is not provided with any new information, we expect α to be close to one. For a given T , $\alpha + \gamma_T$ is the weight attached to prior information, which corresponds to $1 - G$ from Equation (1). The coefficient γ_T should be more negative for more informative treatments. If the treatment is informative, the respondent should place less weight on x_i^{pre} and more weight on the treatment. The "level effects" (Coibion et al., 2023) β_T can be positive or negative.

We expect $T = 3$ to be particularly informative. If the respondent receives information on dissent in the Governing Council, she should most strongly revise her expectations.

Hence, we expect γ_{dis} to be the most negative of the three γ coefficients. We follow Coibion et al. (2024) and Coibion et al. (2023) and estimate the equation using robust Huber regressions.

In every wave of the Bundesbank Survey, participants are asked about their point estimate of the inflation rate for either the next five or the next ten years (question CQ006A/B). This question is included in the core part of the survey, i.e., before our submitted questions, ensuring that their responses are not influenced by our interventions. This setup allows us to examine whether the effect of our treatments varies depending on whether respondents have well-anchored long-term inflation expectations. We define inflation expectations as anchored if a participant’s five- or ten-year estimate corresponds exactly to 2%, the ECB’s inflation target.⁹ We estimate:

$$\begin{aligned}
x_i^{post} = & c + \alpha x_i^{pre} + \sum_{T=2}^4 \beta_T I_i^{(T)} + \sum_{T=2}^4 \gamma_T \left(I_i^{(T)} \times x_i^{pre} \right) + \delta anc_i + \alpha^{anc} (x_i^{pre} \times anc_i) \\
& + \sum_{T=2}^4 \beta_T^{anc} \left(I_i^{(T)} \times anc_i \right) + \sum_{T=2}^4 \gamma_T^{anc} \left(I_i^{(T)} \times x_i^{pre} \times anc_i \right) + \omega_i,
\end{aligned} \tag{3}$$

where anc_i is a dummy variable equal to one if respondent i has anchored long-term inflation expectations and zero otherwise. For individuals in group T with unanchored expectations, the level effect is β_T and γ_T represents the weight these respondents place on prior expectations relative to the control group. Thus, for those with anchored expectations, the level effect becomes $\beta_T + \beta_T^{anc}$ and the weight they place on their prior relative to the control group equals $\gamma_T + \gamma_T^{anc}$.

6 Results

We report our results for two alternative corrections of outliers. In the first correction, we drop responses that put 100% probability on either of the outer bins in the treatment stage, i.e. more than 12% inflation or deflation, and responses of inflation or deflation of more than 100% in the pre-treatment stage. Further, we noticed there are participants whose responses are inconsistent. There are some respondents in the control group who have expectations from the pre-treatment stage that are very

⁹Although the robust mean of the 5-year inflation expectations is slightly higher than that of the 10-year expectations, the difference is not statistically significant. Therefore, we consider it appropriate to combine these expectations into a common anchoring variable.

different from those from the treatment stage. This may be due to general difficulties with formulating density forecasts in the treatment stage or by simply confusing inflation and deflation. Inconsistent answers could bias our results. Therefore, we additionally exclude respondents in our second correction whose change in the interquartile range between the pre-treatment and the treatment stage is larger than the 95th percentile of all changes. While we present the main results for both alternative corrections of outliers, we use only the first outlier correction in the discussion of the pre-treatment results in the next subsection.

6.1 Pre-treatment expectations

We study the determinants of the prior expectations of our respondents in Table (3).¹⁰ For that purpose, we employ Huber robust regression models, using the interquartile range, the standard deviation, and the mean expectation from the pre-treatment stage as dependent variables.

We group the participants based on age, household income, and city size. A dummy variable is equal to one if a respondent belongs to a particular group and zero otherwise. The first group for each variable is used as the omitted variable. Additionally, we employ separate dummy variables for the employment status and residency in East Germany. While we build groups for the variables age, income, and city size, we include household size and years of education as continuous variables. Since participants are not asked directly about the number of years of their education, but rather about their highest level of educational attainment and their level of vocational training or university degree, we must calculate the years of education. We follow the procedure outlined in the German Socio-Economic Panel conducted by the German Institute for Economic Research (DIW Berlin).¹¹ We will later employ all these variables as controls in our baseline regressions.

Our findings, see Table (3), indicate that the interquartile range and standard deviation are significantly smaller for respondents aged 35 and older compared to younger respondents. Hence, older respondents express less uncertainty in their inflation expectations. The absolute value of the coefficients increases for older respondent groups. Moreover, uncertainty is significantly higher for respondents living in larger households and those residing in East Germany.

Regarding mean inflation expectations, age groups comprising respondents aged 45 and older are associated with higher mean expectations relative to younger groups.

¹⁰We present further descriptive statistics in Table (4).

¹¹You can find this procedure here: https://www.diw.de/documents/dokumentenarchiv/17/diw_01.c.410636.de/pgen-v28.png.

Furthermore, the coefficients for household size and East Germany are significantly positive. The coefficient for years of education is marginally significant and negative, though the value is relatively small. All other coefficients are insignificant.

6.2 Predictability of treatment status

We report our findings both with and without the aforementioned set of respondent-specific control variables. Since the allocation of respondents into groups is random, the group-specific coefficients should not be systematically affected by respondent-specific characteristics. Table (5) examines whether the allocation into groups can be predicted by a probit regression based on our control variables, which should not be the case if the randomization has been successful.

The table reveals that only for those respondents that receive the *una* treatment, two groups within the city variable are significant at the 10% level. All other control variable remain insignificant. We take this as strong evidence that the treatment has indeed been randomly assigned. Nonetheless, we will demonstrate in our baseline results that the inclusion of control variables does not significantly alter our findings. However, incorporating control variables into our regression specification should reduce the standard errors of the estimates.

6.3 The effect on the uncertainty of forecasters

We now present the main results of the paper: the effects of the treatments on the uncertainty of forecasters measured as either the interquartile range or the standard deviation of the individual probabilistic forecasts. As dissent in the GovC should primarily affect the second moments of forecasts, these results are our baseline findings. In the second step, we report the results for the mean of the individual probabilistic forecasts. That is, we assess whether dissent effects first moments as well.

In our baseline model, we employ the interquartile range of the individual probabilistic distribution for each forecaster i as our dependent variable x_i^{post} . x_i^{pre} is the corresponding interquartile range from the pre-treatment stage. The estimated coefficients are reported in Table (6). We obtain an estimated α which is positive and relatively close to one. All four estimates of α are larger than 0.8. Hence, respondents in the control group put a large weight on their prior.

For our research question, the estimated β and γ coefficients are the most interesting coefficients. We find that β_{una} and β_{dis} are significantly positive across all four specifications, with β_{dis} consistently larger than β_{una} . This difference, however, lacks

statistical significance.

It is important to note that using only the level effects, i.e. the estimated β coefficients, to discuss the average effect of our treatments is insufficient. Since we also account for prior uncertainty, the total effect of treatment T on uncertainty can be expressed as

$$\frac{\partial x_i^{post}}{\partial I_i^{(T)}} = \beta_T + \gamma_T x_i^{pre}. \quad (4)$$

Hence, β_T represents the unconditional effects of the treatment on uncertainty, i.e. the impact on respondents with zero prior uncertainty. In that case, a positive β_T can be interpreted as the corresponding treatment T increasing uncertainty. To gauge the full effect of the treatment, we need to look at the unconditional effect plus the effect conditional on prior uncertainty. Informative treatments exhibit a negative γ_T , implying that despite a positive β_T treatments can reduce uncertainty for respondents with high prior uncertainty.

Generally, the treatment effect is positive if $\beta_T + \gamma_T x_i^{pre} > 0$. With $\gamma_T < 0$, this implies a critical level of

$$x_{i,crit,T}^{pre} < -\frac{\beta_T}{\gamma_T}. \quad (5)$$

For specification I in Table (6), the critical value for the treatment *una* is 0.54, and 42.04% of respondents have a prior level of uncertainty below this value. This implies that the treatment *una* increases uncertainty for this proportion of respondents, while it reduces uncertainty for the remaining 60%. Importantly, no significant differences exist between the treatments *una* and *dis*.¹² This observation remains consistent across other specifications as well, with both treatments *una* and *dis* increasing uncertainty for around 40% of the participants. The last rows of each regression table reports the critical value of prior uncertainty at which the effect of the treatments change sign.

Figure (2) illustrates that there is a critical value for prior expectations where the sign of the treatment effect changes. Each of the three lines maps pre- into post-treatment uncertainty for the *una* and *dis* treatment as well as the control group.

The estimated γ coefficients show the impact of the information treatment on the dependent variable conditional on the extent of uncertainty in the pre-treatment stage. For all four specifications, the estimated γ_{una} and γ_{dis} are significantly negative.

¹²The treatment effect of T is given by $\beta_T + \gamma_T x_i^{pre}$. Consequently, the test for differences between treatments *una* and *dis* assesses whether $\beta_{una} - \beta_{dis} + (\gamma_{una} - \gamma_{dis})x^{pre} = 0$. Since this difference depends on prior uncertainty, we test the hypothesis for various levels of prior uncertainty. As shown in Table (8), the hypothesis *una* = *dis* cannot be rejected for all specifications where the interquartile range is the dependent variable.

This means that individuals with a high degree of uncertainty in the pre-treatment stage relative to the control group are particularly sensitive to the provision of information. Put differently, the treatments are informative such that respondents put significantly less weight on their prior. The information about dissent in the GovC is more informative than information about a unanimous vote. We can reject the hypothesis $\gamma_{una} = \gamma_{dis}$ at high levels of significance. Upon receiving information about a unanimous vote despite opposing views among council members, survey respondents increase rather than reduce the weight attached to the distribution in the pre-treatment stage. Hence, this treatment is the least informative. We can conclude that dissenting votes in the GovC have an effect on inflation uncertainty among households. The households subject to information about dissent put less weight on their interquartile range from the first-stage relative to the control group. We also include a large set of control variables ranging from gender, age, income to the employment status and the years of education. The results show that older survey participants aged 45 and above exhibit less uncertainty about the evolution of inflation compared to their younger counterparts. Additionally, respondents residing in cities with over 5,000 inhabitants display greater uncertainty than those living in cities with fewer than 5,000 inhabitants. Participants with a higher number of years of schooling submit a significantly wider distribution of inflation expectations. Hence, inflation uncertainty increases with education. Most of the other control variables remain insignificant.


In an alternative regression model, we use the standard deviation of individual forecast distributions as the dependent variable. The results are shown in Table (7). The results are very similar to the baseline findings in Table (6). Among the "level effects", the information about dissent in the GovC is most powerful. This effect remains stable across all four alternative specifications reported in the table. In line with the baseline model, both unanimous decision and dissent treatments lead to an increase in uncertainty for approximately 40% of the respondents.

Again, the information about dissent is most informative across treatments. Agents receiving information about dissent revise their individual probability distribution more strongly compared to all other treatments. Hence, our finding remains unchanged: information about dissent and unanimity is informative for households, with dissenting views carrying greater informative value. There is, however, no significant difference between the average effects of the two treatments.¹³ Both treatments increase uncertainty for 40% respondents. As expected, all findings remain unchanged if we include control variables.


¹³We can only reject the hypothesis $una = dis$ for specification I (see Table (8)).

Figure 2: The effects of the treatments on inflation uncertainty


(a) Regression results Table (6, I)




(b) Regression results Table (6, III)



(c) Regression results Table (7, I)




(d) Regression results Table (7, III)



Notes: The graphs show visualizations of our regression results for the control group and the informative treatments (*una*, *dis*) as presented in Table (6) and Table (7). The horizontal (vertical) axis scales inflation uncertainty before (after) the informative treatments.

Figure 3: The effects of the treatments on inflation uncertainty: anchored vs. unanchored expectations



Notes: The graphs visualize our regression results for respondents with anchored and unanchored expectations, as presented in Table (9, II). We only include coefficients that are significant at the 10% level. The horizontal (vertical) axis scales inflation uncertainty before (after) the informative treatments.

Table (9) shows the regression results when we estimate equation (3) to study whether our results depend on the anchoring of long-term inflation expectations. We use the interquartile range and the standard deviation as dependent variables and employ two methods of outlier correction. For unanchored respondents ($anc_i = 0$), the effect is very similar to our baseline results: the level effects for *una* and *dis* are significantly positive, whereas no significant effect is observed for the *unadis* treatment. Further, only the treatments *una* and *dis* can be considered informative, as the corresponding γ coefficients are significantly negative. Figure (3) depicts the overall treatment effects resulting from specification I. As in the baseline model, information about a unanimous decision and dissent increases inflation uncertainty for 40% of the unanchored respondents, with the difference between the treatments not being statistically significant.

However, when examining participants with anchored expectations, the picture is different: they appear to react only to information about a unanimous decision. Specifically, $\gamma_{una} + \gamma_{una}^{anc}$ is significantly negative at the 10% level while $\gamma_{dis} + \gamma_{dis}^{anc}$ is not significantly different from zero. Since the level effect ($\beta_{una} + \beta_{una}^{anc}$) is insignificant, the sign of the overall effect for these participants does not depend on prior uncertainty. Figure (3) illustrates that unanimity leads to less uncertainty for survey participants with anchored long-term expectations.

7 Robustness

The baseline regression model compares second moments of individual distributions from the treatment-state to the distribution in the first stage. The first stage distribution, however, is based on the assumption of the inflation distribution being described by an isosceles triangle between the minimum and maximum future inflation rate. We now relax the assumption that the individual triangular distributions in the pre-treatment stage are symmetric. Thus, we allow for a skewed distributions. We replace the assumed isosceles triangle by two alternatives. The details of these computations are explained in the online appendix.

First, we draw on a question on the individual probabilistic inflation forecast asked earlier in the sample (question CM004). This question is routinely included in each wave of the BOP-HH survey. We assume the individual prior distributions in our pre-treatment stage to have the same skewness as the distributions in question CM004. The advantage of using answers to this question to infer the skewness of the distributions of inflation expectations is that the design of the answer categories is identical to our survey question: participants can allocate 100 points to ten bins of future inflation rates. The drawback is that the verbal framing of the question differs slightly from our question. Participants are not provided with information about current inflation and the recent policy decision and are asked to provide a probabilistic assessment of inflation over the next 12 months rather than over the next one to two years as in our survey questions. Hence, there are advantages and disadvantages of using answers to this question in order to infer the skewness of the prior distribution of expectations. We need to keep these limitations into account when interpreting the results. As we interpret the coefficients relative to the control group, however, the slight difference in the framing of the questions should not matter much.

In the second alternative, we assume the prior distribution between the minimum and maximum to have the average skewness of the control group. The control group is not provided any new information in the treatment stage, such that the skewness of their prior and posterior distributions should not differ. The advantage of this approach is consistency: we use the skewness based on exactly the same survey design. The disadvantage is that we have to rely on average skewness rather than the skewness of the individual distribution.

Table (10) shows the results for the interquartile range when we impose the skewness from survey question CM004. Our key findings remain unchanged: the estimated β_{dis} is significantly positive across all four specifications and is larger than β_{una} or β_{unadis} . Moreover, participants put less weight on their prior distribution as the information

about dissent is highly informative. Receiving information about dissent reduces the impact of the prior distribution much more than receiving information about a unanimous vote. For the standard deviation of the individual inflation distribution, see Table (11), the results also remain unchanged. Again, the "level effect" of dissent is positive and highly statistically significant. The "slope effect", i.e. the estimated γ coefficients, are significantly negative for participants provided with information about dissenting votes.

We now turn to the second alternative to model the prior distribution based on the mean skewness of the control group. Table (12) contains the estimated coefficients when the interquartile range is used as a dependent variable. Information about dissent raises participants' inflation uncertainty. This treatment is more effective than the provision of information about a unanimous vote. Again we find that participants put less weight on their prior distribution when receiving information about dissent compared to information about unanimity. These findings also prevail when we replace the dependent variable by the standard deviation of individual distributions as shown in Table (13).

We can conclude that the key results of this paper remain robust with respect to the modeling of the pre-treatment stage.¹⁴ For individual skewness from survey question CM004 or for average skewness from the control group, dissenting votes raise inflation uncertainty for about 40% of the survey participants.

8 Conclusions

Over the past decades, central banks have intensified their communication with professional market participants. Recently, central banks go beyond financial professionals and also try to reach private households. One element of the vast array of information provided by central banks is the vote on monetary policy. In the case of the ECB, the voting outcome is not officially announced. Rather, the ECB president characterizes her or his assessment of the views among committee members.

In this paper, we studied the impact of information about the vote in the GovC on inflation expectations. We run a survey with an RCT in order to elicit the causal effect of unanimity or dissent, respectively, on households' distributions of expected future inflation. Our first result suggests that the vote is very informative for households. Households revise their subjective distribution of inflation upon receiving information about the vote. A second result is that dissent raises households' uncertainty about

¹⁴This robustness also holds for our model that partitions participants into anchored and unanchored expectations.

inflation for the 40% of households with the lowest uncertainty before the information treatment. We further find that unanimity does not reduce the uncertainty of all households around the inflation path, but only for the 60% of households with high pre-treatment uncertainty. A unanimous decision unambiguously lowers inflation uncertainty for participants with anchored long-run inflation expectations. Hence, the vote in the GovC is a significant determinant of inflation uncertainty. To the extent that individual inflation uncertainty translates into uncertainty about the real interest rate, higher inflation uncertainty has real economic consequences.

Our findings underline the difficulty of designing the optimal degree of transparency about the vote. Whether or not information about the vote increases or reduces inflation uncertainty depends on the prior degree of uncertainty before the information is received. Relative to remaining opaque about the vote, releasing *any* information about the vote, i.e. either information about unanimity or dissent, can indeed reduce uncertainty for households with high prior uncertainty.

The current way the ECB communicates the vote in the GovC should be overhauled. Whether or not the ECB president mentions the vote in the GovC during the press conference and how she assesses the views of GovC members lies in the discretion of the president. This discretionary nature adds to the uncertainty about the vote. In addition, households and market participants have to wait for the release of the Monetary Policy Accounts three weeks after the meeting in order to get a glimpse of the potential reasons for dissenting views.

References

- Anderes, M., Rathke, A., Streicher, S. and Sturm, J.-E. (2021). The role of ECB communication in guiding markets, *Public Choice* **186**: 351–383.
- Armantier, O., Bruine De Bruin, W., Potter, S., Topa, G., Van Der Klaauw, W. and Zafar, B. (2013). Measuring inflation expectations, *Annual Review of Economics* **5**(1): 273–301.
- Armantier, O., Topa, G., Van der Klaauw, W. and Zafar, B. (2017). An overview of the Survey of Consumer Expectations, *Economic Policy Review* (23-2): 51–72.
- Baerg, N. R. and Krainin, C. (2022). Divided committees and strategic vagueness, *European Journal of Political Economy* **75**: 102240.
- Binder, C. C. (2017). Measuring uncertainty based on rounding: New method and application to inflation expectations, *Journal of Monetary Economics* **90**: 1–12.
- Blinder, A. S. (2007). Monetary policy by committee: Why and how?, *European Journal of Political Economy* **23**(1): 106–123.
- Brouwer, N. and de Haan, J. (2022). The impact of providing information about the ECB’s instruments on inflation expectations and trust in the ECB: Experimental evidence, *Journal of Macroeconomics* **73**: 103430.
- Coibion, O., Georgarakos, D., Gorodnichenko, Y., Kenny, G. and Weber, M. (2024). The effect of macroeconomic uncertainty on household spending, *American Economic Review* **114**: 645–677.
- Coibion, O., Georgarakos, D., Gorodnichenko, Y. and Weber, M. (2023). Forward Guidance and Household Expectations, *Journal of the European Economic Association* **21**(5): 2131–2171.
- Coibion, O., Gorodnichenko, Y., Kumar, S. and Pedemonte, M. (2020). Inflation expectations as a policy tool?, *Journal of International Economics* **124**: 103297.
- Coibion, O., Gorodnichenko, Y. and Weber, M. (2022). Monetary policy communications and their effects on household inflation expectations, *Journal of Political Economy* **130**(6): 1537–1584.
- Conrad, C., Enders, Z. and Glas, A. (2022). The role of information and experience for households’ inflation expectations, *European Economic Review* **143**: 104015.
- De Fiore, F., Lombardi, M. J. and Schuffels, J. (2021). Are households indifferent to monetary policy announcements?, *unpublished* (Bank for International Settlements).
- Ehrmann, M. and Fratzscher, M. (2009). Explaining monetary policy in press conferences, *International Journal of Central Banking* **5**: 41–84.

- Ehrmann, M. and Fratzscher, M. (2013). Dispersed communication by central bank committees and the predictability of monetary policy decisions, *Public Choice* **157**(1-2): 223–244.
- Ehrmann, M. and Wabitsch, A. (2022). Central bank communication with non-experts—A road to nowhere?, *Journal of Monetary Economics* **127**: 69–85.
- Engelberg, J., Manski, C. F. and Williams, J. (2009). Comparing the point predictions and subjective probability distributions of professional forecasters, *Journal of Business & Economic Statistics* **27**(1): 30–41.
- Fernandes, C. M. (2025). Ecb communication as a stabilization and coordination device: Evidence from ex ante inflation uncertainty, *International Journal of Central Banking* **21**(1): 83–135.
- Georgarakos, D., Gorodnichenko, Y., Coibion, O. and Kenny, G. (2024). The causal effects of inflation uncertainty on household’s beliefs and actions, *NBER Working Paper No. 33014*.
- Gerlach-Kristen, P. and Meade, E. E. (2010). 2010-16 Is there a limit on FOMC dissents?, *American University report*.
- Gerlach-Kristen, P. (2004). Is the MPC’s Voting Record Informative about Future UK Monetary Policy?, *The Scandinavian Journal of Economics* **106**(2): 299–313.
- Haldane, A., Macaulay, A. and McMahon, M. (2021). The three E’s of central-bank communication with the public, *unpublished* (University of Oxford).
- Hoffmann, M., Moench, E., Pavlova, L. and Schulte-frankfeld, G. (2022). Would households understand average inflation targeting?, *Journal of Monetary Economics* **129**: S52–S66.
- Horváth, R., Smidkova, K. and Zápál, J. (2012). Central banks’ voting records and future policy, *International Journal of Central Banking* **8**(4): 1–19.
- Hubert, P. (2015). Do Central Bank Forecasts Influence Private Agents? Forecasting Performance versus Signals, *Journal of Money, Credit and Banking* **47**(4): 771–789.
- Jain, M. and Sutherland, C. S. (2020). How do central bank projections and forward guidance influence private-sector forecasts?, *International Journal of Central Banking* **16**(5): 179–219.
- Kostyshyna, O. and Petersen, L. (2024). The effect of inflation uncertainty on household expectations and spending, *NBER Working Paper No. 32939*.
- Kryvtsov, O. and Petersen, L. (2021). Central bank communication that works: Lessons from lab experiments, *Journal of Monetary Economics* **117**: 760–780.
- Lamla, M. J. and Vinogradov, D. V. (2019). Central bank announcements: Big news for little people?, *Journal of Monetary Economics* **108**: 21–38.

- Lewis, D. J., Makridis, C. and Mertens, K. (2020). Do monetary policy announcements shift household expectations?, *FRB of New York Staff Report* (897).
- Lustenberger, T. and Rossi, E. (2020). Does central bank transparency and communication affect financial and macroeconomic forecasts?, *International Journal of Central Banking* (16): 153–201.
- Madeira, C. and Madeira, J. (2019). The effect of FOMC votes on financial markets, *Review of economics and statistics* **101**(5): 921–932.
- Manski, C. F. (2018). Survey measurement of probabilistic macroeconomic expectations: Progress and promise, *NBER Macroeconomics Annual* **32**: 411–471.
- Neuenkirch, M. (2013). Predicting Bank of England’s asset purchase decisions with MPC voting records, *Applied Economics Letters* **20**(13): 1275–1278.
- Rast, S. (2021). Central bank communication with the general public: survey evidence from Germany, *unpublished* (European University Institute).
- Riboni, A. and Ruge-Murcia, F. (2014). Dissent in monetary policy decisions, *Journal of Monetary Economics* **66**: 137–154.
- Rieder, K. (2022). Monetary policy decision-making by committee: why, when and how it can work, *European Journal of Political Economy* **72**: 102091.
- Ruge-Murcia, F. and Riboni, A. (2017). Collective versus individual decision-making: A case study of the Bank of Israel Law, *European Economic Review* **93**: 73–89.
- Thornton, D. L. and Wheelock, D. C. (2014). Making sense of dissents: a history of FOMC dissents, *Federal Reserve Bank of St. Louis Review* **96**(3): 213–227.
- Tillmann, P. (2021). Financial markets and dissent in the ECB’s Governing Council, *European Economic Review* **139**: 103848.

Appendix: Tables

Table 2: Average effects

	Interquartile range	Standard deviation
β_{una}	0.012 [0.040]	0.001 [0.025]
β_{dis}	-0.006 [0.040]	-0.007 [0.024]
β_{unadis}	0.029 [0.040]	0.006 [0.025]
Female	-0.020 [0.030]	-0.012 [0.018]
Age (under 25 omitted)		
25-34	-0.037 [0.122]	0.011 [0.075]
35-44	-0.128 [0.118]	-0.059 [0.072]
45-54	-0.221 [0.116*]	-0.103 [0.071]
55-64	-0.232 [0.115*]	-0.105 [0.070]
65-74	-0.279 [0.115**]	-0.130 [0.070*]
over 74	-0.254 [0.120**]	-0.108 [0.073]
Income (under 1,500 omitted)		
1,500-3,000	0.139 [0.064**]	0.062 [0.039]
3,000-6,000	0.156 [0.065**]	0.070 [0.040*]
over 6,000	0.127 [0.071*]	0.043 [0.043]
City (under 5,000 omitted)		
5,000-20,000	0.098 [0.047**]	0.060 [0.029**]
20,000-100,000	0.131 [0.047***]	0.081 [0.029***]
100,000-500,000	0.178 [0.053***]	0.118 [0.033***]
over 500,000	0.109 [0.052**]	0.078 [0.031**]
Household size	0.018 [0.016]	0.014 [0.010]
Employed	-0.052 [0.046]	-0.023 [0.028]
Years of education	0.028 [0.009***]	0.011 [0.006**]
East Germany	-0.022 [0.039]	0.001 [0.024]
# obs.	2,338	2,338

Notes: We regress changes in the interquartile range and the standard deviation between the pre-treatment and treatment stages on a dummy for each treatment and demographic control variables. We drop responses that put 100% probability on the outer bins (more than 12% inflation/deflation) and responses of more than 100% deflation and inflation for the min and max expectations. The estimates are based on Huber robust regressions and also include a constant. ***, ** and * denote significance at 1%, 5% and 10% levels.

Table 3: Pre-expectations

	Interquartile range	Standard deviation
Female	−0.020 [0.013]	−0.014 [0.009]
Age (under 25 omitted)		
25-34	−0.021 [0.056]	−0.005 [0.039]
35-44	−0.108 [0.053**]	−0.063 [0.037*]
45-54	−0.118 [0.052**]	−0.069 [0.036*]
55-64	−0.109 [0.051**]	−0.065 [0.036*]
65-74	−0.209 [0.052***]	−0.132 [0.036***]
over 75	−0.210 [0.054***]	−0.130 [0.038***]
Income (under 1,500 omitted)		
1,500-3,000	−0.045 [0.028]	−0.027 [0.020]
3,000-6,000	−0.041 [0.029]	−0.028 [0.020]
over 6,000	−0.046 [0.032]	−0.035 [0.022]
City (under 5,000 omitted)		
5,000-20,000	0.017 [0.022]	0.011 [0.015]
20,000-100,000	−0.006 [0.022]	−0.003 [0.015]
100,000-500,000	0.022 [0.024]	0.012 [0.017]
over 500,000	0.024 [0.024]	0.015 [0.016]
Household size	0.017 [0.008**]	0.011 [0.005**]
Employed	0.002 [0.003]	0.001 [0.002]
Years of education	0.004 [0.004]	0.003 [0.003]
East Germany	0.045 [0.077**]	0.031 [0.012**]
# obs.	2,520	2,518

Notes: Dependent variables are the interquartile range and the standard deviation from the pre-treatment stage. Independent variables are respondent-specific factors, which are further explained in the text. We drop responses that put 100% probability on the outer bins (more than 12% inflation/deflation) and responses of more than 100% deflation and inflation for the min and max expectations. The estimates are based on Huber robust regressions and also include a constant. ***, ** and * denote significance at 1%, 5% and 10% levels.

Table 4: Descriptive statistics for household characteristics

	Percentiles				
	Mean	P25	P50	P75	Std. Dev.
Female	0.440	0	0	1	0.496
Age	56.858	46	58	70	14.840
Income					
under 1,500	0.068	0	0	0	0.252
1,500 - 3,000	0.298	0	0	1	0.457
3,000 - 6,000	0.379	0	0	1	0.485
over 6,000	0.176	0	0	0	0.381
City					
under 5,000	0.142	0	0	0	0.349
5,000 - 20,000	0.257	0	0	1	0.437
20,000 - 100,000	0.278	0	0	1	0.448
100,000 - 500,000	0.150	0	0	0	0.357
over 500,000	0.173	0	0	0	0.378
Household size	2.227	2	2	3	1.047
employed	0.536	0	1	1	0.499
Years of education	11.379	10	13	13	1.683
East Germany	0.165	0	0	0	0.372

Notes: This table presents summary statistics of the socio-economic variables we use in our analysis.

Table 5: Predictability of treatment status

	control	una	dis	unadis
Female	0.077 [0.056]	0.019 [0.057]	-0.036 [0.057]	-0.059 [0.057]
Age (under 25 omitted)				
25-34	-0.019 [0.243]	0.144 [0.245]	-0.050 [0.228]	-0.087 [0.243]
35-44	0.051 [0.230]	0.243 [0.234]	-0.322 [0.220]	0.038 [0.229]
45-54	0.054 [0.227]	0.093 [0.231]	-0.165 [0.215]	0.022 [0.227]
55-64	0.176 [0.224]	-0.004 [0.229]	-0.275 [0.213]	0.110 [0.224]
65-74	0.141 [0.225]	0.128 [0.230]	-0.307 [0.214]	0.048 [0.225]
over 74	0.187 [0.233]	0.156 [0.237]	-0.351 [0.223]	0.015 [0.233]
Income (under 1,500 omitted)				
1,500-3,000	-0.118 [0.118]	0.084 [0.122]	0.130 [0.122]	-0.087 [0.119]
3,000-6,000	-0.111 [0.120]	0.077 [0.124]	0.079 [0.124]	-0.034 [0.122]
over 6,000	-0.112 [0.133]	0.111 [0.138]	-0.054 [0.138]	0.071 [0.135]
City (under 5,000 omitted)				
5,000-20,000	0.092 [0.092]	-0.187 [0.090*]	0.035 [0.092]	0.067 [0.092]
20,000-100,000	0.075 [0.092]	-0.206 [0.090*]	0.062 [0.091]	0.076 [0.092]
100,000-500,000	0.140 [0.102]	-0.123 [0.101]	-0.029 [0.103]	0.018 [0.104]
over 500,000	0.053 [0.101]	-0.099 [0.098]	-0.047 [0.101]	0.100 [0.100]
Household Size	0.042 [0.032]	-0.017 [0.032]	-0.017 [0.032]	-0.010 [0.032]
Employed	0.036 [0.087]	-0.054 [0.089]	0.059 [0.088]	-0.043 [0.087]
Years of education	0.018 [0.018]	-0.010 [0.018]	0.014 [0.018]	-0.023 [0.018]
East Germany	0.090 [0.074]	-0.080 [0.076]	-0.033 [0.075]	0.023 [0.074]
# obs.	2,495	2,495	2,495	2,495

Notes: The table reports results for estimating a linear-probability regression with each treatment T as the dependent variable. The estimates also include a constant. ***, ** and * denote significance at the 1%, 5% and 10% level.

Table 6: The effect on the interquartile range in the baseline model

	I	II	III	IV
α	0.812 [0.033***]	0.831 [0.033***]	0.815 [0.033***]	0.828 [0.032***]
β_{una}	0.163 [0.049***]	0.193 [0.049***]	0.168 [0.050***]	0.195 [0.048***]
β_{dis}	0.216 [0.052***]	0.240 [0.052***]	0.234 [0.054***]	0.265 [0.050***]
β_{unadis}	-0.146 [0.050***]	-0.049 [0.050]	-0.082 [0.054]	-0.023 [0.050]
γ_{una}	-0.309 [0.042***]	-0.384 [0.041***]	-0.315 [0.048***]	-0.388 [0.041***]
γ_{dis}	-0.401 [0.049***]	-0.452 [0.048***]	-0.407 [0.049***]	-0.492 [0.048***]
γ_{unadis}	0.296 [0.045***]	0.141 [0.044***]	0.291 [0.052***]	0.104 [0.046**]
Female		-0.019 [0.028]		-0.017 [0.028]
Age (under 25 omitted)				
25-34		-0.040 [0.119]		-0.061 [0.117]
35-44		-0.137 [0.114]		-0.114 [0.113]
45-54		-0.220 [0.112**]		-0.228 [0.111**]
55-64		-0.248 [0.111**]		-0.249 [0.110**]
65-74		-0.313 [0.111***]		-0.323 [0.110***]
over 74		-0.290 [0.116**]		-0.298 [0.114***]
Income (under 1,500 omitted)				
1,500-3,000		0.078 [0.061]		0.067 [0.062]
3,000-6,000		0.105 [0.063*]		0.096 [0.062]
over 6,000		0.068 [0.069]		0.063 [0.068]
City (under 5,000 omitted)				
5,000-20,000		0.089 [0.046*]		0.090 [0.044**]
20,000-100,000		0.123 [0.046***]		0.124 [0.044***]
100,000-500,000		0.167 [0.052***]		0.165 [0.050***]
over 500,000		0.107 [0.050**]		0.102 [0.048**]
Household size		0.023 [0.016]		0.019 [0.015]
Employed		-0.047 [0.044]		-0.054 [0.043]
Years of education		0.030 [0.009***]		0.030 [0.009***]
East Germany		-0.001 [0.038]		0.001 [0.036]
below critical value (<i>una</i>)	42.04%	41.89%	39.94%	41.89%
below critical value (<i>dis</i>)	40.87%	40.87%	38.18%	40.87%
# obs.	2,520	2,338	2,388	2,218

Notes: Dependent variable is the posterior interquartile range. For specifications (I) and (II) we drop responses that put 100% probability on the outer bins (more than 12% inflation/deflation) and responses of more than 100% deflation and inflation for the min and max expectations. For (III) and (IV) we additionally exclude respondents whose change in the interquartile range is larger than the 95th percentile. The estimates are based on Huber robust regressions and also include a constant. ***, ** and * denote significance at 1%, 5% and 10% levels.

Table 7: The effect on the standard deviation in the baseline model

	I	II	III	IV
α	0.789 [0.033***]	0.814 [0.033***]	0.795 [0.033***]	0.825 [0.032***]
β_{una}	-0.002 [0.035]	0.135 [0.035***]	0.114 [0.034***]	0.142 [0.037***]
β_{dis}	0.122 [0.036***]	0.156 [0.036***]	0.150 [0.038***]	0.184 [0.035***]
β_{unadis}	-0.044 [0.035]	-0.013 [0.035]	-0.036 [0.038]	-0.001 [0.035]
γ_{una}	0.014 [0.042]	-0.381 [0.041***]	-0.312 [0.045***]	-0.403 [0.041***]
γ_{dis}	-0.320 [0.049***]	-0.428 [0.049***]	-0.390 [0.056***]	-0.487 [0.048***]
γ_{unadis}	0.127 [0.045***]	0.072 [0.044]	0.115 [0.058**]	0.049 [0.047]
Female		-0.015 [0.020]		-0.016 [0.020]
Age (under 25 omitted)				
25-34		-0.016 [0.083]		-0.025 [0.083]
35-44		-0.094 [0.080]		-0.087 [0.080]
45-54		-0.144 [0.078*]		-0.144 [0.079*]
55-64		-0.160 [0.078**]		-0.158 [0.078**]
65-74		-0.207 [0.077***]		-0.213 [0.078***]
over 75		-0.193 [0.081**]		-0.197 [0.081**]
Income (under 1,500 omitted)				
1,500-3,000		0.042 [0.043]		0.035 [0.044]
3,000-6,000		0.070 [0.044]		0.065 [0.044]
over 6,000		0.037 [0.048]		0.034 [0.038]
City (under 5,000 omitted)				
5,000-20,000		0.058 [0.032*]		0.060 [0.031*]
20,000-100,000		0.084 [0.032***]		0.085 [0.031***]
100,000-500,000		0.120 [0.036***]		0.119 [0.035***]
over 500,000		0.076 [0.035**]		0.072 [0.034**]
Household size		0.018 [0.011]		0.015 [0.011]
Employed		-0.032 [0.031]		-0.037 [0.031]
Years of education		0.020 [0.006***]		0.020 [0.006***]
East Germany		0.006 [0.026]		0.005 [0.026]
below critical value (<i>una</i>)	42.04%	41.70%	39.94%	41.70%
below critical value (<i>dis</i>)	40.87%	39.07%	38.18%	39.07%
# obs.	2,520	2,338	2,388	2,218

Notes: Dependent variable is the posterior standard deviation. For specifications (I) and (II) we drop responses that put 100% probability on the outer bins (more than 12% inflation/deflation) and responses of more than 100% deflation and inflation for the min and max expectations. For (III) and (IV) we additionally exclude respondents whose change in the interquartile range is larger than the 95th percentile. The estimates are based on Huber robust regressions and also include a constant. ***, ** and * denote significance at the 1%, 5% and 10% level.

Table 8: Coefficient restrictions for the baseline results

	Interquartile range				Standard deviation			
	I	II	III	IV	I	II	III	IV
$H_0 : una = dis \text{ for } x_{p25}^{pre}$	0.556	0.539	0.464	0.335	0.062	0.615	0.529	0.408
$H_0 : una = dis \text{ for } x_{p50}^{pre}$	0.961	0.875	0.939	0.800	0.652	0.839	0.910	0.734
$H_0 : una = dis \text{ for } x_{p75}^{pre}$	0.462	0.724	0.343	0.590	0.004	0.888	0.635	0.708

Notes: This table reports p -values for the hypothesis that the treatment effects of *una* and *dis* in the baseline specifications are equal. Formally, we test whether $\beta_{una} - \beta_{dis} + (\gamma_{una} - \gamma_{dis})x^{pre} = 0$. We test this for different percentiles of prior uncertainty.

Table 9: The effect on uncertainty (anchoring)

	Interquartile range		Standard deviation	
	I	II	III	IV
α	0.833 [0.033***]	0.831 [0.032***]	0.816 [0.033***]	0.828 [0.033***]
β_{una}	0.186 [0.053***]	0.187 [0.052***]	0.123 [0.037***]	0.135 [0.037***]
β_{dis}	0.270 [0.056***]	0.298 [0.054***]	0.187 [0.039***]	0.207 [0.038***]
β_{unadis}	-0.050 [0.055]	-0.022 [0.054]	-0.010 [0.038]	0.004 [0.038]
γ_{una}	-0.364 [0.042***]	-0.366 [0.041***]	-0.363 [0.042***]	-0.385 [0.042***]
γ_{dis}	-0.501 [0.049***]	-0.540 [0.049***]	-0.495 [0.049***]	-0.540 [0.049***]
γ_{unadis}	0.147 [0.045***]	0.106 [0.047**]	0.064 [0.045]	0.039 [0.048]
δ	0.041 [0.152]	0.59 [0.143]	0.011 [0.106]	0.027 [0.101]
α^{anc}	-0.085 [0.256]	-0.105 [0.241]	-0.073 [0.256]	-0.103 [0.244]
β_{una}^{anc}	0.070 [0.194]	0.061 [0.184]	0.066 [0.136]	0.051 [0.130]
β_{dis}^{anc}	-0.341 [0.218]	-0.349 [0.206*]	-0.225 [0.153]	-0.233 [0.146]
β_{unadis}^{anc}	-0.006 [0.216]	-0.037 [0.205]	-0.037 [0.151]	-0.055 [0.145]
γ_{una}^{anc}	-0.153 [0.306]	-0.146 [0.289]	-0.145 [0.307]	-0.120 [0.293]
γ_{dis}^{anc}	0.693 [0.367*]	0.683 [0.346**]	0.667 [0.369*]	0.670 [0.351*]
γ_{unadis}^{anc}	-0.004 [0.372]	0.059 [0.356]	0.120 [0.374]	0.170 [0.361]
Controls	✓	✓	✓	✓
Coefficient restrictions				
$\alpha + \alpha^{anc}$	0.749 [0.253***]	0.725 [0.238***]	0.743 [0.254***]	0.724 [0.242***]
$\beta_{una} + \beta_{una}^{anc}$	0.255 [0.186]	0.248 [0.176]	0.193 [0.131]	0.186 [0.124]
$\beta_{dis} + \beta_{dis}^{anc}$	-0.072 [0.211]	-0.051 [0.199]	-0.039 [0.148]	-0.026 [0.141]
$\beta_{unadis} + \beta_{unadis}^{anc}$	-0.056 [0.209]	-0.059 [0.198]	-0.047 [0.146]	-0.051 [0.141]
$\gamma_{una} + \gamma_{una}^{anc}$	-0.517 [0.303*]	-0.512 [0.286*]	-0.508 [0.305*]	-0.505 [0.290*]
$\gamma_{dis} + \gamma_{dis}^{anc}$	0.192 [0.364]	0.143 [0.343]	0.172 [0.365]	0.130 [0.347]
$\gamma_{unadis} + \gamma_{unadis}^{anc}$	0.143 [0.369]	0.166 [0.353]	0.185 [0.371]	0.209 [0.358]
# obs.	2,338	2,218	2,388	2,218

Notes: Dependent variables are the posterior interquartile range (I and II) and the standard deviation (III and IV). For all specifications we drop responses that put 100% probability on the outer bins (more than 12% inflation/deflation) and responses of more than 100% deflation and inflation for the min and max expectations. For (II) and (IV) we additionally exclude respondents whose change in the interquartile range is larger than the 95th percentile. The estimates are based on Huber robust regressions and also include a constant. ***, ** and * denote significance at 1%, 5% and 10% levels.

Table 10: The effect on the interquartile range with skewness from question CM004

	I	II	III	IV
α	0.628 [0.040***]	0.809 [0.030***]	0.788 [0.031***]	0.811 [0.030***]
β_{una}	0.079 [0.054]	0.192 [0.050***]	0.167 [0.050***]	0.192 [0.048***]
β_{dis}	0.163 [0.057***]	0.268 [0.049***]	0.264 [0.048***]	0.294 [0.048***]
β_{unadis}	-0.133 [0.056**]	0.008 [0.050]	-0.037 [0.050]	0.031 [0.049]
γ_{una}	-0.144 [0.051***]	-0.368 [0.039***]	-0.308 [0.045***]	-0.370 [0.039***]
γ_{dis}	-0.294 [0.059***]	-0.490 [0.040***]	-0.477 [0.041***]	-0.531 [0.040***]
γ_{unadis}	0.260 [0.057***]	0.034 [0.041]	0.094 [0.045**]	-0.001 [0.042]
Female		-0.022 [0.029]		-0.023 [0.028]
Age (under 25 omitted)				
25-34		-0.046 [0.119]		-0.069 [0.118]
35-44		-0.151 [0.114]		-0.146 [0.114]
45-54		-0.236 [0.113**]		-0.244 [0.112**]
55-64		-0.264 [0.112**]		-0.264 [0.111**]
65-74		-0.338 [0.112***]		-0.350 [0.111***]
over 75		-0.319 [0.117***]		-0.327 [0.115***]
Income (under 1,500 omitted)				
1,500-3,000		0.083 [0.062]		0.072 [0.062]
3,000-6,000		0.108 [0.063*]		0.100 [0.063]
over 6,000		0.073 [0.070]		0.068 [0.069]
City (under 5,000 omitted)				
5,000-20,000		0.100 [0.046**]		0.098 [0.045**]
20,000-100,000		0.127 [0.046***]		0.128 [0.045***]
100,000-500,000		0.169 [0.052***]		0.167 [0.051***]
over 500,000		0.114 [0.050**]		0.109 [0.049**]
Household size		0.024 [0.016]		0.020 [0.016]
Employed		-0.053 [0.045]		-0.054 [0.044]
Years of education		0.033 [0.009***]		0.033 [0.009***]
East Germany		0.002 [0.038]		0.004 [0.037]
below critical value (<i>una</i>)	n.a.	41.32%	41.32%	41.32%
below critical value (<i>dis</i>)	39.19%	39.19%	39.19%	39.19%
# obs.	2,520	2,388	2,388	2,2181

Notes: Dependent variable is the posterior interquartile range. We assume the prior distributions to have the same skewness as in question CM004. For specifications (I) and (II) we drop responses that put 100% probability on the outer bins (more than 12% inflation/deflation) and responses of more than 100% deflation and inflation for the min and max expectations. For (III) and (IV) we additionally exclude respondents whose change in the interquartile range is larger than the 95th percentile. The estimates are based on Huber robust regressions and also include a constant. ***, ** and * denote significance at the 1%, 5% and 10% level.

Table 11: The effect on the standard deviation with skewness from question CM004

	I	II	III	IV
α	0.787 [0.033***]	0.811 [0.032***]	0.793 [0.033***]	0.822 [0.032***]
β_{una}	-0.001 [0.035]	0.136 [0.035***]	0.114 [0.034***]	0.144 [0.035***]
β_{dis}	0.130 [0.037***]	0.167 [0.036***]	0.159 [0.038***]	0.189 [0.035***]
β_{unadis}	-0.040 [0.035]	-0.008 [0.035]	-0.031 [0.054]	0.003 [0.034]
γ_{una}	0.101 [0.042]	-0.379 [0.041***]	-0.311 [0.044***]	-0.399 [0.041***]
γ_{dis}	-0.341 [0.047***]	-0.444 [0.046***]	-0.417 [0.044***]	-0.499 [0.046***]
γ_{unadis}	0.114 [0.044**]	0.060 [0.044*]	0.102 [0.056*]	0.003 [0.035]
Female		-0.016 [0.020]		-0.017 [0.020]
Age (under 25 omitted)				
25-34		-0.016 [0.083]		-0.025 [0.083]
35-44		-0.094 [0.080]		-0.088 [0.080]
45-54		-0.145 [0.079*]		-0.146 [0.079*]
55-64		-0.162 [0.078**]		-0.159 [0.078**]
65-74		-0.210 [0.078***]		-0.214 [0.078***]
over 75		-0.196 [0.081***]		-0.198 [0.081***]
Income (under 1,500 omitted)				
1,500-3,000		0.041 [0.043]		0.034 [0.044]
3,000-6,000		0.069 [0.044]		0.063 [0.044]
over 6,000		0.036 [0.048]		0.033 [0.048]
City (under 5,000 omitted)				
5,000-20,000		0.059 [0.032*]		0.061 [0.031**]
20,000-100,000		0.084 [0.032***]		0.086 [0.031***]
100,000-500,000		0.120 [0.036***]		0.119 [0.036***]
over 500,000		0.078 [0.055**]		0.074 [0.034**]
Household size		0.018 [0.011]		0.015 [0.011]
Employed		-0.032 [0.031]		-0.037 [0.031]
Years of education		0.020 [0.006***]		0.020 [0.006***]
East Germany		0.006 [0.026]		0.006 [0.026]
below critical value (<i>una</i>)	41.70%	41.70%	41.70%	41.70%
below critical value (<i>dis</i>)	39.82%	39.87%	39.18%	39.82%
# obs.	2,520	2,518	2,388	2,391

Notes: Dependent variable is the posterior standard deviation. We assume the prior distributions to have the same skewness as in question CM004. For specifications (I) and (II) we drop responses that put 100% probability on the outer bins (more than 12% inflation/deflation) and responses of more than 100% deflation and inflation for the min and max expectations. For (III) and (IV) we additionally exclude respondents whose change in the interquartile range is larger than the 95th percentile. The estimates are based on Huber robust regressions and also include a constant. ***, ** and * denote significance at the 1%, 5% and 10% level.

Table 12: The effect on the interquartile range with skewness from the control group

	I	II	III	IV
α	0.812 [0.033***]	0.831 [0.032***]	0.815 [0.033***]	0.830 [0.032***]
β_{una}	0.163 [0.049***]	0.193 [0.050***]	0.168 [0.050***]	0.195 [0.048***]
β_{dis}	0.216 [0.052***]	0.240 [0.052***]	0.234 [0.054***]	0.265 [0.050***]
β_{unadis}	-0.146 [0.056***]	-0.049 [0.050]	-0.082 [0.054]	-0.023 [0.050]
γ_{una}	-0.309 [0.042***]	-0.383 [0.041***]	-0.315 [0.048***]	-0.388 [0.041***]
γ_{dis}	-0.401 [0.049***]	-0.452 [0.048***]	-0.434 [0.056***]	-0.492 [0.048***]
γ_{unadis}	0.296 [0.045***]	0.141 [0.044]	0.184 [0.058**]	0.104 [0.047**]
Female		-0.016 [0.029]		-0.017 [0.028]
Age (under 25 omitted)				
25-34		-0.040 [0.119]		-0.061 [0.117]
35-44		-0.137 [0.114]		-0.134 [0.113]
45-54		-0.220 [0.112**]		-0.228 [0.111**]
55-64		-0.248 [0.111**]		-0.249 [0.110**]
65-74		-0.313 [0.111***]		-0.326 [0.110***]
over 75		-0.287 [0.116***]		-0.297 [0.114***]
Income (under 1,500 omitted)				
1,500-3,000		0.078 [0.062]		0.067 [0.061]
3,000-6,000		0.105 [0.063*]		0.096 [0.062]
over 6,000		0.068 [0.069]		0.063 [0.068]
City (under 5,000 omitted)				
5,000-20,000		0.089 [0.046*]		0.089 [0.044**]
20,000-100,000		0.124 [0.046***]		0.124 [0.044***]
100,000-500,000		0.166 [0.052***]		0.166 [0.050***]
over 500,000		0.107 [0.050**]		0.102 [0.048**]
Household size		0.023 [0.016]		0.019 [0.015]
Employed		-0.047 [0.044]		-0.054 [0.043]
Years of education		0.030 [0.009***]		0.030 [0.009***]
East Germany		-0.001 [0.038]		0.004 [0.037]
below critical value (<i>una</i>)	41.29%	41.29%	41.29%	41.29%
below critical value (<i>dis</i>)	38.92%	38.92%	38.92%	38.92%
# obs.	2,520	2,518	2,388	2,391

Notes: Dependent variable is the posterior interquartile range. We assume the prior distributions to have the average skewness of our control group. For specifications (I) and (II) we drop responses that put 100% probability on the outer bins (more than 12% inflation/deflation) and responses of more than 100% deflation and inflation for the min and max expectations. For (III) and (IV) we additionally exclude respondents whose change in the interquartile range is larger than the 95th percentile. The estimates are based on Huber robust regressions and also include a constant. ***, ** and * denote significance at the 1%, 5% and 10% level.

Table 13: The effect on the standard deviation with skewness from the control group

	I	II	III	IV
α	0.789 [0.033***]	0.813 [0.033***]	0.795 [0.033***]	0.825 [0.033***]
β_{una}	-0.002 [0.035]	0.135 [0.035***]	0.114 [0.034***]	0.142 [0.034***]
β_{dis}	0.122 [0.036***]	0.159 [0.036***]	0.149 [0.038***]	0.184 [0.035***]
β_{unadis}	-0.044 [0.035]	-0.013 [0.035]	-0.036 [0.038]	-0.000 [0.035]
γ_{una}	0.139 [0.042]	-0.381 [0.041]	-0.312 [0.045***]	-0.403 [0.041***]
γ_{dis}	-0.320 [0.049***]	-0.428 [0.048***]	-0.391 [0.056***]	-0.487 [0.048***]
γ_{unadis}	0.127 [0.045**]	0.072 [0.044*]	0.115 [0.058*]	0.049 [0.047]
Female		-0.015 [0.020]		-0.016 [0.020]
Age (under 25 omitted)				
25-34		-0.016 [0.083]		-0.025 [0.083]
35-44		-0.093 [0.080]		-0.087 [0.080]
45-54		-0.144 [0.078*]		-0.145 [0.079*]
55-64		-0.160 [0.078**]		-0.157 [0.078**]
65-74		-0.207 [0.078***]		-0.213 [0.077***]
over 75		-0.193 [0.081**]		-0.196 [0.081**]
Income (under 1,500 omitted)				
1,500-3,000		0.042 [0.043]		0.035 [0.044]
3,000-6,000		0.070 [0.044]		0.065 [0.044]
over 6,000		0.037 [0.048]		0.034 [0.048]
City (under 5,000 omitted)				
5,000-20,000		0.058 [0.032*]		0.060 [0.031*]
20,000-100,000		0.084 [0.032***]		0.085 [0.031***]
100,000-500,000		0.120 [0.036***]		0.119 [0.035***]
over 500,000		0.076 [0.035**]		0.072 [0.034**]
Household size		0.018 [0.011]		0.015 [0.011]
Employed		-0.032 [0.031]		-0.037 [0.031]
Years of education		0.020 [0.006***]		0.197 [0.006***]
East Germany		0.006 [0.026]		0.005 [0.026]
below critical value (<i>una</i>)	41.70%	41.70%	41.70%	41.70%
below critical value (<i>dis</i>)	39.97%	39.97%	39.97%	39.97%
# obs.	2,520	2,518	2,388	2,391

Notes: Dependent variable is the posterior standard deviation. We assume the prior distributions to have the average skewness of our control group. For specifications (I) and (II) we drop responses that put 100% probability on the outer bins (more than 12% inflation/deflation) and responses of more than 100% deflation and inflation for the min and max expectations. For (III) and (IV) we additionally exclude respondents whose change in the interquartile range is larger than the 95th percentile. The estimates are based on Huber robust regressions and also include a constant. ***, ** and * denote significance at the 1%, 5% and 10% level.

Appendices

Appendix A Skewness

In this appendix, we explain how we calculate the skewness and the percentiles when we drop the assumption that the individual triangular distributions in the pre-treatment stage are symmetric. We measure the skewness of the distributions in question CM004 and of the control group as

$$Skew_i = \frac{p75_i + p25_i - (2 \times p50_i)}{p75_i - p25_i}, \quad (\text{A.1})$$

where $p25_i$, $p50_i$ and $p75_i$ are the 25th, 50th and 75th percentiles of the distribution of participant i , respectively.

In general, the symmetric or skewed triangular distribution is defined by its lower limit a , the upper limit b and the mode c . The cumulative distribution function can be expressed as

$$F(x) = \begin{cases} 0 & x \leq a \\ \frac{(x-a)^2}{(b-a)(c-a)} & a < x \leq c \\ 1 - \frac{(b-x)^2}{(b-a)(b-c)} & c < x < b \\ 1 & b \leq x. \end{cases} \quad (\text{A.2})$$

Based on the individual values for a_i , b_i and c_i , the 25th, 50th and 75th percentiles of the individual distributions are given by

$$F_i(0.25) \equiv p25_i = a_i + 0.5\sqrt{(b_i - a_i)(c_i - a_i)}, \quad (\text{A.3})$$

$$F_i(0.5) \equiv p50_i = \begin{cases} a_i + \sqrt{\frac{(b_i - a_i)(c_i - a_i)}{2}} & c \geq \frac{a_i + b_i}{2} \\ b_i - \sqrt{\frac{(b_i - a_i)(b_i - c_i)}{2}} & c \leq \frac{a_i + b_i}{2}, \end{cases} \quad (\text{A.4})$$

and

$$F_i(0.75) \equiv p75_i = b_i - 0.5\sqrt{(b_i - a_i)(b_i - c_i)}. \quad (\text{A.5})$$

Inserting (A.3), (A.4) and (A.5) in (A.1) yields

$$Skew_i = \frac{1}{N_i} \begin{cases} M_i - 2 \left(a_i + \sqrt{\frac{(b_i - a_i)(c_i - a_i)}{2}} \right) & c_i \geq \frac{a_i + b_i}{2} \\ M_i - 2 \left(b_i - \sqrt{\frac{(b_i - a_i)(b_i - c_i)}{2}} \right) & c_i \leq \frac{a_i + b_i}{2} \end{cases} \quad (\text{A.6})$$

with $M_i \equiv b - 0.5\sqrt{(b_i - a_i)(b_i - c_i)} + a_i + 0.5\sqrt{(b_i - a_i)(c_i - a_i)}$ and $N_i \equiv b_i - 0.5\sqrt{(b_i - a_i)(b_i - c_i)} - (a_i + 0.5\sqrt{(b_i - a_i)(c_i - a_i)})$.

According to (A.6), skewness now depends on a_i , b_i and c_i . Because solving for c_i is not straightforward, we minimize the squared deviation between $Skew_i$ and either the skewness from question CM004 or the average skewness of the control group, depending on the specification.

We solve

$$\min_{a_i < c_i < b_i} \left[S_i(a_i, b_i, c_i) - \tilde{S}_i \right]^2 \quad (\text{A.7})$$

or

$$\min_{a_i < c_i < b_i} \left[S_i(a_i, b_i, c_i) - \tilde{S}_{control} \right]^2, \quad (\text{A.8})$$

respectively, where $\tilde{S}_i$ denotes the skewness of respondent i 's distribution based on question CM004. $\tilde{S}_{control}$ is the average skewness of the control group. Knowing the optimal value for c_i , the asymmetric distribution is defined and we can calculate the standard deviation and the interquartile range as well as the mean.

Appendix B Additional tables

Tables (B.1) and (B.2) report the results using OLS and Huber regressions, respectively, of the outcome variables on treatment dummies with and without control variables.

Table (B.3) presents the results from a regression of the interquartile range from the treatment stage on a dummy for each treatment for different values of prior uncertainty and demographic variables.

Table B.1: The effects on post-treatment variables (OLS regressions)

	Interquartile range		Standard deviation	
β_{una}	-0.010 [0.148]	-0.074 [0.156]	-0.130 [0.096]	-0.050 [0.100]
β_{dis}	-0.008 [0.147]	-0.058 [0.153]	-0.003 [0.096]	-0.037 [0.099]
β_{unadis}	-0.001 [0.150]	0.010 [0.154]	-0.013 [0.096]	-0.005 [0.098]
Female		0.417 [0.128***]		0.274 [0.082***]
Age (under 25 omitted)				
25-34		0.299 [0.447]		0.146 [0.296]
35-44		-0.122 [0.427]		-0.134 [0.287]
45-54		-0.322 [0.401]		-0.246 [0.273]
55-64		-0.498 [0.412]		-0.379 [0.278]
65-74		-0.591 [0.396]		-0.445 [0.272]
over 74		-0.790 [0.401**]		-0.547 [0.278]
Income (under 1,500 omitted)				
1,500-3,000		-0.533 [0.358]		-0.354 [0.226]
3,000-6,000		-0.585 [0.349*]		-0.389 [0.221*]
over 6,000		-0.768 [0.350**]		-0.492 [0.224**]
City (under 5,000 omitted)				
5,000-20,000		-0.148 [0.199]		-0.075 [0.127]
20,000-100,000		-0.103 [0.209]		-0.053 [0.132]
100,000-500,000		0.175 [0.253]		0.132 [0.159]
over 500,000		-0.114 [0.224]		-0.053 [0.143]
Household size		-0.002 [0.059]		-0.001 [0.039]
Employed		-0.014 [0.208]		-0.000 [0.132]
Years of education		-0.198 [0.039***]		-0.130 [0.025***]
East Germany		0.218 [0.162]		0.147 [0.105]
# obs.	2,520	2,338	2,520	2,338

Notes: We regress the interquartile range and the standard deviation from the treatment stage on a dummy for each treatment and demographic control variables. We drop responses that put 100% probability on the outer bins (more than 12% inflation/deflation) and responses of more than 100% deflation and inflation for the min and max expectations. The estimates are based on OLS regressions and also include a constant. ***, ** and * denote significance at 1%, 5% and 10% levels.

Table B.2: The effects on post-treatment variables (Huber robust regressions)

	Interquartile range		Standard deviation	
β_{una}	-0.029 [0.043]	-0.021 [0.040]	-0.025 [0.030]	-0.017 [0.029]
β_{dis}	-0.013 [0.043]	-0.006 [0.040]	-0.008 [0.030]	-0.006 [0.029]
β_{unadis}	-0.028 [0.044]	-0.026 [0.040]	-0.020 [0.030]	-0.003 [0.029]
Female		-0.026 [0.030]		-0.019 [0.021]
Age (under 25 omitted)				
25-34		-0.074 [0.123]		-0.033 [0.087]
35-44		-0.243 [0.118**]		-0.171 [0.084**]
45-54		-0.335 [0.116***]		-0.222 [0.083***]
55-64		-0.378 [0.115***]		-0.255 [0.082***]
65-74		-0.511 [0.115***]		-0.344 [-0.082***]
over 74		-0.488 [0.120***]		-0.333 [0.085***]
Income (under 1,500 omitted)				
1,500-3,000		0.049 [0.064]		0.022 [0.046]
3,000-6,000		0.056 [0.065]		0.036 [0.046]
over 6,000		0.032 [0.071]		0.012 [0.051]
City (under 5,000 omitted)				
5,000-20,000		0.114 [0.047**]		0.079 [0.034**]
20,000-100,000		0.125 [0.047***]		0.087 [0.034**]
100,000-500,000		0.215 [0.053***]		0.150 [0.038***]
over 500,000		0.118 [0.052**]		0.085 [0.037**]
Household size		0.031 [0.017*]		0.024 [0.012**]
Employed		-0.047 [0.046]		-0.031 [0.033]
Years of education		0.030 [0.009***]		0.018 [0.007***]
East Germany		0.008 [0.039]		0.017 [0.028]
# obs.	2,520	2,338	2,520	2,338

Notes: We regress the interquartile range and the standard deviation from the treatment stage on a dummy for each treatment and demographic control variables. We drop responses that put 100% probability on the outer bins (more than 12% inflation/deflation) and responses of more than 100% deflation and inflation for the min and max expectations. The estimates are based on Huber robust regressions and also include a constant. ***, ** and * denote significance at 1%, 5% and 10% levels.

Table B.3: Heterogeneity Analysis

	Prior uncertainty		Gender		Age		Income		City		Household size		Employed		Years of education		East Germany	
	top 25%	bottom 25%	male	female	top 25%	bottom 25%	top 25%	bottom 25%	top 25%	bottom 25%	top 25%	bottom 25%	yes	no	top 25%	bottom 25%	yes	no
OLS regressions																		
β_{una}	0.324 [0.637]	0.033 [0.276]	-0.137 [0.197]	0.035 [0.258]	-0.534 [0.371]	-0.944 [0.314]	0.090 [0.203]	-0.128 [0.294]	-0.290 [0.239]	0.101 [0.200]	-0.249 [0.276]	-0.091 [0.193]	0.141 [0.226]	-0.291 [0.214]	-0.020 [0.193]	0.095 [0.668]	-0.483 [0.389]	0.021 [0.166]
β_{dis}	-0.074 [0.575]	0.248 [0.293]	-0.333 [0.290]	0.332 [0.290]	-0.618 [0.381]	-0.123 [0.271]	-0.093 [0.203]	-0.034 [0.311]	0.230 [0.385]	0.142 [0.237]	0.234 [0.342]	-0.035 [0.189]	-0.052 [0.181]	-0.028 [0.265]	-0.224 [0.147]	0.179 [0.700]	-0.682 [0.440]	0.037 [0.159]
β_{unadis}	0.634 [0.522]	0.212 [0.273]	-0.089 [0.193]	0.190 [0.267]	-0.113 [0.453]	0.010 [0.270]	-0.117 [0.173]	0.025 [0.310]	0.140 [0.313]	0.455 [0.242]	-0.109 [0.235]	0.010 [0.188]	-0.031 [0.181]	0.098 [0.256]	-0.183 [0.146]	0.017 [0.674]	-0.513 [0.394]	0.090 [0.165]
Huber regressions																		
β_{una}	-0.336 [0.293]	0.003 [0.046]	-0.040 [0.053]	0.009 [0.064]	-0.007 [0.066]	-0.050 [0.107]	-0.035 [0.085]	0.118 [0.181]	-0.141 [0.103]	0.079 [0.091]	0.022 [0.150]	-0.144 [0.091]	-0.016 [0.062]	-0.066 [0.049]	0.004 [0.060]	-0.020 [0.110]	-0.169 [0.121]	-0.002 [0.043]
β_{dis}	-0.500 [0.297 ⁺]	0.009 [0.047]	-0.008 [0.032]	0.000 [0.064]	0.071 [0.069]	-0.182 [0.105]	-0.101 [0.088]	-0.068 [0.171]	-0.166 [0.107]	-0.009 [0.094]	0.035 [0.141]	-0.162 [0.088]	-0.015 [0.061]	-0.009 [0.051]	-0.044 [0.059]	0.144 [0.110]	-0.224 [0.118 ⁺]	0.023 [0.043]
β_{unadis}	-0.195 [0.293]	-0.001 [0.046]	-0.015 [0.052]	0.014 [0.065]	0.012 [0.068]	0.032 [0.109]	-0.066 [0.084]	-0.162 [0.166]	-0.036 [0.104]	-0.045 [0.097]	0.123 [0.144]	-0.058 [0.088]	0.036 [0.062]	-0.048 [0.050]	0.011 [0.060]	0.036 [0.107]	-0.136 [0.119]	0.006 [0.043]

Notes: We regress the interquartile range from the treatment stage on a dummy for each treatment for different values of prior uncertainty and demographic variables. We drop responses that put 100% probability on the outer bins (more than 12% inflation/deflation) and responses of more than 100% deflation and inflation for the min and max expectations. The estimates also include a constant and control variables. ***, ** and * denote significance at 1%, 5% and 10% levels.

Appendix C Screenshots of the survey

Figure C.1: Pre-Treatment stage

P1908 | 2021_019 | Inflation expectations [min, max] | inflexp1_[a-b]

Respondent group: all

Range of valid values: -100.0 to +100.0

Info box "ECB Governing Council": The Governing Council of the ECB is the institution's main decision-making body. It consists of the six members of the ECB's Executive Board as well as the governors of the national central banks of the 19 euro area Member States.

Info box "Policy rate": The policy rate is the rate of interest at which commercial banks can borrow money from the central bank or deposit it there. In the euro area, the European Central Bank (ECB) is responsible for setting the policy rate, by means of which it can steer economic conditions, inflation and the exchange rate, amongst other things.

The following questions ask again about your inflation expectations given different scenarios.

Assume that the European Central Bank (ECB) is aiming for an annual inflation rate of 2% over the medium term. Please also assume that the inflation rate is 1% in 2021. The ECB Governing Council (i) decides to keep the policy rate (i) at 0%.

Question: In your opinion, how high will the inflation rate be at least over the next one to two years? And at most?

Note: If you assume there will be deflation, please enter a negative value. Values may have one decimal place.

a At least [percent]

b At most [percent]

Notes: The graph shows a screen shot of the question asked in the pre-treatment stage. The full survey is available at <https://www.bundesbank.de/resource/blob/763860/e6fb9fca831da8b44ed3c498b634cece/mL/questionnaire-19-data.png>.

Figure C.2: Treatment group $T = 1$ (control group)

P1909A | 2021_019 | Inflation expectations probabilistic – POST | inflexp_post_[a-j]

Respondent group: all

Input filter: drandom3 = 1

Info box "ECB Governing Council": The Governing Council of the ECB is the institution's main decision-making body. It consists of the six members of the ECB's Executive Board as well as the governors of the national central banks of the 19 euro area Member States.

Info box "Policy rate": The policy rate is the rate of interest at which commercial banks can borrow money from the central bank or deposit it there. In the euro area, the European Central Bank (ECB) is responsible for setting the policy rate, by means of which it can steer economic conditions, inflation and the exchange rate, amongst other things.

Continue to assume that the ECB is aiming for an annual inflation rate of 2% over the medium term. Please also assume that the inflation rate is 1% in 2021. The ECB Governing Council (i) decides to keep the policy rate (i) at 0%.

QUESTION: In your opinion, how likely is it that the rate of inflation will change as follows over the next one to two years?

Note: The aim of this question is to determine how likely you think it is that something specific will happen in the future. You can rate the likelihood on a scale from 0 to 100, with 0 meaning that an event is completely unlikely and 100 meaning that you are absolutely certain it will happen. Use values between the two extremes to moderate the strength of your opinion. Please note that your answers to the categories have to add up to 100.

- a The rate of deflation (opposite of inflation) will be 12% or higher.
- b The rate of deflation (opposite of inflation) will be between 8% and 12%.
- c The rate of deflation (opposite of inflation) will be between 4% and 8%.
- d The rate of deflation (opposite of inflation) will be between 2% and 4%.
- e The rate of deflation (opposite of inflation) will be between 0% and 2%.
- f The rate of inflation will be between 0% and 2%.
- g The rate of inflation will be between 2% and 4%.
- h The rate of inflation will be between 4% and 8%.
- i The rate of inflation will be between 8% and 12%.
- j The rate of inflation will be 12% or higher.

Notes: The graph shows a screen shot of the question asked from treatment group $T = 1$ (control group= in the treatment stage. The full survey is available at <https://www.bundesbank.de/resource/blob/763860/e6fb9fca831da8b44ed3c498b634cece/mL/questionnaire-19-data.png>.

Figure C.3: Treatment group $T = 2$ (unanimity)

P1909B | 2021_019 | Inflation expectations probabilistic – POST | inflexp_post_[a-j]

Respondent group: all

Input filter: drandom3 = 2

Info box "ECB Governing Council": The Governing Council of the ECB is the institution's main decision-making body. It consists of the six members of the ECB's Executive Board as well as the governors of the national central banks of the 19 euro area Member States.

Info box "Policy rate": The policy rate is the rate of interest at which commercial banks can borrow money from the central bank or deposit it there. In the euro area, the European Central Bank (ECB) is responsible for setting the policy rate, by means of which it can steer economic conditions, inflation and the exchange rate, amongst other things.

Continue to assume that the ECB is aiming for an annual inflation rate of 2% over the medium term. Please also assume that the inflation rate is 1% in 2021. The ECB Governing Council (i) decides to keep the policy rate (i) at 0%. The ECB President informs the media that this was a unanimous decision.

QUESTION: In your opinion, how likely is it that the rate of inflation will change as follows over the next one to two years?

Note: The aim of this question is to determine how likely you think it is that something specific will happen in the future. You can rate the likelihood on a scale from 0 to 100, with 0 meaning that an event is completely unlikely and 100 meaning that you are absolutely certain it will happen. Use values between the two extremes to moderate the strength of your opinion. Please note that your answers to the categories have to add up to 100.

- a The rate of deflation (opposite of inflation) will be 12% or higher.
- b The rate of deflation (opposite of inflation) will be between 8% and 12%.
- c The rate of deflation (opposite of inflation) will be between 4% and 8%.
- d The rate of deflation (opposite of inflation) will be between 2% and 4%.
- e The rate of deflation (opposite of inflation) will be between 0% and 2%.
- f The rate of inflation will be between 0% and 2%.
- g The rate of inflation will be between 2% and 4%.
- h The rate of inflation will be between 4% and 8%.
- i The rate of inflation will be between 8% and 12%.
- j The rate of inflation will be 12% or higher.

Notes: The graph shows a screen shot of the question asked from treatment group $T = 2$ (unanimity) in the treatment stage. The full survey is available at <https://www.bundesbank.de/resource/blob/763860/e6fb9fca831da8b44ed3c498b634cece/mL/questionnaire-19-data.png>.

Figure C.4: Treatment group $T = 3$ (dissent)

P1909C | 2021_019 | Inflation expectations probabilistic – POST | inflexp_post_[a-j]

Respondent group: all

Input filter: drandom3 = 3

Info box "ECB Governing Council": The Governing Council of the ECB is the institution's main decision-making body. It consists of the six members of the ECB's Executive Board as well as the governors of the national central banks of the 19 euro area Member States.

Info box "Policy rate": The policy rate is the rate of interest at which commercial banks can borrow money from the central bank or deposit it there. In the euro area, the European Central Bank (ECB) is responsible for setting the policy rate, by means of which it can steer economic conditions, inflation and the exchange rate, amongst other things.

Continue to assume that the ECB is aiming for an annual inflation rate of 2% over the medium term. Please also assume that the inflation rate is 1% in 2021. The ECB Governing Council (i) decides to keep the policy rate (i) at 0%. The ECB President informs the media that this was a majority decision, i.e. there were dissenting votes.

QUESTION: In your opinion, how likely is it that the rate of inflation will change as follows over the next one to two years?

Note: The aim of this question is to determine how likely you think it is that something specific will happen in the future. You can rate the likelihood on a scale from 0 to 100, with 0 meaning that an event is completely unlikely and 100 meaning that you are absolutely certain it will happen. Use values between the two extremes to moderate the strength of your opinion. Please note that your answers to the categories have to add up to 100.

- a The rate of deflation (opposite of inflation) will be 12% or higher.
- b The rate of deflation (opposite of inflation) will be between 8% and 12%.
- c The rate of deflation (opposite of inflation) will be between 4% and 8%.
- d The rate of deflation (opposite of inflation) will be between 2% and 4%.
- e The rate of deflation (opposite of inflation) will be between 0% and 2%.
- f The rate of inflation will be between 0% and 2%.
- g The rate of inflation will be between 2% and 4%.
- h The rate of inflation will be between 4% and 8%.
- i The rate of inflation will be between 8% and 12%.
- j The rate of inflation will be 12% or higher.

Notes: The graph shows a screen shot of the question asked from treatment group $T = 3$ (dissent) in the treatment stage. The full survey is available at <https://www.bundesbank.de/resource/blob/763860/e6fb9fca831da8b44ed3c498b634cece/mL/questionnaire-19-data.png>.

Figure C.5: Treatment group $T = 4$ (unanimity despite different opinions)

P1909D | 2021_019 | Inflation expectations probabilistic – POST | inflexp_post_[a-i]

Respondent group: all

Input filter: drandom3 = 4

Info box "ECB Governing Council": The Governing Council of the ECB is the institution's main decision-making body. It consists of the six members of the ECB's Executive Board as well as the governors of the national central banks of the 19 euro area Member States.

Info box "Policy rate": The policy rate is the rate of interest at which commercial banks can borrow money from the central bank or deposit it there. In the euro area, the European Central Bank (ECB) is responsible for setting the policy rate, by means of which it can steer economic conditions, inflation and the exchange rate, amongst other things.

Continue to assume that the ECB is aiming for an annual inflation rate of 2% over the medium term. Please also assume that the inflation rate is 1% in 2021. The ECB Governing Council (i) decides to keep the policy rate (i) at 0%. The ECB President informs the media that this was a unanimous decision despite opinions differing.

QUESTION: In your opinion, how likely is it that the rate of inflation will change as follows over the next one to two years?

Note: The aim of this question is to determine how likely you think it is that something specific will happen in the future. You can rate the likelihood on a scale from 0 to 100, with 0 meaning that an event is completely unlikely and 100 meaning that you are absolutely certain it will happen. Use values between the two extremes to moderate the strength of your opinion. Please note that your answers to the categories have to add up to 100.

- a The rate of deflation (opposite of inflation) will be 12% or higher.
- b The rate of deflation (opposite of inflation) will be between 8% and 12%.
- c The rate of deflation (opposite of inflation) will be between 4% and 8%.
- d The rate of deflation (opposite of inflation) will be between 2% and 4%.
- e The rate of deflation (opposite of inflation) will be between 0% and 2%.
- f The rate of inflation will be between 0% and 2%.
- g The rate of inflation will be between 2% and 4%.
- h The rate of inflation will be between 4% and 8%.
- i The rate of inflation will be between 8% and 12%.
- j The rate of inflation will be 12% or higher.

Notes: The graph shows a screen shot of the question asked from treatment group $T = 4$ (unanimity despite different opinions) in the treatment stage. The full survey is available at <https://www.bundesbank.de/resource/blob/763860/e6fb9fca831da8b44ed3c498b634cece/mL/questionnaire-19-data.png>.

Essay V

Inflation Expectations and a Conservative Central Banker: A Natural Experiment

This paper was presented at the following workshops and international conferences:

- i. VfS Ausschuss für Geldtheorie und Geldpolitik (Frankfurt)

Date: March 2026

Presenter: Peter Tillmann

- ii. 2nd Kiel-CEPR Monetary Policy Conference (Berlin)

Date: May 2026

Presenter: Peter Tillmann

Acknowledgement

We thank Mamadou Lamine Barry for help with the Nexis data. Winnie Coleman kindly shared her data. We are grateful to Lena Dräger, Sandra Eickmeier, David Finck, Sören Karau (discussant), Giang Nghiem, Davide Romelli and participants at the VfS Ausschuss für Geldtheorie und Geldpolitik (Frankfurt, 2026) and the 2nd Kiel-CEPR Monetary Policy Conference (Berlin, 2026) for helpful discussions. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Deutsche Bundesbank or the Eurosystem.

Inflation Expectations and a Conservative Central Banker: A Natural Experiment

Moritz Grebe*

Peter Tillmann[†]

Abstract

This paper studies the response of inflation expectations to an exogenous shift in central bank preferences. We use the unexpected announcement of the resignation of Jens Weidmann, the president of the Deutsche Bundesbank, in 2021 as a natural experiment constituting a rare case of an exogenous shift in the composition of the ECB's Governing Council. As a member of the Governing Council, he was a known policy hawk and a vocal critic of the ECB's asset purchases. Our evidence from survey data suggests that the news about the resignation causes a strong increase in individual inflation uncertainty of German households and a significant fall in the level of trust in the ECB. The effect on the mean inflation expectations remains inconclusive. Thus, the shift in policy preferences exerts a strong effect on second moments, while the first moment effect remains weak.

Keywords: Monetary policy delegation, central bank governor, household survey, inflation uncertainty, trust

JEL classification: E52, E43, E32

*Deutsche Bundesbank, Email: moritz.grebe@bundesbank.de

[†]Justus Liebig University Giessen, Email: peter.tillmann@wirtschaft.uni-giessen.de

1 Introduction

A substantial literature argues that a conservative central banker, a policymaker with a large relative weight on inflation stabilization in her loss function, reduces average inflation. Appointing a known inflation hawk to lead an independent central bank, the argument goes, contains inflation expectations. This addresses the inflation bias of discretionary monetary policy and raises social welfare (Kydland and Prescott, 1977; Barro and Gordon, 1983; Rogoff, 1985). Appointing a conservative central banker even remains socially beneficial in New-Keynesian models in which the economy does not suffer from an inflation bias (Clarida et al., 1999). In these models, a central banker with a high relative weight on stabilizing inflation in his loss function reduces the variance of inflation. In this case, a conservative central banker attenuates the stabilization bias of discretionary monetary policy and thus raises welfare. In either case, appointing a weight-conservative central banker has an impact on inflation expectations.

The benefits of a conservative central banker are generally shared by researchers and policymakers (Constâncio, 2011). However, the empirical evidence on the causal effect of a conservative policymaker on the expectations of the public remains slim. The reason for this is that exogenous variation in the preferences of central bankers is rare, while such variation is exactly what is needed to identify a causal effect on inflation expectations. The empirical literature reviewed below uses cross-country data to establish a link between central bank conservatism and/or central bank independence and the resulting inflation rate, though much of the evidence does not allow for a causal interpretation.

In this paper, we estimate the causal effect of an exogenous change in the degree of conservatism of the central banker on the inflation expectations of German households. We use a natural experiment: the announcement of the resignation of Jens Weidmann, president of the Deutsche Bundesbank and a member of the Governing Council of the European Central Bank (ECB), in October 2021. As a known policy hawk and a frequent critic of ECB policy, his resignation for personal reasons was perceived as a likely shift of the composition of the Governing Council towards less inflation aversion. Importantly, the resignation was completely unexpected as Jens Weidmann was re-appointed for his second term only two years earlier. Below, we show that there is indeed no sign the public expected the resignation. In addition, this event is informative because the new German government, chaired by the Social Democrats after 16 years of coalition governments chaired by the Christian Democrats, did not immediately announce a successor, thus fueling expectations of a dovish tilt. The newly formed government announced Joachim Nagel as the successor of Jens

Weidmann on December 20, 2021. The resignation does not coincide with a shift in the ECB's policy framework, i.e. its mandate, its inflation target or its independence. The news of the resignation falls into the October wave of the Bundesbank Online Panel - Households, a large and well-established online survey among German households. We use the answers to this survey, in particular to the questions on households' inflation expectations and their level of trust in the ECB, to estimate the causal effect of the news. Comparing survey responses submitted before the news with the responses after the news allows us to identify the effect of the resignation. Thus, the "treatment" in this case is the timing of the survey response.¹

Our results show no clear effect of the news on expectations about the level of inflation. This applies to the mean of the individual forecast distributions, the individual point estimates and the qualitative expectations of future inflation. This result is surprising given the fact that the literature emphasizes the transmission channel through the moderation of inflation expectations. However, we find strong and highly significant effects on the second moments of inflation expectations, i.e. on inflation uncertainty measured by the width of the individual forecast distributions. Both the interquartile range and the standard deviation of individual forecasts are significantly higher for respondents who submit their answers after the announcement of the resignation. In addition, we show that the level of trust in the ECB to be able to deliver price stability is significantly lower for households who complete the survey after the resignation. All results remain robust with respect to different estimation methods and different (sub)sets of survey respondents. Our results emphasize the role of inflation uncertainty and trust as key variables sensitive to a perceived shift in policy preferences. This is in contrast to the literature which focuses on first moments. Our interpretation of the results is that inflation expectations remain well anchored, even in the short run, while the uncertainty about the path of inflation towards the anchor increases.

A large literature has established a firm link between central bank independence and low inflation (Romelli, 2022) in the cross section of countries. In addition, a number of papers study the cross-country impact of changes in central bank leadership on economic outcomes (Göhlmann and Vaubel, 2007; Dreher et al., 2008, 2010; Moser and Dreher, 2010; Kuttner and Posen, 2010; Carvalho et al., 2017; Ioannidou et al., 2022). Both strands of the literature face two challenges. First, the identification of exogenous variation in either the institutional design of central banks or the governorship of central banks is difficult. This hampers the estimation of causal effects. Second, it is difficult to distinguish the impact of cross-country differences in

¹This design is similar to Dräger, Gründler and Potrafke (2025) and De Fiore et al. (2025).

the leadership of central banks on inflation from cross-country differences in central bank independence. This is why Berger and Woitek (2005) advocate single-country evidence on the impact of central bank preferences on economic outcomes.² Our paper addresses both points: we provide causal evidence and can isolate the change in governorship from other potential institutional changes.

The paper also draws on the literature on the delegation and the institutional design of monetary policy (Backus and Driffill, 1985; Vickers, 1986; Walsh, 1995), in particular on the role of uncertainty about the preferences of the central banker on the economic outcomes of the delegation of monetary policy (Beetsma and Jensen, 1998; Alesina and Gatti, 1995; Sibert, 2002; Hahn, 2016; Schaumburg and Tambalotti, 2007). While we can quantify the impact of an anticipated change in the composition of the Governing Council with a likely tilt towards less monetary conservatism, we cannot separate the effect from the remaining uncertainty about the preferences of the next Bundesbank president.

Our paper is also related to Neuenkirch and Tillmann (2016), who quantify the impact of “superstar” central bankers (e.g. “Super Mario” Draghi). They adopt a matching approach in order to construct a counterfactual and find that good central bankers do indeed face a more favorable output-inflation trade-off.³ While their paper analyzes the impact of acclaimed central bank governors, our paper concentrates on a shift in the preferences of the central bank committee.

Finally, this paper is inspired by the recent interest in household expectations from survey data and their reactions to monetary policy decisions, central bank communication or other events (Coibion et al., 2020, 2024, 2022; Hoffmann et al., 2022; Dräger, Lamla and Pfajfar, 2025; Dräger, Gründler and Potrafke, 2025). This literature also shows that an increase in uncertainty about future inflation, which is our main result, has detrimental effects on household spending (Fischer et al., 2024; Kostyshyna and Petersen, 2024).

The remainder of this paper is structured as follows. Section 2 sheds light on the circumstances around the resignation. Section 3 introduces the household survey and the identification. The empirical framework and the main results are documented in section 4. Section 5 uses an illustrative model to rationalize our findings. Section 6 concludes.

²Berger and Woitek (2005) measure changes in the preferences of the rate-setting board of the Bundesbank before the creation of the euro area on economic outcomes using voting records of policymakers.

³See also Mehrling et al. (2007) for a discussion of the personalities of central bankers.

2 A natural experiment: Jens Weidmann's resignation

On October 20, 2021, Jens Weidmann, the President of the Deutsche Bundesbank, resigned unexpectedly. He resigned for personal reasons. The Bundesbank issued a press release at 10:45 am in both German and English and also posted the news on Twitter.⁴

The Bundesbank is the central bank of the largest economy of the euro area. Its governor is a member of the ECB's Governing Council and arguably one of the very influential voices on the formulation of monetary policy in the euro area.⁵ As a known hawk on monetary policy, or a "conservative central banker" in the sense of Rogoff (1985), Jens Weidmann embodied the stability orientation of the German public. His controversies with Mario Draghi, the ECB president, about unconventional monetary policy measures and policies directed towards containing sovereign yield spreads of individual euro area countries were widely appreciated in the German media, making Jens Weidmann a very popular and well-known Bundesbank president (Tillmann and Walter, 2019).⁶ When Christine Lagarde succeeded Mario Draghi, the relationship between the president of the Bundesbank and the ECB leadership improved. The Financial Times comments on his resignation as follows:

"The president of the Bundesbank has been one of the most vocal critics of the ultra-loose monetary policy pursued by the European Central Bank, where he fought an often lonely battle against its bond buying and negative interest rate policies."

In addition, the Financial Times interprets the resignation as a shift towards a less hawkish policy stance:

"As the longest-serving member of the ECB's governing council, and one of its most "hawkish" voices who upheld the Bundesbank's traditionally orthodox views on monetary policy that draw on Germany's experience of hyperinflation in the 1920s, his departure opens the door to a potential further shift in favour of the ECB's "dovish" majority."

⁴This is the Press Release. The Twitter posts are shown in the appendix.

⁵The voting right rotates. The governors of the central banks of Germany, France, Italy, Spain and the Netherlands, the five largest economies of the euro area, share four voting rights on the Governing Council. The remaining governors share 11 voting rights.

⁶Grebe and Tillmann (2025) show the consequences of dissent in the Governing Council on the inflation expectations of German households.

The resignation was not just covered by specialized media such as the Financial Times, but was also widely exploited by the tabloid press. The appendix provides an example from bild.de, the largest German tabloid newspaper, claiming: “Disaster for our money!: Germany’s guardian of the stable euro gives up”.


We investigate the impact of the resignation on the inflation expectations of German households. The resignation is a very informative natural experiment. This is because the news about Jens Weidmann’s departure corresponds to an exogenous change in the monetary policy preferences of the Governing Council. It is a rare case of a shift in the composition of the council not due to a term in office coming to its end or an expected resignation after months of public debates.

Jens Weidmann was appointed President of the Deutsche Bundesbank in 2011 by Chancellor Angela Merkel and re-appointed for another eight-year tenure in 2019. Three points make his resignation particularly useful to study the effects of exogenous shifts in monetary policy preferences. First, the resignation was completely unexpected. As we will show below, the German public was taken by surprise on October 20. Second, the resignation, to become effective on December 31, was announced three weeks after the general election. It fell into the time of the negotiations about forming a new coalition government in Germany. Arguably, Jens Weidmann timed his resignation in order to allow the coalition parties to include his succession in their negotiations. Hence, there was no immediate successor announced because not even the new government was sworn in. The newly formed government announced the successor of Jens Weidmann on December 20, 2021. The fact that the new coalition government was expected to be led by the Social Democrats suggests that the likely successor of Jens Weidmann will be less hawkish. Third, the resignation does not coincide with a shift in the ECB’s policy framework, i.e. its mandate, its inflation target or its independence, which could undermine the identification.

At the time of the announced resignation, the latest available inflation rate for Germany was 4.1% (year-on-year change in the Consumer Price Index in September 2021). The latest available inflation rate in the euro area was 3.4% (year-on-year change in the Harmonized Index of Consumer Prices in September 2021).

We now turn to the response of financial markets to the news of the resignation. Figure (1) shows the intraday yields on 10-year German bunds, 10-year Italian government bonds and the spread between them in a narrow window from 10:30 am to 11:00 am on October 20. Both yields should reflect the future path of short-term interest rates. In addition, Italian yields incorporate a risk premium, which should be sensitive to changes in market perceptions of future ECB measures in support

Figure 1: Bond yields and spreads around the announcement of the resignation



Notes: The dashed vertical line highlights the time of the press release at 10:45 am CET.

of highly indebted member states. The announced resignation of Jens Weidmann, a vocal critic of these measures, should shift the balance in the Governing Council towards more support. German bond yields exhibit no discernible reaction. Yields on Italian bonds fall, though the magnitude of the reaction is tiny. As a consequence, the yield spread of Italy against Germany falls by less than one basis point. Overall, these responses suggest that the impact of the news on first moments, i.e. mean inflation and interest rate expectations, is minuscule. However, we will show below that the news triggered a large change in second moments, i.e. the uncertainty about future inflation.

3 Data and identification


3.1 The BOP-HH survey

We use data from wave 22 of the BOP-HH, which was conducted in October 2021. The BOP-HH is a well-established online panel of German households.⁷ A key property of wave 22 is that October 20, the day of the announced resignation, falls into the survey period. As explained below, this allows us to compare responses before and after the news of the resignation. The time stamps on each response provide us with the day the response was submitted, but not the exact time. Hence, some of the responses submitted on October 20 might have been completed before the news about the resignation broke on 10:45 am. Figure (2) shows the number of responses over the days of the survey.

More than 3,000 respondents completed the survey on October 18 and 19, i.e. before

⁷Details about the survey are available at <https://www.bundesbank.de/en/bundesbank/research/survey-on-consumer-expectations>. The survey is conducted in German. Here, we use the translation available on the Bundesbank website.

Figure 2: Number of survey responses per day



Notes: The figure shows the number of survey responses per day in October 2021 of wave 22 of the BOP-HH. The dashed vertical line highlights October 20, the day Jens Weidmann announced his resignation.

the resignation, and around 1,000 respondents completed the survey on October 20 or after. In the empirical analysis below we only include uninterrupted responses. The sample period includes the release of the flash CPI inflation figure (4.5%) for October and the regular meeting of the ECB's Governing Council, both on October 28, 2021. Importantly, the week between October 21 and October 28 covers the quiet period before the Governing Council meeting during which members are not allowed to make public comments which could influence the expectations about monetary policy. On October 28, the ECB decided to leave interest rates unchanged, which was in line with market expectations. As a result, market interest rates barely reacted.⁸ Hence, the post-resignation subsample is not contaminated by other important news about monetary policy.

In the following, we explain the specific questions from wave 22, which are particularly relevant for our analysis. We also explain how we transform the responses in order to make them operational for our empirical estimation.⁹ Table (1) reports the descriptive statistics of the six dependent variables.

Probabilistic expectations We use the responses to the following question: "In your opinion, how likely is it that the rate of inflation will change as follows over the next twelve months?" The respondents are given bins over which they have

⁸According to the EA-MPD of Altavilla et al. (2019), interest rates on German two- and ten-year bonds increase by 0.45 and 1.25 basis points, respectively, during the event window.

⁹The English version of the original survey questions is available at Deutsche Bundesbank.

to allocate 100 points in total. The bins range from deflation of 12% or more to inflation of 12% or more. The width of the outer bins is four percentage points, while the width of the inner bins is two percentage points. The number of bins and the varying width of bins is identical to the Survey of Consumer Expectations run by the Federal Reserve Bank of New York. We show the exact formulation of the questions in the appendix. In order to analyze the responses, we need to fit a probability distribution to each individual answer. We follow the procedure of Engelberg et al. (2009) and Manski (2018), which we explain in the appendix. This provides us with the individual mean inflation expectation as well as two measures of individual inflation uncertainty: the interquartile range and the standard deviation of the distribution. The interquartile range is more robust with respect to outliers compared to the standard deviation of individual forecasts.

Inflation point expectation Respondents receive factual information on the definition of inflation and deflation. They are asked: “Do you think inflation or deflation is more likely over the next twelve months?”. If they choose inflation (deflation), they are then asked: “What do you think the rate of inflation (deflation) will roughly be over the next twelve months?” Respondents can enter a value with one decimal place. We calculate the point expectation as

$$E_i(\pi) = \begin{cases} E_i(\pi^{infl}) & \text{if } \mathbb{I}_i^{infl} = 1 \\ -E_i(\pi^{defl}) & \text{if } \mathbb{I}_i^{defl} = 1, \end{cases} \quad (1)$$

where $\mathbb{I}_i^{infl}$ ($\mathbb{I}_i^{defl}$) is equal to one if respondent i chooses inflation (deflation) and zero otherwise and $E_i(\pi^{infl})$ and $E_i(\pi^{defl})$ are the expected rates of inflation or deflation.

Qualitative inflation expectations Respondents are asked: “What developments do you expect in the [the inflation rate] over the next twelve months?”. They can choose between “will decrease significantly”, “will decrease slightly”, “will stay roughly the same”, “will increase slightly” and “will increase significantly”. We code the “increase (decrease) strongly” answers as a value of 2 (-2), the “increase (decrease) slightly” answers as a value of 1 (-1) and the “stay the same” answer as a value of 0.

Trust The survey participants are asked to provide their level of trust in the ECB to be able to deliver price stability from a value of 0 (do not trust at all) to a value of 10 (trust entirely). They could also choose the answer “I am unfamiliar with the ECB”. In our regressions, we use the values from 0 to 10. Survey rounds before or after wave 22 in October 2021 do not include a question about trust. We interpret

trust following Aikman et al. (2025). These authors argue that trust reflects the expectation that the “central bank will act in the public interest and possess the technical competence to achieve its objectives.” It is a broader notion than credibility and includes also informal norms and social capital.¹⁰

Table 1: Descriptive statistics

	# obs.	mean	median	std.	min.	max.
inflation point expectation	3295	5.01	4.00	7.21	-70.00	100.00
qualitative inflation expectation	3931	1.21	1.00	0.90	-2.00	2.00
mean inflation expectation	3536	3.80	3.27	3.38	-25.00	25.00
interquartile range	3536	2.42	1.67	2.66	0.04	29.82
standard deviation	3536	1.66	1.20	1.72	0.03	18.37
trust	3887	4.41	5.00	2.58	0.00	10.00

Notes: The table shows descriptive statistics for the dependent variables. We include only the uninterrupted survey responses.

Socio-demographic information Survey participants also provide detailed information about their socio-demographic status, which we use as control variables in our estimations below. We include a dummy that is one if the respondent identifies as male and zero otherwise. Another dummy variable is one if the respondent is employed and zero otherwise. The household size and the age of the respondent are taken as such without further transformations. We transform the information on school education and professional education into one variable that reflects the overall number of years of schooling. Separate dummy variables indicate whether the household income is below 2,000 EUR, between 2,000 and 5,000 Euro and between 5,000 and 8,000 Euro, thus leaving households with an income of more than 8,000 Euro as the omitted category. We include information about the residency of households. The separate dummy variables indicate whether the household lives in the East, the South and the North of Germany leaving the West as the omitted category. Two

¹⁰See Eickmeier and Petersen (2024) for an in-depth analysis of the determinants and the consequences of trust in the ECB. Ehrmann (2024) provides a survey of the importance of trust for monetary policy. Christelis et al. (2020) show that trust in the central bank helps anchor inflation expectations.

additional dummy variable indicate whether the household lives in a medium city or a big city leaving small cities as the omitted category. Finally, a dummy variable is one if the respondent identifies as single and zero otherwise.

3.2 Identification


We estimate the causal effect of the news about the resignation on inflation expectations and trust in the ECB. Specifically, we compare survey responses completed before October 20 and responses completed on October 20 or after. Our identifying assumption is that (i) the resignation is unexpected and (ii) absent the treatment status, individuals in the treatment and control group are statistically indistinguishable.

Let us discuss each assumption in turn. Figure (3) shows the daily number of visits of the Wikipedia page on Jens Weidmann (in German) as well as the weekly Google search volume on "Jens Weidmann" in Germany. If the resignation was expected, we should see fluctuations in both series before the news. Both number are indistinguishable from zero before October 20 and explode as the news breaks on President Weidmann's announced resignation. In panel (c), the figure shows the number of newspaper and magazine articles in German containing the word "Jens Weidmann" for each day of the year 2021. We include all articles we can find on the Nexis database. We can clearly see a massive increase in reporting about Jens Weidmann on October 20 and the subsequent days. There are zero articles on October 18 and 19, but 71 and 73 on October 20 and 21, respectively. Over the full year 2021, we find 1161 relevant articles, 630 of which appeared after October 20. None of the articles published before October 20 speculates about a resignation of the Bundesbank president. All three panels of the figure support the notion that the announced resignation was indeed unexpected. There were no rumours about a possible resignation in the days and weeks before October 20.

As regards the second part of the identifying assumption, we need to show that the treatment status is unrelated to socio-demographic factors. Suppose for example that well educated households complete the survey late in the survey period, i.e. after the resignation, while less educated households complete the survey early. If both types of households differ in their response of inflation expectations to the news, this would bias our findings. We estimate a Probit regression in which we "explain" the treatment status, i.e. the survey being completed within h days after 10/20, on the control variables for different values of h . To save space, we report the results in Table (2) for $h = 2, 7, 10$ only.

With the exception of the employment status for $h = 10$ days, all estimated coefficients remain statistically insignificant. These results suggest that the treatment status

Figure 3: Web traffic, search volume and newspaper reporting around Jens Weidmann's resignation



Notes: The figure shows the daily number of visits of Jens Weidmann's Wikipedia page (panel a), the weekly Google search volume for "Jens Weidmann" in Germany (panel b) and the daily number of German newspaper and magazine articles containing "Jens Weidmann" (panel c) recorded in the Nexis database. The search volume is indexed to 100 in the week with the highest number. The orange dashed line highlights October 20. The gray dashed line indicates the October meeting of the ECB's Governing Council.

cannot be predicted from the socio-demographic control variables. This supports our identifying assumption. In the appendix, we show that oil prices and Covid-restrictions do not confound the comparison between the treatment and the control group. Oil prices and the intensity of lockdowns exhibit very little fluctuations in October 2021.

4 Empirical evidence

In this section, we provide causal evidence on the impact of monetary conservatism on inflation expectations. The announced resignation of Jens Weidmann provides us

Table 2: Determinants of treatment status

	window size		
	$h = 2$	$h = 7$	$h = 10$
constant	−1.050 (0.276***)	−0.641 (0.251**)	−0.575 (0.247**)
male	0.005 (0.049)	−0.010 (0.045)	−0.019 (0.044)
HHsize	0.015 (0.030)	0.013 (0.027)	0.036 (0.026)
big city	−0.007 (0.071)	0.047 (0.065)	−0.006 (0.064)
medium city	−0.021 (0.054)	−0.015 (0.049)	−0.035 (0.048)
age	0.001 (0.002)	0.001 (0.002)	0.001 (0.002)
income _{<2K}	0.040 (0.146)	0.049 (0.133)	0.058 (0.131)
income _{2K–5K}	0.031 (0.126)	0.008 (0.116)	0.014 (0.114)
income _{5K–8K}	−0.097 (0.133)	−0.102 (0.122)	−0.087 (0.120)
east	−0.007 (0.073)	0.039 (0.067)	0.104 (0.066)
north	0.008 (0.073)	0.007 (0.067)	0.011 (0.066)
south	−0.098 (0.061)	−0.035 (0.056)	−0.028 (0.055)
years of edu	0.006 (0.008)	−0.003 (0.007)	−0.004 (0.007)
single	−0.143 (0.082*)	−0.074 (0.073)	−0.034 (0.072)
employed	0.057 (0.061)	0.116 (0.056*)	0.118 (0.055**)
method	Probit	Probit	Probit
# obs.	3747	3747	3747

Notes: The dependent variable is one if the survey was completed within h days after October 20 and 0 otherwise. The regression includes uninterrupted responses only. A significance level of 10%, 5% and 1% is indicated by *, **, ***, respectively.

with a natural experiment.

4.1 Model

We regress the dependent variable, i.e. the mean expected inflation rate, the point forecast, the qualitative forecast, the uncertainty of households about inflation and the degree of trust in the ECB, on a dummy variable that is one for each respondent submitting her answers in a window of h days beginning October 20, 2021 and zero otherwise. Thus, the treatment is the timing of the survey responses.

We also include a constant and the socio-economic control variables explained in the previous section. Since we are interested only in the effect of President Weidmann’s resignation on the dependent variable, we report the coefficient on the dummy variable only. Each regression includes only the responses which were submitted without interruptions. This is because an interruption could potentially violate our identifying assumption. Thus, the regression equation reads

$$y_i = \alpha + \beta \mathbb{D}_{i,h}^{Weidmann} + \Gamma' X_i + \varepsilon_i, \quad (2)$$

where y_i is the dependent variable for respondent i and $\mathbb{D}_{i,h}^{Weidmann}$ is the treatment dummy. The vector X_i contains all control variables. For a survey completed within h days after October 20, the treatment dummy equals one, i.e.

$$\mathbb{D}_{i,h}^{Weidmann} = \begin{cases} 1 & \text{if survey is completed within } h \text{ days after Oct. 20} \\ 0 & \text{otherwise} \end{cases} \quad (3)$$


for $h = 1, \dots, 12$. The first window only includes responses submitted on October 20, the last and widest window includes all submissions between October 20 to October 31. The advantage of this approach is that the number of observations remains unchanged as we increase the window size h . We estimate this equation using robust Huber regressions, which are standard in the literature on survey-based inflation expectations, and OLS regressions with robust standard errors. The estimates of the Huber regressions are less affected by outliers. Some of the dependent variables are measured in discrete steps, e.g. the qualitative expectation of inflation or the level of trust in the ECB. For these variables we also provide results from an Ordered Probit regression. Naturally, the magnitude of the coefficient estimates for the least squares estimation and the Probit estimation are not directly comparable. All regressions include the full set of control variables introduced earlier. As our interest lies on the impact of the news about the announced resignation, we show the estimated β coefficients only.

4.2 Baseline results

Figure (4) shows the reaction of individual mean inflation expectations to the news about the resignation of the Bundesbank president. The estimates of the Huber regressions show a significantly lower mean expectations of households who submit their responses after the resignation. This holds for each possible window length. For the $h = 5$ day window, the mean expected inflation rate is about 0.25 percentage points lower than for pre-resignation respondents. The sign of this response is

counter-intuitive. In contrast, the OLS results remain insignificant. Put differently, the reaction of mean inflation is either counter-intuitive or indistinguishable from zero. Below, we show that the reactions of the point estimates of future inflation and the qualitative assessment of future inflation exhibit notably different responses.

Figure 4: Response of mean inflation expectations to the news about the resignation




Notes: The dots are the point estimates from regression (2). The bars reflect 90% confidence bands.

We now move from the mean of the individual distributions to the uncertainty about future inflation perceived by households. Figure (5) shows the estimates for the interquartile range of the individual distributions. Respondents who complete the survey three days after the resignation or later perceive a significantly higher uncertainty about future inflation than the respondents who submitted their answers before the resignation. Both the Huber regressions and the OLS regressions provide highly significant and strong effects. Households submitting their responses in the $h = 6$ day window after the resignation perceive inflation uncertainty to be about 0.12 points (Huber) or 0.36 points (OLS) higher. This is an economically sizable effect. The sample mean (standard deviation) of the interquartile range is 2.42 (2.66). The results remain qualitatively unchanged if we use the standard deviation of the individual inflation distribution as a measure of forecast uncertainty, see Figure (6). The plot of the response coefficients for each horizon h reveals a hump-shaped pattern. Apparently, households need a few days to receive and process the news about the resignation until they update their expectations. For both the interquartile range and the standard deviation, the OLS estimates are higher than the Huber estimates, which implies that outliers have a strong effect on the OLS estimates.


Figure (7) depicts the impact on the point forecasts of inflation. For each window size

Figure 5: Response of inflation uncertainty (interquartile range) to the news about the resignation



Notes: The dots are the point estimates from regression (2). The bars reflect 90% confidence bands.


Figure 6: Response of inflation uncertainty (standard deviation) to the news about the resignation



Notes: The dots are the point estimates from regression (2). The bars reflect 90% confidence bands.

h , the estimated coefficients remain indistinguishable from zero. The insignificant impact of the news on the point forecast is also consistent with the absence of a market response to the news as shown in Figure (1). Thus, the point forecasts and the mean of the individual forecast distributions exhibit different responses to the news about the resignation.

Figure 7: Response of point inflation expectations to the news about the resignation




Notes: The dots are the point estimates from regression (2). The bars reflect 90% confidence bands.

The reactions of the qualitative expectations of future inflation are notably different from both the mean inflation expectation and the point estimates. Figure (8) shows that households expect a significant increase in inflation as a result of the resignation. The response is significantly positive for responses submitted in the first week after the resignation. The increase in expected inflation is in line with the theoretical predictions for an intermediate horizon discussed before.

Finally, we assess whether the news about the resignation of the Bundesbank president affects the trust of German households in the ECB. The estimates are shown in Figure (9). All three estimation methods, i.e. Ordered Probit, Huber and OLS, yield a consistent result: trust significantly deteriorates for households which submit their responses within four or more days after the resignation. The effect remains highly statistically significant and economically large for all window sizes. Again, plot is hump-shaped suggesting that the maximum impact of the news on trust occurs only a few days. Hence, the unexpected resignation does not only raise uncertainty about the future inflation outlook, but also undermines the trust in the ECB.

In the appendix, we re-estimate the baseline model but (1) include only respondents which have a reasonable perception of current inflation, i.e. an inflation perception between 0% and 10%, (2) exclude households with zero trust in the ECB and (3) include additional control variables at a daily frequency. These control variables capture the tightness of lockdowns in October 2021 and the level of economic activity. All results remain qualitatively unchanged. Wave 22 of the BOP-HH also asks respondents whether they are aware of the recent strategy review of the ECB. In

Figure 8: Response of qualitative inflation expectations to the news about the resignation



Notes: The dots are the point estimates from regression (2). The bars reflect 90% confidence bands.


another alternative specification, we let the treatment dummy interact with a dummy that reflects whether households are aware of the strategy review or not. This interaction term remains insignificant.¹¹ Another alternative specification leaves out survey responses submitted on October 20. The drop in trust and the increase in uncertainty become even stronger than in the baseline model. Finally, we also use data from a survey conducted by Coleman and Nautz (2023a,b). This survey asks only one question, i.e. about medium-term inflation, and the formulation of the question is different from the BOP-HH. We find an increase in the expected inflation rate as a consequences of the announced resignation.

4.3 The role of information about the ECB

Not all survey participants are equally well informed about the mandate and the strategy of the ECB and the role of the Bundesbank president in the Eurosystem. Ideally, we would distinguish between informed and non-informed respondents. Unfortunately, the survey wave in October 2021 does not include a question about economic literacy or general knowledge about the ECB. The only information we could use is the response to the survey question on the recently revised ECB strategy. Each respondent is asked: “Are you aware, e.g. from the media, that the European Central Bank (ECB) has introduced a new monetary policy strategy?” The response

¹¹We do not report these results.

Figure 9: Response of trust in the ECB to the news about the resignation



Notes: The dots are the point estimates from regression (2). The bars reflect 90% confidence bands.

is either yes or no. This question is very specific, but the only opportunity for us to distinguish between informed and uninformed households. We re-estimate the model for inflation uncertainty measured by the interquartile range, but let the treatment interact with a dummy variable ($aware_i$) that is one for informed households and zero otherwise.

Table 3: Interaction with awareness about ECB strategy

	window size							
	$h = 2$		$h = 5$		$h = 8$		$h = 11$	
$\mathbb{D}_{i,h}^{Weidmann}$	-0.043 (0.137)	0.010 (0.058)	0.196 (0.129)	0.116 (0.051**)	0.231 (0.123*)	0.137 (0.049***)	0.251 (0.122**)	0.138 (0.048***)
$aware_i$	-0.067 (0.108)	0.016 (0.045)	-0.118 (0.111)	0.018 (0.048)	-0.093 (0.113)	0.033 (0.049)	-0.087 (0.114)	0.031 (0.050)
$\mathbb{D}_{i,h}^{Weidmann} \times aware_i$	0.488 (0.282*)	0.057 (0.103)	0.474 (0.234**)	0.025 (0.090)	0.361 (0.227)	-0.032 (0.088)	0.330 (0.224)	-0.023 (0.087)
method	OLS	Huber	OLS	Huber	OLS	Huber	OLS	Huber
controls	yes	yes	yes	yes	yes	yes	yes	yes
# obs.	3369	3369	3369	3369	3369	3369	3369	3369

Notes: The dependent variable is the interquartile range. The model also includes a constant.

We show the results in Table (3). For the robust Huber estimations, all results remain unchanged. For window sizes of five days or more, households receiving the news about the resignation perceive a significantly higher uncertainty about inflation. Hence, allowing for informed or uninformed households to respond differently does not change our results. The OLS results are broadly consistent with this, but offer

one additional insight: for window sizes of $h = 2$ and $h = 5$ days, the interaction term is significant, but not the treatment itself. Hence, for these windows, the increase in inflation uncertainty stems from the informed households.

4.4 Updating inflation expectations

In the previous subsection, we found that households' point forecasts for inflation do not respond to the news about the announced resignation. One reason for this could be that the true effect is nonlinear, which results in insignificant estimates when estimated in a linear model. If a household expects a relatively high inflation even before receiving the news, the response could be small or even negative. If a participant foresees a relatively low level of inflation, he might want to revise his expectations upwards in light of the news.

In order to investigate this potential nonlinearity, we concentrate on the point forecasts submitted in the October wave of the survey, $\pi_{i,Oct}^{point}$, but limit the sample to those households which submitted a point forecast also in the September wave, $\pi_{i,Sep}^{point}$. We include $\pi_{i,Sep}^{point}$ as an additional explanatory variable. In addition, we let $\pi_{i,Sep}^{point}$ interact with the treatment dummy. The coefficient β_0 reflects the unconditional effect of the treatment, while the coefficient β_2 reflects the moderating effect of inflation expectations in the previous survey round. The estimated model is

$$\pi_{i,Oct}^{point} = \alpha + \beta_0 \mathbb{D}_{i,h}^{Weidmann} + \beta_1 \pi_{i,Sep}^{point} + \beta_2 (\mathbb{D}_{i,h}^{Weidmann} \times \pi_{i,Sep}^{point}) + \Gamma' X_i + \varepsilon_i. \quad (4)$$

Table (4) reports the estimated coefficients for selected window sizes. Due to the nonlinear nature of the effect, we cannot show the results graphically. The sample size, i.e. the number of households answering both the September and the October survey, is 1241. We find a consistent pattern: the unconditional effect of the news about the resignation is positive and highly statistically significant. Ranging between 1.03 ($h = 2$) and 0.62 ($h = 11$), the effect is also economically large. Past inflation expectations enter with a positive coefficient. Importantly, the coefficient on the interaction term is negative. Hence, the impact of the news about the resignation becomes smaller for higher levels of inflation expectations in September. The opposite signs of the estimated β_0 and β_2 allow us to calculate a critical inflation forecast from the September survey, at which the effect of the news on the October forecast changes sign. Consider the $h = 2$ window. Households expecting future inflation to be below 4.15% in September raise their point forecast in October when receiving the news of Jens Weidmann's resignation. Households expecting inflation to be above the critical rate in the September survey reduce their expectation when receiving

the news. Households with low forecasts raising their inflation expectations in light of the news is in line with our economic intuition. However, it is not clear why households with relatively high expectations in the September round should reduce their expectations. We find this pattern only for the point forecasts, but not for the mean of the individual forecast distributions.

Table 4: Updating inflation expectations

	window size			
	$h = 2$	$h = 5$	$h = 8$	$h = 11$
$\mathbb{D}_{i,h}^{Weidmann}$	1.029 (0.139***)	1.144 (0.124***)	0.591 (0.115***)	0.622 (0.112***)
$\pi_{i,Sep}^{point}$	0.395 (0.005***)	0.400 (0.006***)	0.396 (0.006***)	0.401 (0.006***)
$\mathbb{D}_{i,h}^{Weidmann} \times \pi_{i,Sep}^{point}$	-0.248 (0.018***)	-0.299 (0.017***)	-0.162 (0.016***)	-0.163 (0.016***)
critical $\pi_{i,Sep}^{point}$	4.15%	3.83%	3.65%	3.81%
method	Huber	Huber	Huber	Huber
controls	yes	yes	yes	yes
# obs.	1241	1241	1241	1241

Notes: The table shows the estimated coefficient from regression (4). The model also includes a constant.

5 An illustrative model

Our main result is that the news about the resignation of Jens Weidmann changes the public's uncertainty about inflation, but leaves the mean expectations unaffected. We present a simple model that can rationalize this result. Consider the reduced-form Phillips curve

$$\pi_{t+1} = \pi^* + \theta^j u_t, \quad (5)$$

where π^* is the inflation target. The economy is hit by a shock u_t with $u_t \sim \mathcal{N}(0, \sigma_u^2)$. The central banker reacts to the shock such that the pass-through of the shock to inflation is θ^j , which differs across central bankers.

There are two types of central bankers $j \in \{H, D\}$ reflecting hawks (H) and doves (D), which share the same inflation target. The two types have the same policy objective, but differ in policy implementation. As a result, the expected inflation

rate is independent from the type

$$E(\pi_{t+1}|j) = \pi^*. \quad (6)$$

Hawkish central bankers place a large weight on inflation stabilization and forcefully respond to shocks. Their reaction function is simple. As a result, the pass-through of the shock is weak. Dovish central bankers place less weight on inflation stabilization and respond to signals other than the shock. Their reaction function is more complex. As a result, the pass-through of the shock is stronger. We can summarize this as follows, where the pass-through coefficient for type j is an average response $\bar{\theta}$ plus a type-specific distortion η^j that is orthogonal to the structural shock

$$\theta^j = \bar{\theta} + \eta^j \quad (7)$$

with

$$\text{Var}(\eta^D) > \text{Var}(\eta^H). \quad (8)$$

It follows that

$$\text{Var}(\pi_{t+1} | j) = (\bar{\theta}^2 + \text{Var}(\eta^j)) \sigma_u^2 \quad (9)$$

The inflation rate is more predictable under the hawkish central banker. This model suggests that households should not change their mean inflation forecasts after receiving their news about a possible shift from a policy hawk to a dove. This is in line with our evidence and also consistent with the missing market reaction around the news. The increase in the variance of expected inflation should correspond to an increase in individual inflation uncertainty of survey participants. Thus, we should expect a higher inflation uncertainty relative to the control group of households which do not receive the news. This is exactly what we find in the data.

6 Conclusions

A shift in the monetary policy preferences of central bank governors and, as a consequence, of central bank committees should change the expectations of households. To establish a causal effect, this shift needs to be unexpected. The announced resignation of Jens Weidmann from the presidency of the Deutsche Bundesbank is a rare case of a surprise change in the average preferences of the ECB's Governing Council. Since Weidmann was an outspoken policy hawk and the newly elected German government was led by the Social Democrats, the likely successor picked by the new government was more dovish.

In this paper, we showed that households adjust their inflation expectations as a response to the news about the resignation. We find a strong increase in individual uncertainty about future inflation and a deterioration of trust in the ECB. The response of households' mean inflation expectations remains weak or insignificant. Inflation expectations remain well anchored even in the short-run, while the uncertainty around inflation increases.

One conclusion from our findings is that personalities matter for households' expectations, even in an institutional environment with a highly independent central bank, a clear mandate for price stability and a credible inflation target. Appointing a known inflation hawk contains inflation uncertainty and raises trust. Furthermore, the preferences of an outspoken and prominent policymaker matter although monetary policy decisions are made by a committee, not a single decision maker. Moreover, the preferences of the central banker matter, at least in his or her home country, even if the policymaker is in a minority position in the Governing Council.

References

- Aikman, D., Monti, F. and Zhang, S. (2025). In the Fed we trust? Measuring trust in central banking and its effects on the macroeconomy, *unpublished*.
- Alesina, A. and Gatti, R. (1995). Independent central banks: low inflation at no cost?, *The American Economic Review (Papers and Proceedings)* **85**: 196–2000.
- Altavilla, C., Brugnolini, R. S., Gürkaynak, R., Motto, R. and Ragusa, G. (2019). Measuring euro area monetary policy, *Journal of Monetary Economics* **108**: 162–179.
- Backus, D. and Driffill, J. (1985). Inflation and reputation, *American Economic Review* **75**: 530–538.
- Barro, R. J. and Gordon, D. B. (1983). Rules, discretion and reputation in a model of monetary policy, *Journal of Monetary Economics* **12**: 101–121.
- Beetsma, R. M. W. J. and Jensen, H. (1998). Inflation targets and contracts with uncertain central bank preferences, *Journal of Money, Credit and Banking* **30**: 384–403.
- Berger, H. and Woitek, U. (2005). Does conservatism matter? A time-series approach to central bank behaviour, *The Economic Journal* **505**: 745–766.
- Carvalho, C., Flórido, T. and Zilberman, E. (2017). Transitions in central bank leadership, *Available at SSRN 2997879*.
- Christelis, D., Georgarakos, D., Jappelli, T. and van Rooij, M. (2020). Trust in the central bank and inflation expectations, *International Journal of Central Banking* **16**: 1–37.
- Clarida, R., Galí, J. and Gertler, M. (1999). The science of monetary policy: a New Keynesian perspective, *Journal of Economic Literature* **37**: 1661–1707.
- Coibion, O., Georgarakos, D., Gorodnichenko, Y., Kenny, G. and Weber, M. (2024). The effect of macroeconomic uncertainty on household spending, *American Economic Review* **114**: 645–677.
- Coibion, O., Gorodnichenko, Y., Kumar, S. and Pedemonte, M. (2020). Inflation expectations as a policy tool?, *Journal of International Economics* **124**: 103297.
- Coibion, O., Gorodnichenko, Y. and Weber, M. (2022). Monetary policy communications and their effects on household inflation expectations, *Journal of Political Economy* **130**(6): 1537–1584.
- Coleman, W. and Nautz, D. (2023a). Inflation expectations, inflation target credibility, and the COVID-19 pandemic: evidence from Germany, *Journal of Money, Credit and Banking* **55**: 1937–1953.
- Coleman, W. and Nautz, D. (2023b). Inflation target credibility in times of high inflation, *Economics Letters* **222**: 110930.

- Constâncio, V. (2011). Laudatio for Professor Kenneth Rogoff, recipient of the Deutsche Bank Prize in Financial Economics, *September 11, 2011*.
- De Fiore, F., Lombardi, M. and Pierres-Tejada, A. (2025). Fiscal stimulus plans and households' expectations, *unpublished*.
- Dräger, L., Gründler, K. and Potrafke, N. (2025). Political shocks and inflation expectations: evidence from the 2022 Russian invasion of Ukraine, *Journal of International Economics* **153**: 104029.
- Dräger, L., Lamla, M. J. and Pfajfar, D. (2025). How to limit the spillover from an inflation surge to inflation expectations?, *Journal of Monetary Economics* **144**: 103546.
- Dreher, A., Sturm, J.-E. and de Haan, J. (2008). Does high inflation cause central bankers to lose their job? Evidence based on a new data set, *European Journal of Political Economy* **24**: 778–787.
- Dreher, A., Sturm, J.-E. and de Haan, J. (2010). When Is a central bank governor replaced? Evidence based on a new data set, *Journal of Macroeconomics* **32**: 766–781.
- Ehrmann, M. (2024). Trust in central banks, *ECB Working Paper* (No. 3006).
- Eickmeier, S. and Petersen, L. (2024). Toward a holistic approach to central bank trust, *NBER Working Paper* (No. 32716).
- Engelberg, J., Manski, C. F. and Williams, J. (2009). Comparing the point predictions and subjective probability distributions of professional forecasters, *Journal of Business & Economic Statistics* **27**(1): 30–41.
- Fischer, J. J., Herler, C. and Schnattinger, P. (2024). The effect of inflation uncertainty on household spending, *Bank of England, unpublished*.
- Göhlmann, S. and Vaubel, R. (2007). The educational and occupational background of central bankers and its effect on inflation: an empirical analysis, *European Economic Review* **51**: 925–941.
- Grebe, M. and Tillmann, P. (2025). Household expectations and dissent among policymakers, *European Journal of Political Economy* **86**: 102638.
- Hahn, V. (2016). Designing monetary policy committees, *Journal of Economic Dynamics and Control* **65**: 47–67.
- Hoffmann, M., Moench, E., Pavlova, L. and Schulte Frankenfeld, G. (2022). Would households understand average inflation targeting?, *Journal of Monetary Economics* **129**: S52–S66.
- Ioannidou, V., Kokas, S., Lambert, T. and Michaelides, A. (2022). (In)dependent central banks, *CEPR Discussion Paper* (No. 17802).

- Kostyshyna, O. and Petersen, L. (2024). The effect of inflation uncertainty on household expectations and spending, *NBER Working Paper* (No. 32939).
- Kuttner, K. and Posen, A. (2010). Do markets care who chairs the central bank?, *Journal of Money, Credit and Banking* **42**: 347–371.
- Kydland, F. E. and Prescott, E. C. (1977). Rules rather than discretion: the inconsistency of optimal plans, *Journal of Political Economy* **85**: 473–492.
- Manski, C. F. (2018). Survey measurement of probabilistic macroeconomic expectations: Progress and promise, *NBER Macroeconomics Annual* **32**: 411–471.
- Mehrling, P., Moss, L. S., Pixley, J. and Tavlas, G. S. (2007). What if the leader of the central bank told hilarious jokes and did card tricks? A panel of experts, *American Journal of Economics and Sociology* **66**: 863–887.
- Moser, C. and Dreher, A. (2010). Do markets care about central bank governor changes? Evidence from emerging markets, *Journal of Money, Credit and Banking* **42**: 1589–1612.
- Neuenkirch, M. and Tillmann, P. (2016). Does a good central banker make a difference?, *Economic Inquiry* **54**: 1541–1560.
- Rogoff, K. (1985). The optimal degree of commitment to an intermediate monetary target, *The Quarterly Journal of Economics* **100**: 1169–1189.
- Romelli, D. (2022). The political economy of reforms in central bank design: evidence from a new dataset, *Economic Policy* **37**: 641–688.
- Schaumburg, E. and Tambalotti, A. (2007). An investigation of the gains from commitment in monetary policy, *Journal of Monetary Economics* **54**: 302–324.
- Sibert, A. (2002). Monetary policy with uncertain central bank preferences, *European Economic Review* **46**: 1093–1109.
- Tillmann, P. and Walter, A. (2019). The effect of diverging communication: The case of the ECB and the Bundesbank, *Economics Letters* **176**: 68–74.
- Vickers, J. (1986). Signalling in a model of nonetary policy with incomplete information, *Oxford Economic Papers* **38**: 443–455.
- Walsh, C. E. (1995). Optimal contracts for central bankers, *The American Economic Review* **85**: 150–167.

Appendices

Appendix A Official Bundesbank communication

Figure (A.1) shows the official tweets, i.e. the official messages posted on X (formerly Twitter) announcing the resignation of Jens Weidmann, in German and English.

Figure A.1: Tweets from @bundesbank announcing the resignation



Notes: Both tweets were posted at 10:45 on October 20, 2021, i.e. at the same time of the press release.

Jens Weidmann also released a letter to the Bundesbank staff.¹² In her official response, President Lagarde stresses her close personal relations with Jens Weidmann, but does not comment his policy preferences.¹³

¹²See <https://www.bundesbank.de/en/press/press-releases/letter-from-the-president-to-bundesbank-staff-877834>.

¹³See <https://www.ecb.europa.eu/press/pr/date/2021/html/ecb.pr211020~d413bec716.en.html>.

Appendix B Coverage by tabloid media

Jens Weidmann’s resignation was covered widely by all sections of the German media. Figure (B.2) gives an example from bild.de, Germany’s largest tabloid paper. The headlines says (translation by DeepL): “Disaster for our money!: Germany’s guardian of the stable euro gives up”. The subtitle says: “In the midst of monster inflation, of all times: Jens Weidmann wanted to protect our money from devaluation – but the SPD and Green parties have other plans.”

Figure B.2: Article on bild.de



Notes: Appeared on October 20, 2021, on bild.de, Germany’s largest tabloid newspaper.

Appendix C Individual probability distributions

To fit individual distributions based on respondents’ answers to probability questions, we follow the well-established approach introduced by Engelberg et al. (2009). Respondents allocate probability mass across ten bins. Depending on how many of these bins are used (i.e., receive non-zero weight), we apply different distributional assumptions.

One bin: If the respondent assigns all probability to a single bin, we assume a symmetric triangular distribution bounded by the left and right limit of that bin. The mode is set at the midpoint.

Two adjacent bins: If two adjacent bins are used, we determine the bounds of a symmetric triangular distribution such that the cumulative distribution function (CDF) matches the relative weight placed in the first bin. Let x denote the boundary between the two bins and p_{left} the cumulative probability to the left of x (i.e., in the first bin). Depending on which bin has more weight, either the lower or upper bound is fixed at the bin limit, and the other bound

is computed to satisfy the implied cumulative probability. Specifically, for a symmetric triangular distribution with bounds a and b (mode $c = (a + b)/2$), the bounds are determined as follows:

$$\text{if the first bin dominates (fix } a): \quad b = a + \frac{x - a}{\sqrt{p_{\text{left}}/2}},$$

$$\text{if the second bin dominates (fix } b): \quad a = \frac{x - \sqrt{p_{\text{left}}/2} b}{1 + \sqrt{p_{\text{left}}/2}}.$$

If both bins have equal weights, a is set to the left edge of the first bin and b to the right edge of the second bin.


Three or more bins: If more than two bins are used, we fit a four-parameter generalized beta distribution $\beta(a, b, l, r)$. The parameters are chosen to minimize the sum of squared deviations between the empirical cumulative distribution and the theoretical CDF.

From the fitted distributions, we compute the mean, the standard deviation and the interquartile range for each respondent.

Appendix D Other developments in October 2021


Our identification rests on the assumption that the expectations of the treatment group and the control group are equally affected by other macroeconomic factors. We do not want that other factors confound the comparison between both groups. To major forces driving inflation expectations in 2021 are changes in oil prices and changes in the intensity of Covid-restrictions. Figure (D.3) plots the price of Brent in euro. We clearly spot the strong price increase in 2020 and 2021. However, during October 2021 prices remain relatively stable. Thus, we can rule out that the treatment group faces systematically higher oil prices than the control group. Figure (D.4) shows the evolution of the stringency index for Germany, which measures the intensity of restrictions during the pandemic. Again, we spot drastic changes since the spring of 2020. Importantly, the index exhibits only minuscule change in October 2021. There are no systematic difference in the lockdown intensity between the treatment and the control group. Below, we also use the stringency index as a control variable in the regressions.

Figure D.3: The evolution of oil prices



Notes: Price of one barrel of Brent in euro. The dashed vertical line highlights October 20, the day Jens Weidmann announced his resignation.

Figure D.4: The evolution of Covid-restrictions



Notes: The stringency index measures the severity of restrictions in Germany during the Covid-19 pandemic. The dashed vertical line highlights October 20, the day Jens Weidmann announced his resignation.


Appendix E Additional control variables

The resignation of Jens Weidmann falls into the later stage of the Covid-19 pandemic. In order to include information about the state of economic activity and the intensity of lockdowns, which could both impact households' inflation expectations as well as their uncertainty about future inflation, we add two variables to the set of controls. Both are available at the daily frequency. The first variable is the Stringency Index for

Germany, which measures the intensity of lockdown measures. The second variable is the Google Mobility Index for Germany, which uses data on cell phone locations to measure workplace-related mobility.

Figures (E.5), (E.6) and (E.7) show the estimated results for this extended model. We show the results for the mean expected inflation rates, the uncertainty about inflation and the level of trust in the ECB only. Compared to the results from the main part of the paper, all results remain qualitatively unaffected.

Figure E.5: Response of mean inflation expectations to the news about the resignation (with additional controls)




Notes: The dots are the point estimates from regression (2). The point estimates from the baseline model are highlighted by the star. The bars reflect 90% confidence bands.

Appendix F “Reasonable” households only

A common feature of household surveys of inflation expectations is that households have very heterogeneous knowledge about actual inflation. Grossly misjudging the prevailing level of inflation could lead to extreme values of expected inflation and other variables. The BOP-HH survey explicitly asks participants about the current inflation rate. In this subsection, we re-estimate the empirical model but include only those households with reasonable perceptions of inflation, i.e. households which think actual inflation at the time of the survey is between 0% and 10%. As mentioned earlier, the latest available CPI inflation rate for Germany in October 2021 was 4.1%. We show the results for mean inflation, Figure (F.8), the interquartile range, (F.9) and the level of trust, Figure (F.10), only. All results remain qualitatively unchanged compared to the baseline findings.

Figure E.6: Response of inflation uncertainty (interquartile range) to the news about the resignation (with additional controls)



Notes: The dots are the point estimates from regression (2). The point estimates from the baseline model are highlighted by the star. The bars reflect 90% confidence bands.

Appendix G Excluding households with zero trust


Figure (G.11) shows a histogram for the level of trust in the ECB for responses submitted before and after the resignation. Around 12% (before) and 14% (after) of the respondents say they do not trust the ECB at all.

We want to rule out that the results on trust are driven by the respondents with zero trust in the ECB. Therefore, we re-estimate the model but exclude all respondents with zero trust either before or after the news about the resignation. The results are shown in Figure (G.12). While the impact of the news on the level of trust is a bit weaker than in the baseline results, the significantly negative effect is still visible.

Appendix H Excluding October 20

Each of the estimation windows in the baseline estimation includes the responses submitted on October 20. Since we do not have the exact time stamp of each submission, this could include responses submitted before 10:45, the time of the news. In addition, many households might have received the news in the evening or even the next morning. Therefore, we estimate the model but exclude all responses submitted on October 20 from the treatment group and the control group. Thus, the horizon h^* now includes $h - 1$ days. Figures (H.13) to (H.15) shows the resulting estimates. The impact of the news about the resignation is somewhat stronger than

Figure E.7: Response of trust in the ECB to the news about the resignation (with additional controls)



Notes: The dots are the point estimates from regression (2). The point estimates from the baseline model are highlighted by the star. The bars reflect 90% confidence bands.

in the baseline case, in particular for trust in the ECB.

Appendix I The announcement of the successor


The government announced Joachim Nagel as the new president of the Bundesbank on December 2021. The announcement day falls into wave 24 of the BOP-HH survey. However, only 102 out of 3365 respondents submitted their answers before December 20. Figure (I.16) shows the distribution of the survey responses of wave 24. The number of responses submitted before the announcement is too small to conduct an empirical analysis in order to estimate whether households responded to the announcement of Jens Weidmann's successor.

His succession did not come as a surprise to informed households as Joachim Nagel's name had been mentioned in the press before as a likely successor. Furthermore, Joachim Nagel was not a household name (unlike Jens Weidmann). Thus, we do not expect households to respond to the news.

Appendix J Evidence from an alternative survey

We corroborate our main results obtained from the BOP-HH survey using survey data from Coleman and Nautz (2023b) and Coleman and Nautz (2023a). These authors run a daily survey among German households conducted by Civey and ask:

Figure F.8: Response of mean inflation expectations to the news about the resignation (reasonable responses only)




Notes: The dots are the point estimates from regression (2). The point estimates from the baseline model are highlighted by the star. The bars reflect 90% confidence bands.

“In what range do you think the annual inflation rate will be over the medium term?”. Households can choose between inflation being “(A) clearly above 2%”, “(B) slightly above 2%”, “(C) below, but close to 2%”, “(D) clearly below 2%” or “(N) do not know”. This formulation matches the definition of the ECB’s inflation target (“below, but close to 2%”) before its revision in 2021.

We use the survey responses submitted between October 15 and October 31, 2021 to remain consistent with the BOP-HH data used in the main part of the paper. In this period, 2,100 respondents completed the survey. Figure (J.17) shows the distribution of the responses over the days. Many responses are dated October 15 and October 29 and few responses are available for the remaining days. In particular, we have few observations immediately before and after October 20.

We transform answer (A) into a value of +2, answer (B) into a value of +1, answer (C) into a value of 0 and answer (D) into a value of -1. This provides us with an indicator of individual inflation expectations. We do not use answers (N) at this point. However, we construct a separate dummy variable that is one if a respondent chooses answer (D) and zero otherwise. As we lack a better measure in this survey, we interpret this variable as a measure of inflation uncertainty. Of course, the “don’t know” answer also reflects a respondent’s attention to inflation or his economic literacy. In addition, the survey contains a rich set of sociodemographic information about age, education, regional purchasing power etc., which we transform into control variables following the construction of the variables in the main part of the paper.


Figure F.9: Response of inflation uncertainty (interquartile range) to the news about the resignation (reasonable responses only)



Notes: The dots are the point estimates from regression (2). The point estimates from the baseline model are highlighted by the star. The bars reflect 90% confidence bands.


We estimate the model from the main part of the paper using the Civey survey. For both variables, the expected inflation rate and the “don’t know” dummy variable, the robust least squares regressions deliver error messages, probably due to multicollinearity. This leaves us with the OLS estimates and the (Ordered) Probit regressions for the two variables. Figure (J.18) shows the estimated coefficients for the expected inflation rate. For both estimation methods, we find a significant increase of expected inflation for windows larger than $h = 3$ days. Figure (J.19) shows the reaction of the “don’t know” dummy variable to the news about the resignation. This variable is a noisy proxy for the degree of inflation uncertainty. We find that uncertainty responds positively to the news on October 20. The response is particularly significant for larger windows, which is in line with the distribution of the survey responses shown in Figure (J.17). The smaller windows simply do not include enough observations to find a significant effect.

Figure F.10: Response of trust in the ECB to the news about the resignation (reasonable responses only)




Notes: The dots are the point estimates from regression (2). The point estimates from the baseline model are highlighted by the star. The bars reflect 90% confidence bands.

Figure G.11: Distribution of trust in the ECB




Notes: The figure shows a histogram of the level of trust across all uninterrupted survey responses. Trust can range from 0 (no trust) to 10 (full trust).

Figure G.12: Response of trust in the ECB to the news about the resignation (excluding no trust)




Notes: The dots are the point estimates from regression (2). The point estimates from the baseline model are highlighted by the star. The bars reflect 90% confidence bands.

Figure H.13: Response of mean inflation expectations to the news about the resignation (excluding October 20)




Notes: The dots are the point estimates from regression (2). The point estimates from the baseline model are highlighted by the star. The bars reflect 90% confidence bands.

Figure H.14: Response of inflation uncertainty (interquartile range) to the news about the resignation (excluding October 20)




Notes: The dots are the point estimates from regression (2). The point estimates from the baseline model are highlighted by the star. The bars reflect 90% confidence bands.

Figure H.15: Response of trust in the ECB to the news about the resignation (excluding October 20)




Notes: The dots are the point estimates from regression (2). The point estimates from the baseline model are highlighted by the star. The bars reflect 90% confidence bands.

Figure I.16: Number of survey responses per day (wave 24 - December)




Notes: The figure shows the number of survey responses per day in December 2021 and January 2022, respectively, of wave 24 of the BOP-HH. The dashed vertical line highlights December 20, the day the government announced Joachim Nagel as the new president.

Figure J.17: Number of survey responses per day (Civey survey)




Notes: The figure shows the number of survey responses per day in October 2021 in the Civey-survey conducted by Coleman and Nautz (2023a,b). The dashed vertical line highlights October 20, the day Jens Weidmann announced his resignation.

Figure J.18: Response of inflation expectations to the news about the resignation (Civey survey)



Notes: The dots are the point estimates from regression (2). The bars reflect 90% confidence bands.

Figure J.19: Response of inflation uncertainty to the news about the resignation (Civey survey)



Notes: The dots are the point estimates from regression (2). The bars reflect 90% confidence bands.

Essay VI

Prices at Prime Time: How TV News Shape Inflation Attention in Germany

This paper was presented at the following workshops and international conferences:

- i. ECONDAT spring meeting: Economics with nontraditional data and analytical tools (London)

Date: June 2025

Presenter: Moritz Grebe

Acknowledgement

We thank Andreas Schmidt and Sönke Treu from the NDR Archive for their support and for welcoming us in Hamburg. We also thank Andreas Joseph for discussing an early version of the paper. Dimitrios Kanelis, Emanuel Mönch, Peter Tillmann, and participants at the ECONDAT 2025 Spring Meeting in London provided valuable feedback. We are grateful to Friederike Schimmel, Laurenz Burggraf, and Nicole Panz for excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Deutsche Bundesbank or the Eurosystem.

Prices at Prime Time: How TV News Shape Inflation Attention in Germany

Brenton Joey Bruns* Moritz Grebe†

Abstract

We provide new long-run evidence on inflation coverage in television news and its transmission to households. We study the *Tagesschau*, Germany's main evening news program, and construct two datasets on inflation coverage: a historical archive dataset from 1956 to 2024 and a video-based dataset from 1989 to 2024. These data measure the amount, prominence, direction, and sentiment of inflation news. We show that *Tagesschau* coverage reacts nonlinearly to inflation, increasing much more strongly once inflation exceeds a threshold. Since 1999, this threshold lies between 2.1% and 2.8%, close to the ECB's inflation target. Linking these data to household surveys, IV local-projection estimates suggest persistent effects on household attention, inflation perceptions, and inflation expectations. Household-level evidence shows that news on the previous day raises attention and that belief responses depend on whether reports cover rising or falling prices. Overall, television news is an important channel through which inflation information reaches households.

Keywords: inflation expectations, household attention, media coverage, television news, Tagesschau

JEL Classification: D83, D84, E31, E52

*Justus Liebig University Giessen, Email: brenton.joey.bruns@wirtschaft.uni-giessen.de

†Deutsche Bundesbank, Email: moritz.grebe@bundesbank.de

1 Introduction

Household attention to inflation is central to the formation of inflation perceptions and expectations, and therefore to the transmission of monetary policy. Yet attention varies substantially over time: it appears to be low in stable low-inflation environments, but rises when inflation becomes high or volatile (Bracha and Tang, 2025; Weber et al., 2025; Pfäuti, 2025; Buelens, 2023). Less is known, however, about the channels through which inflation becomes salient to households.

This paper studies television news as one such channel. Media coverage can make inflation visible as a public issue and select which price developments receive attention. Existing work has mostly focused on newspapers, digital news, and search behavior.¹ Television news has received much less attention, although it remains an important source of information for households.

This is particularly true for Germany. In the Bundesbank Online Panel Households (BOP-HH), television/radio was still the most frequently mentioned source of economic information in December 2024 (Figure (1)). Television news is therefore a relevant channel through which economic information may reach households beyond those who actively search for it. Unlike newspapers or online news, a television broadcast gives viewers less control over which individual topics they are exposed to once they watch the program.

We focus on the *Tagesschau*, Germany’s main evening news program. Its 8 p.m. edition reaches millions of viewers, summarizes the day’s most important news in a short format, and is widely seen as a credible source of information (Newman et al., 2025). These features make the *Tagesschau* a useful setting to study how inflation news reaches households.

A central contribution of the paper is measurement. We construct two new datasets on inflation coverage in the *Tagesschau*. The first is a long-run archive dataset from 1956 to 2024. It is based on NDR archive metadata and records inflation-related reports, including their frequency, duration, and position within the broadcast. We combine these dimensions into a historical *Tagesschau* attention index. The second dataset uses all available videos of the 8 p.m. broadcasts from September 1989 to August 2024. We transcribe the videos, identify inflation-related reports, and classify their content by price direction and sentiment. These data allow us to measure both the amount and the content of inflation coverage over a long period.

We document three main findings. First, inflation coverage in the *Tagesschau* varies substantially over time and across inflation episodes. Two periods stand out in the

¹See, e.g., Lamla and Lein (2014), Dräger (2015), and Arndt (2024) for evidence based on newspapers; Macaulay and Song (2023) for digital news; and Buelens (2023) for search behavior.

long-run archive data: the oil-price shocks of the 1970s and the post-Covid inflation surge. In both episodes, inflation received particularly strong television coverage. Other periods with elevated inflation, such as the 1980s and 1990s, received less coverage.

Second, *Tagesschau* coverage reacts nonlinearly to inflation. Using a threshold model, we find that media attention rises much more strongly once inflation exceeds a certain level. For the full sample, the estimated threshold is about 5.4%. In the Bundesbank period, from January 1960 to December 1998, it is about 5.7%. In the ECB period, from January 1999 to July 2024, the threshold is much lower, between 2.1% and 2.8%, close to the ECB's inflation target. Above the estimated threshold, the response of inflation coverage to inflation is roughly twice as large in the ECB period as in the Bundesbank period. This suggests that inflation coverage has become more sensitive to inflation developments over time.

Third, inflation-related TV coverage is transmitted to households. In aggregate time-series data, months with more *Tagesschau* coverage are also months in which households pay more attention to inflation and report higher inflation perceptions. To address reverse causality and omitted changes in the inflation environment, we use an instrumental variable local-projections approach based on inflation surprises covered in *Tagesschau* reports about preliminary inflation releases. The results suggest economically sizable and persistent effects on household outcomes. The response is strongest for inflation perceptions, which increase substantially for about one year, while household attention responds more persistently over the full 24-month horizon. Inflation expectations also rise, but less persistently. Household-level evidence from the BOP-HH points in the same direction: inflation-related news coverage on the previous day is associated with higher household attention, while belief responses depend on the direction of the news. During the recent high-inflation period covered by the survey, reports about falling prices are associated with lower inflation perceptions and expectations, whereas reports about rising prices do not significantly move beliefs.

Our paper contributes to different strands of the literature. First, it relates to work on attention to inflation and information frictions in expectation formation. A large theoretical literature studies how costly information processing and limited attention affect expectation formation and macroeconomic dynamics.² Empirically, much of the evidence on limited attention is indirect: forecasts respond incompletely to new information, and agents revise their expectations when information is made salient (Carroll, 2003; Coibion and Gorodnichenko, 2012; Cavallo et al., 2017). Relatedly,

²See, e.g., Maćkowiak et al. (2023) for a review of the rational-inattention literature.

Larsen et al. (2021) show that media coverage helps explain time-varying information rigidities in households' inflation expectations. Other work measures attention to inflation more directly. Coibion et al. (2018) show that many firms pay little attention to inflation when it is not important for their business decisions, while Bracha and Tang (2025) document that household inattention falls when inflation rises. Weber et al. (2025) similarly find that households and firms are more attentive and better informed when inflation is high, while Pfäuti (2025) documents a decline in inflation attention after the Great Inflation and relates this decline to lower inflation volatility and persistence.

Several papers find nonlinearities in inflation attention. For the United States, Pfäuti (2023) estimates an attention threshold of around 4% inflation and shows that attention roughly doubles above this level. Using cross-country evidence, Korenok et al. (2026) find that such thresholds typically lie between about 2% and 4%. For the euro area, Buelens (2023) uses Google search data to show that inflation attention rises nonlinearly with inflation and moves across different attention regimes. We contribute to this literature by showing that television coverage of inflation is also nonlinear and may be one channel through which nonlinear household attention arises.

Second, our paper relates to the literature on media coverage of inflation and monetary policy. For Germany, Lamla and Maag (2012) show that more intensive inflation reporting and more heterogeneous news content increase disagreement in households' inflation forecasts. Lamla and Lein (2014) further show that more intensive reporting can improve the accuracy of consumers' inflation expectations, while the tone of reporting may bias expectations. For Sweden, Dräger (2015) finds that media reports help explain the link between actual inflation, inflation perceptions, and inflation expectations, with asymmetric responses to news about rising and falling inflation. Closely related, Chahrour et al. (2025) show that news selection itself is state-dependent: inflation is more likely to be covered when inflation is high, and this selection affects household inflation expectations.

Other papers study the content and framing of inflation news in greater detail. Macaulay and Song (2023) show that U.S. inflation news contains both a financial "Wall Street" narrative and a household-oriented "Main Street" narrative. The latter becomes more prominent during the recent inflation surge and is associated with more pessimistic consumer sentiment. For Germany, Arndt (2024) shows that tabloid inflation news is more closely related to the perceptions and expectations of lower-income and lower-education households, while quality newspapers are more closely related to the beliefs of higher-income and higher-education households. Müller et al.


(2022) document how the framing of price developments in German media changes over time, and Heikkinen and Heimonen (2025) show that media tone helps explain heterogeneity in inflation expectations. A related literature studies media coverage of monetary policy (Berger et al., 2011; Picault et al., 2022). Overall, this literature shows that media do not only transmit macroeconomic information, but also select, frame, and evaluate it.

Third, our paper relates to a small but growing literature on television as a channel through which macroeconomic news reaches households. Hirsch et al. (2023) show that households' inflation expectations respond more strongly to ECB announcements when these decisions receive television news coverage. Recent papers study television coverage of inflation more directly. Xu et al. (2024) find that official television news affects household inflation expectations in China. For the United States, Binder et al. (2025) show that greater television news coverage of inflation causally raises household inflation expectations. Çekin and Polattimur (2025) construct a television news index for major U.S. channels and find that inflation coverage on CNN and Fox News affects household inflation expectations, especially among women, lower-income, and less-educated respondents during periods of rising inflation.

Relative to this literature, we provide long-run evidence on inflation coverage in television news and its transmission to households. While existing work on television and inflation focuses mainly on recent episodes, our archive data allow us to study how inflation coverage evolved in Germany from the postwar period to the recent inflation surge. Combining archive metadata with video transcripts, we measure not only the amount of inflation coverage, but also its prominence and content. We then link this coverage to several stages of household belief formation: attention, perceptions, and expectations.

The paper proceeds as follows. Section 2 describes the *Tagesschau* and introduces our archive and video datasets. Section 3 studies the nonlinear relationship between inflation and *Tagesschau* attention. Section 4 analyzes the transmission of inflation-related TV coverage to households using aggregate time-series evidence, IV local projections, and household-level data. Section 5 concludes.

Figure 1: German households' main sources of economic information




Notes: This figure shows responses to question P6007 from wave 60 (December 2024) of the BOP-HH. Shares are weighted using household weights. Respondents were asked: “What is your main source of information on topics related to the economy?” Up to two items could be selected. “Newspapers” includes online newspapers.

2 The Tagesschau

The *Tagesschau* is Germany’s oldest and most widely watched television news program. Its first edition was broadcast in December 1952, at a time when television was still a new medium and only around 1,000 viewers were able to receive the signal.

Initially, the program aired three times per week. It was expanded to six editions per week in October 1956 and has been broadcast daily since September 1961.³ Over time, additional daily editions were introduced, but the 8 p.m. broadcast remains the central edition. It lasts 15 minutes and summarizes the day’s most important news. Its institutional role in German television is reflected in the fact that many channels schedule their evening prime-time programming for 8:15 p.m., immediately after the *Tagesschau* ends. Despite the growing importance of online media, the 8 p.m. edition has continued to reach an average audience of more than 9 million viewers for many years (Figure (A.1)). This is more than nine times the average circulation of major daily newspapers such as the *Frankfurter Allgemeine Zeitung* (FAZ) or the *Süddeutsche Zeitung* (SZ) (Figure (2)). The *Tagesschau* is not only widely watched, but also widely trusted: it is regularly the most trusted news source in Germany (Figure (3)).

Figure 2: Reach of the *Tagesschau* and other news sources




Notes: This figure compares the average daily audience of the *Tagesschau* in millions of viewers by year with circulation or reach measures for other established news sources. Abbreviations in parentheses indicate the type of news source: television (TV), newspaper (NP), or online. All data are from Statista.

These features make the *Tagesschau* a useful setting for studying the information that enters households’ information sets. As a public-service news program broadcast across the ARD network, the *Tagesschau* provides a relatively common source of

³For a brief overview of the early history of the *Tagesschau*, see https://www.ndr.de/der_ndr/unternehmen/chronik/Vom-Kellerkind-zur-Institution-die-Geschichte-der-Tagesschau,ardaktuell101.html (in German).

news exposure across households. Television news differs from newspapers and online search because viewers do not need to actively search for specific economic information. Instead, they are exposed to economic news as part of a general news program. The short and fixed format of the 8 p.m. edition may also make it more likely that viewers receive and process the information, rather than selectively skimming or searching for content. For these reasons, inflation coverage in the *Tagesschau* is likely to capture a form of macroeconomic information that reaches a broad household audience with relatively low search costs.

Figure 3: In the *Tagesschau* we trust



Notes: This figure shows trust in selected German news brands from the *Reuters Institute Digital News Report 2025* (Newman et al., 2025). Trust is based on the question: “How trustworthy would you say news from the following brands is?” on a scale from 0 (“not at all trustworthy”) to 10 (“completely trustworthy”). Responses from 6–10 are coded as “Trust”, 5 as “Neither”, and 0–4 as “Don’t trust”. Respondents who had not heard of a given brand were excluded. Abbreviations in parentheses indicate the main type of news source: television (TV), newspaper (NP), or online.

2.1 The archive dataset

The first dataset consists of metadata we collect at the NDR archive in Hamburg. We search for broadcasts that can be associated to certain keywords such as "inflation" or "price increase".⁴ We intentionally search not only for reports that focus narrowly on inflation itself, but also for broader reports on price developments. Such reports may also shape households’ inflation beliefs, even though changes in individual prices do not necessarily reflect inflation in the macroeconomic sense.

⁴See table A.1 for the full list.

The search returns archive records with metadata, including short content descriptions, the total length of each broadcast, the length of the relevant segment, and timestamps indicating where the segment appeared in the broadcast.⁵ The archive does not search the entire text of the TV reports, but only the title, short description and tags given to the reports.

We manually remove editions that are clearly unrelated to inflation developments in Germany. For instance, because one of our keywords is the German word for “price” which may also refer to an “award” or “prize”, we exclude broadcasts on prizes or awards. We also exclude reports about inflation developments in other countries.

Overall, we identify 862 inflation-related *Tagesschau* editions covering the period from May 1956 to July 2024. From this dataset we extract three variables that could be informative for the attention of the *Tagesschau* with regard to inflation.

First, we count the *Tagesschau* editions per month. If the *Tagesschau* is more attentive to inflation in one month, the topic should appear in more editions.

Second, we use the duration of the report within a *Tagesschau*. The 8 p.m. edition is limited to 15 minutes. Therefore, more important topics should take up more time in an edition. We calculate the relative duration by dividing the duration of the TV report in seconds by the total duration of the daily edition and calculate the monthly average.

Third, we use the positioning of the report within the daily news. We assume that topics are positioned in the *Tagesschau* based on their importance. The top news comes right at the beginning, less important topics at the end. We divide the start time of the report in seconds by the total duration of the edition. We then calculate one minus this proportion and obtain our measure for the positioning in a news broadcast. A value of 1 means that the *Tagesschau* starts with a report on inflation. The lower the value, the further back the report is in the newscast. Again, we average the values on a monthly frequency.

Before proceeding, we test whether our three measures capture meaningful variation in inflation-related attention, that is, whether each of them is systematically related to the inflation rate. To do so, we regress the monthly standardized attention measures on the inflation rate and a set of macroeconomic controls. Here and throughout the paper, these controls include inflation volatility, measured as the 12-month moving standard deviation of the inflation rate, the unemployment rate, log industrial production, and the short-term interest rate measured as the monthly three-month interbank rate for Germany. Since some of these variables are not available for the full sample period, the analysis starts in January 1960.

⁵An example page of the archive output is shown in Figure (A.2) (in German).

Table 1: *Tagesschau* meta data and the inflation rate

	#Broadcasts	Duration	Positioning
Inflation rate	0.245 [0.045***]	0.086 [0.026***]	0.092 [0.025***]
Inflation volatility	0.289 [0.154*]	0.008 [0.121]	-0.107 [0.128]
Unemployment rate	-0.036 [0.012***]	-0.048 [0.015***]	-0.051 [0.015***]
Log industrial production	0.007 [0.002***]	0.007 [0.002***]	0.005 [0.002***]
Interest rate	-0.053 [0.024***]	-0.037 [0.014**]	-0.029 [0.016*]
#Obs.	771	771	771

Notes: The table reports monthly regressions of standardized *Tagesschau* coverage measures on inflation and controls. #Broadcasts denotes the number of inflation-related editions, Duration their average relative length, and Positioning their average placement in the broadcast, with higher values indicating earlier placement. Robust standard errors are reported in brackets. The regressions also include a constant. ***, ** and * denote significance at the 1%, 5% and 10% level.


Table (1) shows that both the intensive margin of coverage, captured by the number of broadcasts, and the extensive margin of coverage, captured by duration and position, are positively related to inflation. Higher inflation is associated with more broadcasts on the topic, longer inflation-related segments, and more prominent placement within the broadcast. We therefore combine all three into a single historical *Tagesschau* Attention index (TA), defined as

$$TA_t = B_t \frac{(D_t + P_t)}{2}, \quad (1)$$

where B_t denotes the number of *Tagesschau* broadcasts with an inflation-related report in month t . D_t is the monthly average relative duration of these reports, and P_t is the monthly average relative position of the report within the broadcast.

Figure (4) shows the resulting index together with the monthly inflation rate. Two periods stand out: the oil crisis of the 1970s and the post-Covid inflation surge. In both cases, inflation seems to have received particularly strong media attention. By contrast, other periods with relatively high inflation, such as the 1980s and 1990s, received less coverage. In these decades, there were even months without any inflation-related report at all. In the 2000s and 2010s, attention in television news fluctuates more strongly and appears to track inflation developments more closely.

Figure 4: *Tagesschau* attention to inflation




Notes: This figure plots our baseline index measuring the inflation-related attention of the *Tagesschau* together with the monthly inflation rate.

To better understand the content of inflation-related reports, we analyse the short descriptions and titles from the archive. We exclude words that seem to be technical *Tagesschau*-specific terms or internal abbreviations. We also exclude city names, since the descriptions often begin with the location of the event and these words do not tell us much about the content of the report. Finally, we exclude the words "inflation" and "prices", since they appear in almost every description or title. Figure (5) shows word clouds with the ten most frequent remaining words for each decade, together with the index.⁶

In the 1950s and 1960s, the most frequent terms are relatively broad and do not point to one dominant inflation topic. For the 1950s, where we identify only few inflation-related editions, frequent terms include names such as Hans-Christoph Seeböhm, Federal Minister of Transportation until 1966, and Hans-Günther Sohl, a leading figure in German heavy industry. In the 1960s, terms such as “EEC”, “farmers”, and “German Trade Union Confederation” suggest a mix of European integration and sectoral discussions about prices and wages, while “Schiller” refers to Karl Schiller, Federal Minister of Economic Affairs from 1966 to 1972.

⁶The original German words were translated into English. The German version is shown in Figure (A.3).

Figure 5: *Tagesschau* attention and content



Notes: The figure shows the ten most frequent words in the titles and short descriptions of inflation-related *Tagesschau* reports by decade. Technical archive terms, city names, and the words “inflation” and “prices” are excluded. The original German words are translated into English. The line shows the historical *Tagesschau* attention index.

The oil-price episodes are clearly visible in the following two decades. In the 1970s, words such as “OPEC”, “USA”, and “France” point to the international dimension of the first oil crisis. In the 1980s, reports are more directly about energy prices, with frequent terms such as “oil prices”, “petrol prices”, and again “OPEC”. The term “decline” also appears, indicating that some reports covered falling prices.

For the 1990s, the most frequent terms do not directly point to German reunification, although reunification was an important driver of inflation in this period. Instead, reporting contains broader political and institutional references, including Chancellor Helmut Kohl, the federal government, and the European Union. In the 2000s, the euro and the ECB appear among the most frequent terms. The *Deutsche Bahn* also appears, likely reflecting reports about train-fare increases.

In the 2010s, inflation-related reporting is closely related to the euro-area crisis, with frequent references to “ECB”, “euro”, “Mario Draghi”, and “Greece”. In the 2020s, coverage of the recent inflation surge centers on “Ukraine” and “energy prices”, consistent with the energy-price shock following Russia’s war against Ukraine. The terms “Bundestag” and “federal government” also appear frequently.

Since we can search for our keywords only in report descriptions, titles, and tags, but not in the full spoken text, our index likely provides a lower bound of inflation-related coverage. At the same time, it captures reports that are clearly related to inflation

or price dynamics.

To further support the index and to make it possible to analyze the spoken content in more detail, we construct a second dataset.

2.2 The video dataset

Our second dataset consists of all available videos of the *Tagesschau* 8 p.m. broadcasts from September 1989 to August 2024. We obtained the videos from the official *Tagesschau* website.⁷ Videos are missing for the period from March 2004 to March 2007. For this period, the *Tagesschau* archive provided us with all segments that match our inflation-related keywords in the archive search. In contrast to the metadata-based archive dataset described above, we received the video files of these segments.⁸

To extract text from the videos, we use the `openai/whisper-large-v3` model (Radford et al., 2023).⁹

The *Tagesschau* archive also provides us with official transcripts for broadcasts from November 2014 to July 2023. We use these transcripts to evaluate the quality of our own transcripts for the same period. For this comparison, we use the `doc2vec` model (Le and Mikolov, 2014), implemented in Gensim (Řehůřek and Sojka, 2010).¹⁰ The comparison shows that the automatic transcription works well.

We then search the transcripts of all broadcasts for the same keywords as in the archive dataset (Table A.1) to identify inflation-related reports. In the archive dataset, the archive search directly identifies the relevant segment within a broadcast. In the video dataset, by contrast, the keyword search first tells us only whether an inflation-related keyword appears somewhere in the full 15-minute broadcast. We therefore manually identify the relevant segments by watching the corresponding videos. In this step, we also correct transcription errors in the inflation-related segments.

Our search yields 1102 inflation-related reports from September 1989 to August 2024. This number is larger than the total number of reports in the archive dataset

⁷Recent videos are available at https://www.tagesschau.de/tagesschau_20_uhr/ts-78220.html. Older videos are available at <https://www.tagesschau.de/inland/tsvorzwanzigjahren-ts-142.html>.


⁸Strictly speaking, for this period we cannot search the full text of all broadcasts. As discussed above, the archive search is based on matches between our keywords and the short descriptions, titles, and tags of archived segments.

⁹We use the automatic speech recognition pipeline from Hugging Face, available at <https://huggingface.co/openai/whisper-large-v3>. The pipeline is optimized with `FlashAttention 2` (Dao, 2023). We process the audio in 30-second chunks, set the language to German, and do not return timestamps.

¹⁰The model is trained with a vector size of 100, a window size of 5, a minimum count of 1, and 20 epochs.

for a much longer horizon. Figure (6) compares the number of inflation-related broadcasts in the video and archive datasets. The two series are highly correlated, with a correlation coefficient of 0.74, but the video-based series is higher in most months. This is expected, since the video dataset allows us to search the full spoken text of each broadcast. It also supports our interpretation of the archive dataset as a lower bound of inflation coverage.

Figure 6: *Tagesschau* attention to inflation: archive and video datasets



Notes: This figure plots the monthly number of inflation-related *Tagesschau* broadcasts identified in the archive dataset and in the video dataset.

Next, we analyze the spoken text in more detail. Beyond the amount of inflation coverage, the direction of reported price changes and the sentiment of the coverage may also matter for households. We therefore classify each inflation-related report along both dimensions.


First, we use the `gpt-5.2` model to classify whether a report mainly discusses rising, falling, or stable prices.¹¹ Second, we use the same model to classify the tone of each report as positive, neutral, or negative. To ensure the quality of these labels, we manually review all classifications and correct them where necessary.

We aggregate both measures to the monthly level. For price direction, we compute net price direction as the number of reports about rising prices minus the number

¹¹We do not distinguish between rising prices and rising inflation. While these concepts are technically different, we found that this distinction is difficult to apply consistently in practice. Some reports discuss price changes without making a clear statement about aggregate inflation. Moreover, a more detailed classification would make the measure less transparent. Since households are also likely to interpret news about prices as relevant for inflation, we use the broader category of price changes.

of reports about falling prices in a given month. For sentiment, we compute net sentiment as the number of positive reports minus the number of negative reports.¹²

Figure 7: Comparison of content measures with the number of videos



Notes: This figure compares the monthly number of inflation-related *Tagesschau* videos with two content-based measures from the video dataset. Panel (a) shows the net price direction, while panel (b) shows net sentiment.

Figure (7) plots the text-based measures together with the number of inflation-related

¹²We do not normalize these measures by the total number of reports in a month. In months with few inflation-related reports, such normalization would create an index that often takes extreme values and would remove much of the variation in the measure.

videos. The number of videos is positively correlated with net price direction with a correlation coefficient of 0.77 and negatively correlated with net sentiment (-0.58). This indicates that periods with more inflation coverage are typically also periods with more reporting about rising prices and a more negative tone.

3 Tagesschau Attention and Inflation: A Threshold Model

We now study how attention to inflation in the *Tagesschau* varies with the inflation rate. Previous work shows that household attention to inflation can change once inflation exceeds certain thresholds (Pfäuti, 2023; Korenok et al., 2026). Since media coverage is a potential channel through which inflation becomes salient to households, we ask whether media attention itself also changes across inflation regimes. For this exercise, we use our historical *Tagesschau* attention index from the archive dataset, which covers the longest period in our data.

We estimate the following threshold model:

$$TA_t = \alpha + \sum_{k=0}^K \beta_{k+1} \pi_t \mathbb{1}(\gamma_k < \pi_t \leq \gamma_{k+1}) + \delta X_t + \varepsilon_t. \quad (2)$$

Here, TA_t is the standardized *Tagesschau* attention index, and π_t is the inflation rate. The thresholds γ_k divide the inflation rate into $K + 1$ regimes, with $\gamma_0 = -\infty$ and $\gamma_{K+1} = +\infty$. The coefficient β_{k+1} measures the slope of inflation in regime $k + 1$. The vector X_t includes the monthly unemployment rate, inflation volatility, log industrial production, and the monthly interest rate. The term α is a regime-independent constant, and ε_t is the error term.

We estimate the model using Stata’s `threshold` model. The model searches over possible threshold values and chooses the ones that minimize the sum of squared residuals. We restrict candidate thresholds to lie between the 10th and 90th percentiles of the inflation distribution. We allow for up to three thresholds. The number of thresholds is chosen by the Bayesian Information Criterion.¹³

Note that the procedure also allows for no threshold. In that case, the model reduces to a linear relationship between inflation and media attention. We obtain confidence intervals for the estimated thresholds from a bootstrap procedure with 2,000 replications.

¹³For details on Stata’s `threshold` command, see <https://www.stata.com/manuals/tsthreshold.pdf>. A discussion of related threshold procedures in inflation applications can be found in Nautz and Scharff (2012).

Since our sample starts in 1956, it covers two different monetary policy regimes. Until December 1998, monetary policy in Germany was conducted by the Deutsche Bundesbank. Since January 1999, it has been conducted by the ECB. This institutional change may matter for inflation coverage in the *Tagesschau*. Under the ECB regime, inflation is evaluated against an explicit inflation target, which may make deviations from price stability more salient for news coverage.

To study this, we estimate Equation (2) for three samples: the full sample from January 1960 to July 2024, the Bundesbank period from January 1960 to December 1998, and the ECB period from January 1999 to July 2024.

Table 2: Threshold estimates

	Full Sample		Bundesbank		ECB	
	January 1960-July 2024		January 1960-December 1998		January 1999-July 2024	
Thresholds (γ_1)	5.395 (5.014–5.776)	5.395 (5.010–5.780)	5.714 (5.196–6.233)	5.714 (5.238–6.191)	2.839 (2.025–3.653)	2.132 (1.337–2.927)
Region 1 ($\pi_t \leq \gamma_1$)						
Inflation rate	0.043 [0.022*]	0.109 [0.036***]	0.065 [0.022***]	0.032 [0.029]	0.122 [0.069*]	0.130 [0.089]
Region 2 ($\gamma_1 < \pi_t$)						
Inflation rate	0.223 [0.035***]	0.259 [0.048***]	0.211 [0.036***]	0.190 [0.040***]	0.391 [0.077***]	0.391 [0.099***]
Inflation volatility		0.267 [0.160*]		−0.270 [0.158*]		0.142 [0.293]
Unemployment rate		−0.035 [0.012***]		−0.090 [0.017***]		0.025 [0.042]
Log industrial production		0.006 [0.001***]		0.009 [0.003***]		−0.008 [0.008]
Interest rate		−0.057 [0.024**]		−0.030 [0.023]		−0.055 [0.058]
#Obs.	774	774	468	468	306	306

Notes: The table reports estimates of Equation (2) for the full sample, the Bundesbank period, and the ECB period. The dependent variable is the standardized *Tagesschau* attention index. We allow for up to three thresholds and use a trimming parameter of 10%. The 95% confidence intervals for the estimated thresholds are obtained from a bootstrap procedure with 2,000 replications and are reported in parentheses. Robust standard errors are reported in brackets. All specifications include a regime-independent constant. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table (2) reports the results for all three samples, both with and without control variables. Figures A.4 and A.5 show the corresponding threshold estimates graphically without controls. Across all specifications, the model selects one threshold. For the full sample, the estimated threshold is about 5.4%. The Bundesbank-period threshold is slightly higher, at about 5.7%. In the ECB period, the threshold is much lower: 2.8% without controls and 2.1% with controls.

The slope estimates show that the response of inflation coverage is concentrated

above these thresholds. Below the threshold, the inflation coefficients are smaller and lose statistical significance in the Bundesbank and ECB samples once controls are included. Above the threshold, the coefficients are larger and statistically significant throughout. In the full sample, the above-threshold coefficient is 0.223 without controls and 0.259 with controls. Hence, a one-percentage-point increase in inflation above the threshold is associated with an increase in the standardized attention index of about 0.22 to 0.26 standard deviations. In the Bundesbank period, the corresponding coefficients are 0.211 and 0.190, respectively. In the ECB period, the coefficient is 0.391 in both specifications. Overall, the response of inflation coverage above the estimated threshold is roughly twice as large in the ECB period as in the Bundesbank period.

The macroeconomic controls do not change the main pattern. In the full sample and in the Bundesbank period, unemployment is negatively related to inflation coverage, while industrial production is positively related to it. The unemployment coefficient is larger in absolute value during the Bundesbank period, at -0.090 , compared with -0.035 in the full sample. The coefficient on industrial production is also somewhat larger in the Bundesbank period, at 0.009 , compared with 0.006 in the full sample. The other controls show less stable patterns. During the ECB period, none of the macroeconomic controls is statistically significant.

The lower threshold in the ECB period is economically plausible. It is close to the ECB's inflation target, which may have given TV editors a clear benchmark for assessing when inflation developments become newsworthy. This interpretation, however, should be treated with caution. The ECB subsample is shorter, and average inflation is lower than in the Bundesbank period, at 1.81% compared with 3.18%. This lower-inflation environment may mechanically push the estimated threshold down.

At the same time, the descriptive pattern in Figure (4) is consistent with this interpretation. In the ECB period, months without any inflation-related report are less frequent, and coverage appears to respond already at lower inflation rates. The standardized attention index also shows greater within-period variation in the ECB period, with a standard deviation of 1.19, compared with 0.85 in the Bundesbank period.

Overall, the results show that *Tagesschau* coverage reacts nonlinearly to inflation. If television news is an important channel through which inflation news reaches households, this pattern could help explain nonlinearities in household attention to inflation. In the next section, we therefore study whether *Tagesschau* coverage is related to household attention, inflation perceptions, and inflation expectations.

4 Transmission to households

In this section, we study whether inflation-related TV coverage is transmitted to households. We focus on three outcomes: household attention to inflation, inflation perceptions, and inflation expectations. We first analyze these outcomes at the aggregate monthly level and then turn to household-level evidence. We use two household datasets. For the aggregate time-series analysis, we use the German part of the *Business and Consumer Survey* (BCS) of the European Commission. For the micro-level analysis, we use the Bundesbank Online Panel Households (BOP-HH). The BCS is the longest-running household survey for Germany. It is conducted monthly in EU member states and candidate countries. We use the German data starting in 1985. The survey is usually fielded during the first three weeks of the month and includes about 2,000 households. German consumers are asked qualitative questions about how consumer prices changed over the past 12 months (question Q5) and how they expect prices to change over the next 12 months relative to the past 12 months (question Q6).¹⁴

The survey provides monthly balance indices for both questions.¹⁵ We use these indices as measures of household inflation perceptions and inflation expectations. To measure household attention to inflation, we follow Bracha and Tang (2025). Their first measure of inattention is the share of households who answer “don’t know” when asked about inflation perceptions over the past 12 months. We construct this measure from the “don’t know” share in question Q5 of the BCS. Since we interpret inattention as the counterpart of attention, we define household attention as one minus this share.

We use the BOP-HH as a complementary dataset for the micro-level analysis. One advantage of this survey is that it asks households directly for point estimates of inflation. The BOP-HH started with three pilot waves in April, May, and June 2019 and has been conducted monthly since April 2020. It is representative of the German online population aged 16 or older and is usually fielded during the last two weeks of the month.¹⁶

The survey asks households for a point estimate of inflation over the previous 12

¹⁴For the exact wording of the questions and a general overview of the survey, see https://economy-finance.ec.europa.eu/document/download/426aefda-4888-42e7-ac02-1ac85f979b3d_en?filename=bcs_user_guide_Mar_26.pdf.

¹⁵The balance index is defined as $Balance = PP + \frac{P}{2} - \frac{M}{2} - MM$, where PP is the share of households answering “risen a lot” for past prices or “increase more rapidly” for expected prices. Correspondingly, P refers to “risen moderately” or “increase at the same rate”, M to “stayed about the same” or “stay about the same”, and MM to “fallen” or “fall”.

¹⁶For more information about the survey, see <https://www.bundesbank.de/en/bundesbank/research/survey-on-consumer-expectations-bop-hh/survey-on-consumer-expectations-794568>.

months (question CQ002). This allows us to construct the second inattention measure proposed by Bracha and Tang (2025): the gap between perceived inflation and actual inflation. We compute the absolute value of this gap and interpret it as a measure of inattention. To keep the interpretation consistent with our BCS attention measure, we multiply it by -1 , so that higher values indicate more attention. For inflation expectations, we use households' point estimate of inflation over the next 12 months (question CM003).

Before turning to the empirical analysis, we first provide a simple model to show how TV coverage may influence household attention and beliefs.

4.1 A simple model

To guide the empirical analysis, we develop a simple model that links television news coverage to household information and belief formation. In contrast to models in which household attention is an active and endogenous information choice (e.g. as in Pfäuti (2023); Bracha and Tang (2025); Weber et al. (2025)), we treat households' exposure to inflation news as passive and let the informativeness of the signal vary with editorial TV coverage. This is motivated by the particular nature of TV news and the *Tagesschau* in particular: it is a widely watched daily evening news program with a fixed schedule and low marginal attention cost. Therefore in our model variation in household information arises not from optimizing behavior of the households, but from editorial decisions of the *Tagesschau*.

Let π_t denote current inflation. Households do not observe π_t perfectly. Instead, they receive information about inflation through the main evening TV news. They observe a noisy TV signal

$$s_t = \pi_t + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_t^2), \quad (3)$$

where the signal noise variance σ_t^2 depends on editorial TV attention a_t^{TV} . This captures the intensity of editorial coverage devoted to inflation. Higher TV attention makes the signal more precise:

$$\sigma_t^2 = \frac{\bar{\sigma}^2}{1 + \kappa a_t^{TV}}, \quad \kappa > 0. \quad (4)$$

Here, $\bar{\sigma}^2$ denotes baseline noise in the TV signal, and κ measures how strongly additional TV coverage reduces this noise. When a_t^{TV} is higher, σ_t^2 is lower, so the signal is more precise.

Motivated by our findings above, we allow TV attention to respond nonlinearly to

inflation. In particular, editorial attention rises once inflation exceeds a threshold level π^* :

$$a_t^{TV} = a_L^{TV} + a_H^{TV} \mathbf{1}\{\pi_t \geq \pi^*\}, \quad a_H^{TV} > 0. \quad (5)$$

Here, a_L^{TV} is baseline TV attention, while a_H^{TV} is the additional attention when inflation exceeds the threshold. When inflation is high enough, TV attention increases by a_H^{TV} . Inflation then receives more prominent news coverage and the TV signal becomes more informative.

Before observing the signal, households hold prior beliefs about current inflation,

$$\pi_t \sim \mathcal{N}(\bar{\pi}, V_t), \quad (6)$$

where $\bar{\pi}$ is the prior mean and V_t denotes prior uncertainty. Taking V_t as given, households update their beliefs after observing s_t according to Bayes' rule:

$$\hat{\pi}_t = (1 - a_t^{HH})\bar{\pi} + a_t^{HH} s_t, \quad a_t^{HH} = \frac{V_t}{V_t + \sigma_t^2}, \quad (7)$$

where a_t^{HH} denotes household attention, that is, the weight households place on the TV signal relative to their prior. Substituting (4) into the expression for a_t^{HH} yields

$$a_t^{HH} = \frac{V_t(1 + \kappa a_t^{TV})}{V_t(1 + \kappa a_t^{TV}) + \bar{\sigma}^2}, \quad (8)$$

which is increasing in TV attention a_t^{TV} . Greater editorial attention therefore increases the precision of the TV signal and raises household attention in the sense that households place more weight on current inflation news.

To link perceptions to inflation expectations, we assume that households perceive inflation to follow a stationary AR(1) process,

$$\pi_{t+1} = (1 - \rho)\bar{\pi} + \rho\pi_t + u_{t+1}, \quad \rho \in (0, 1), \quad (9)$$

with $u_{t+1} \sim \mathcal{N}(0, \sigma_u^2)$ and $\mathbb{E}_t[u_{t+1}] = 0$. One-period-ahead inflation expectations are then

$$\mathbb{E}_t[\pi_{t+1}] = (1 - \rho)\bar{\pi} + \rho\hat{\pi}_t. \quad (10)$$

Substituting (7) into (10) gives

$$\mathbb{E}_t[\pi_{t+1}] = \bar{\pi} + \rho a_t^{HH} (s_t - \bar{\pi}). \quad (11)$$

Inflation expectations therefore respond more strongly to inflation news when house-

hold attention a_t^{HH} is higher.

The model yields a simple prediction. Once inflation is high enough to trigger greater editorial attention in TV news, the inflation signal becomes more precise. Households then place more weight on inflation news when updating their beliefs. As a result, inflation expectations become more responsive to incoming information.

In the following sections, we use these predictions to guide the empirical analysis. We first estimate reduced-form specifications that relate household attention, perceptions, and expectations to inflation-related TV coverage, conditional on inflation and other macroeconomic controls. We then turn to instrumental-variable specifications and micro-level evidence.

4.2 Simple regressions

We now estimate whether inflation-related TV news is associated with household attention, inflation perceptions, and inflation expectations, conditional on the inflation environment. Specifically, we estimate

$$HH_t = \alpha + \beta TV_t + \gamma \pi_t + \delta X_t + \varepsilon_t, \quad (12)$$

where HH_t denotes the household variable from the BCS and TV_t denotes inflation-related *Tagesschau* coverage. We use two measures of TV coverage: our historical *Tagesschau* attention index and the number of inflation-related videos from the video dataset. The vector X_t includes inflation volatility, the unemployment rate, log industrial production, and the interest rate. The household variables and the TV measures are standardized. Specifications using the historical *Tagesschau* index start in January 1985. Specifications using the number of videos start in January 1989. Table (3) reports the results. Across specifications, TV coverage is positively related to household attention and inflation perceptions. A one-standard-deviation increase in the *Tagesschau* attention index is associated with an increase in household attention of 0.132 standard deviations; using the number of videos, the coefficient is 0.221. For inflation perceptions, the corresponding coefficients are 0.072 and 0.125. For inflation expectations, the results differ across the two TV measures: the coefficient is insignificant for the historical attention index and positive for the number of videos. The macroeconomic controls confirm the strong role of the inflation environment. Inflation is positively related to all three household measures. Other controls also matter in several specifications, but they do not change the positive association between *Tagesschau* coverage and household attention or perceptions.

Table 3: *Tagesschau* coverage and household attention, perceptions, and expectations

<i>Dependent variable:</i>	<i>Tagesschau</i> attention index			Number of videos		
	Attention	Perceptions	Expectations	Attention	Perceptions	Expectations
<i>Tagesschau</i> measure	0.132 [0.034***]	0.072 [0.033**]	−0.047 [0.048]	0.221 [0.033***]	0.125 [0.028***]	0.134 [0.047***]
Inflation	0.116 [0.023***]	0.395 [0.022***]	0.323 [0.049***]	0.100 [0.026***]	0.355 [0.023***]	0.209 [0.043***]
Inflation volatility	0.235 [0.096**]	0.086 [0.085]	−0.744 [0.128***]	0.279 [0.099***]	0.127 [0.093]	−0.750 [0.132***]
Interest rate	0.113 [0.018***]	0.104 [0.015***]	0.151 [0.026***]	0.127 [0.017***]	0.098 [0.016***]	0.147 [0.022***]
Unemployment rate	0.228 [0.018***]	0.139 [0.014***]	0.019 [0.019]	0.250 [0.020***]	0.122 [0.015***]	−0.041 [0.020**]
Log industrial production	−0.004 [0.003]	0.035 [0.002***]	0.004 [0.003]	0.004 [0.003]	0.031 [0.003***]	−0.012 [0.004***]
# Obs.	471	471	471	419	419	419

Notes: The table reports regressions of household attention, inflation perceptions, and inflation expectations on inflation-related *Tagesschau* coverage, as in Equation (12). The first three columns use the standardized *Tagesschau* attention index, while the last three columns use the standardized number of inflation-related *Tagesschau* videos. All dependent variables and *Tagesschau* measures are standardized. The regressions include a constant. Robust standard errors are reported in brackets. ***, ** and * denote significance at the 1%, 5% and 10% level.

Table (4) keeps the same specification but replaces the coverage measure with our text-based measures of price direction and sentiment. Net price direction is positively related to all three household outcomes. A one-standard-deviation increase in net price direction is associated with an increase of 0.180 standard deviations in household attention, 0.089 in inflation perceptions, and 0.191 in inflation expectations. Thus, months in which inflation coverage focuses more on rising prices are also months with higher household attention, perceptions, and expectations.

The sentiment results are also consistent with a role of news content. Since net sentiment is defined as positive minus negative reports, lower values indicate a more negative tone. The coefficients on net sentiment are negative for all three household outcomes and statistically significant for attention and expectations. The coefficient for inflation perceptions is small and not statistically significant.

Overall, the simple regressions show that both the amount and the content of inflation-related TV coverage are related to household outcomes. These regressions, however, are conditional correlations and should not be interpreted causally. Editors may also cover inflation more when households already pay more attention to it. We therefore turn to an IV local-projections approach in the next section.

Table 4: *Tagesschau* price direction, sentiment, and household attention, perceptions, and expectations

<i>Dependent variable:</i>	Net price direction			Net sentiment		
	Attention	Perceptions	Expectations	Attention	Perceptions	Expectations
<i>Tagesschau</i> measure	0.180 [0.038***]	0.089 [0.030***]	0.191 [0.038***]	-0.090 [0.036**]	-0.024 [0.031]	-0.177 [0.036***]
Inflation	0.112 [0.024***]	0.363 [0.023***]	0.185 [0.044***]	0.141 [0.024***]	0.384 [0.023***]	0.192 [0.045***]
Inflation volatility	0.339 [0.100***]	0.160 [0.088*]	-0.713 [0.133***]	0.353 [0.099***]	0.163 [0.089*]	-0.681 [0.129***]
Interest rate	0.114 [0.018***]	0.092 [0.016***]	0.145 [0.022***]	0.110 [0.018***]	0.089 [0.016***]	0.143 [0.022***]
Unemployment rate	0.253 [0.020***]	0.120 [0.015***]	-0.053 [0.020***]	0.259 [0.021***]	0.122 [0.015***]	-0.044 [0.020**]
Log industrial production	0.003 [0.003]	0.030 [0.003***]	-0.014 [0.004***]	0.003 [0.003]	0.029 [0.003***]	-0.013 [0.004***]
# Obs.	415	415	415	415	415	415

Notes: The table reports regressions of household attention, inflation perceptions, and inflation expectations on inflation-related *Tagesschau* coverage, as in Equation (12). The first three columns use the standardized net price direction measure, while the last three columns use the standardized net sentiment measure. All dependent variables and *Tagesschau* measures are standardized. The regressions include a constant. Robust standard errors are reported in brackets. ***, ** and * denote significance at the 1%, 5% and 10% level.

4.3 A monthly LP-IV approach

We now use a local-projections instrumental-variables approach (LP-IV) to estimate the dynamic effects of inflation-related *Tagesschau* coverage on household outcomes. We use two measures of monthly inflation-related *Tagesschau* coverage as endogenous variables: the standardized historical *Tagesschau* attention index and the standardized number of inflation-related videos from the video dataset.

Our instrument is the interaction between the inflation surprise and an indicator for *Tagesschau* coverage of the preliminary inflation release. This approach is related to Binder et al. (2025), who use daily CPI surprises and CPI release dates as instruments for inflation-related media coverage. We adapt this idea to the monthly frequency and focus identification on unexpected inflation news that is actually reported by the *Tagesschau*, rather than on release events or inflation surprises more broadly.

We define the inflation surprise as

$$Z_t = \pi_t - \mathbb{E}_{t-1}[\pi_t], \quad (13)$$

where π_t is realized inflation and $\mathbb{E}_{t-1}[\pi_t]$ is expected inflation formed before the release. We construct $\mathbb{E}_{t-1}[\pi_t]$ using a set of univariate and multivariate forecasting

models.¹⁷ We use the signed surprise rather than its absolute value. The reason is that the direction of the surprise is likely to matter for media attention. Our earlier results show that *Tagesschau* attention increases mainly in periods of higher inflation, which suggests that inflation news is not symmetric in sign. Using the signed surprise allows positive and negative inflation news to have different effects on media coverage.

In Germany, the Federal Statistical Office publishes a preliminary inflation release near the end of each month. The final figures usually follow in the first half of the subsequent month. We focus on *Tagesschau* coverage of the preliminary release because it is the first official publication of monthly inflation and is therefore likely to be the more informative news event.

Let R_t be a dummy equal to one if the *Tagesschau* covered the preliminary inflation release in month t . The first stage is

$$X_t = c + \alpha_1(Z_t \times R_t) + \alpha_2 R_t + \sum_{j=1}^p \Phi_j \mathbf{C}_{t-j} + \varepsilon_t, \quad (14)$$

where X_t denotes monthly inflation-related *Tagesschau* coverage, measured either by the historical attention index or by the number of inflation-related videos, and $\mathbf{C}_{t-j}$ includes lagged controls. We use the longest available sample for each specification. Specifications based on the historical *Tagesschau* attention index start in January 1985, when the BCS household data become available. Specifications based on the video dataset start in September 1989, when the video data begin. Both samples end in July 2024.

The coefficient on the interaction term is positive and statistically significant in both first stages. The Kleibergen–Paap rk Wald F -statistics are 11.9 when using the *Tagesschau* attention index and 11.1 when using the number of inflation-related videos. Both values are above the common rule-of-thumb threshold of 10, indicating moderate instrument strength.

For each horizon $h = 0, 1, \dots, H$, we then estimate

$$Y_{t+h} = \beta_0^h + \beta_1^h \widehat{X}_t + \sum_{j=1}^p \Gamma_j^h \mathbf{C}_{t-j} + u_{t+h}, \quad (15)$$

where Y_t is alternatively household attention, inflation perceptions, or inflation expectations, and $\widehat{X}_t$ is fitted monthly *Tagesschau* coverage from the first stage. The control vector includes 12 lags of the outcome variable, lags of the media measure, and

¹⁷See Appendix C for the procedure.

lags of the unemployment rate, inflation, log industrial production, and the interest rate. Estimating Equation (15) horizon by horizon yields the impulse responses $\{\beta_1^h\}_{h=0}^H$. We estimate the model up to $h = 24$ months. Inference is based on heteroskedasticity- and autocorrelation-robust standard errors.

A key identification concern is that the *Tagesschau* may cover preliminary releases mainly in months with high or surprising inflation. We therefore test whether the release-coverage dummy is related to inflation, inflation surprises, or other macroeconomic variables. To do so, we regress the release-coverage dummy on inflation, the inflation surprise, inflation volatility, the interest rate, the unemployment rate, and log industrial production. We estimate both linear probability models and probit specifications. We report this validity check both for the full archive sample starting in 1960 and for the post-1989 sample used for the video-based specifications.¹⁸

Table 5: Validity check for *Tagesschau* coverage of preliminary inflation releases

	January 1960-July 2024		January 1989-July 2024	
	OLS	Probit	OLS	Probit
Inflation	0.007 [0.009]	0.074 [0.064]	0.024 [0.017]	0.113 [0.072]
Inflation surprise	0.017 [0.035]	0.154 [0.212]	0.032 [0.048]	0.127 [0.215]
Inflation volatility	0.048 [0.047]	0.443 [0.268*]	-0.027 [0.069]	0.068 [0.322]
Interest rate	0.009 [0.006]	0.107 [0.042**]	0.002 [0.012]	0.011 [0.049]
Unemployment rate	0.020 [0.004***]	0.188 [0.028***]	-0.012 [0.009]	-0.007 [0.040]
Log industrial production	0.000 [0.001]	0.014 [0.005***]	-0.008 [0.002***]	-0.034 [0.010***]
# Obs.	696	696	419	419

Notes: The dependent variable is a dummy equal to one if the *Tagesschau* covered the preliminary inflation release in month t . Columns (1) and (3) report linear probability models, while columns (2) and (4) report probit estimates. All specifications include inflation, the inflation surprise, inflation volatility, the interest rate, the unemployment rate, and log industrial production and a constant. Robust standard errors are reported in brackets. ***, ** and * denote significance at the 1%, 5% and 10% level.

Table (5) shows no significant relation between release-related *Tagesschau* coverage and either inflation or the inflation surprise, in both samples and across both

¹⁸For the full sample, the release-coverage dummy is constructed from the archive data. This classification is less precise because the archive texts are short and do not contain the full report. For the post-1989 period, we instead use the video data and manually identify *Tagesschau* reports covering the preliminary inflation release. Since the full text is available in this case, the classification is more transparent and more precise.

specifications. This suggests that coverage of preliminary inflation releases is not mechanically driven by the level of inflation or by the size of the surprise.

We also account for the timing of the household survey. Our measure of household attention is based on the BCS, which is usually fielded during the first three weeks of a month, while preliminary inflation releases typically occur in the last week of the month. We therefore shift the household outcome forward by one month so that the preliminary release falls into the relevant survey window.


Figure (8) plots the estimated β_1^h coefficients. A shock to inflation-related *Tagesschau* attention raises all three household outcomes. Household attention increases almost on impact and remains elevated over the full 24-month horizon. The response reaches about one standard deviation and is slightly larger when we use the historical *Tagesschau* attention index.

Inflation perceptions show the largest response. They increase by about 0.5 standard deviations on impact and reach a maximum of about 1.7 standard deviations when we use the historical attention index. The effect remains statistically significant for up to 15 months. When we use the number of videos instead, the response is smaller but still substantial, with an immediate increase and a maximum effect of about one standard deviation.

Inflation expectations also rise after the shock, but the response is less persistent than for attention and perceptions. When we use the historical attention index, the response is smaller than the response of inflation perceptions. When we use the number of videos, expectations increase by up to about one standard deviation. In both specifications, the response becomes statistically insignificant after about 7 to 10 months.

The stronger response of perceptions is consistent with the simple model above. Preliminary-release reports contain direct information about current inflation, which is closely related to the perception question. Expectations require an additional step: households must translate news about current inflation into beliefs about future inflation. This may explain why expectations react less persistently.

Figure 8: Impulse responses of household attention, inflation perceptions, and inflation expectations to *Tagesschau* attention shocks




Notes: This figure shows local-projection IV estimates of the responses of household attention, inflation perceptions, and inflation expectations to a one-standard-deviation shock in *Tagesschau* attention to inflation, as in Equation (15). Panel (a) uses the standardized *Tagesschau* attention index, while panel (b) uses the standardized number of inflation-related *Tagesschau* videos from our video dataset. The sample starts in January 1985 for the attention-index specifications and in September 1989 for the video specifications. Both samples end in July 2024. Shaded areas indicate 68% and 90% confidence intervals based on heteroskedasticity- and autocorrelation-robust standard errors.

Figure (9) reports the corresponding results for the text-based measures. These specifications replace the TV attention measure with net price direction or net sentiment. The responses to net price direction are similar to the baseline results, which is consistent with the positive correlation between price-direction news and the amount of inflation coverage. Household attention reacts persistently, inflation perceptions show the largest response, and inflation expectations respond less persistently. The magnitudes are also comparable to the baseline estimates.

The results for net sentiment move in the opposite direction. Since net sentiment

is defined as positive minus negative reports, a positive shock corresponds to more positive, or less negative, inflation coverage. Such a shock lowers household attention, inflation perceptions, and inflation expectations.

Figure 9: Impulse responses of household attention, perceptions, and expectations to directional and sentiment-based *Tagesschau* shocks



Notes: The figure shows local-projection IV estimates of the responses of household attention, inflation perceptions, and inflation expectations to a one-standard-deviation shock in inflation-related *Tagesschau* coverage. Panel (a) uses the standardized net price direction measure, while panel (b) uses the standardized net sentiment measure. The sample starts in September 1989 and ends in July 2024. Shaded areas indicate 68% and 90% confidence intervals based on heteroskedasticity- and autocorrelation-robust standard errors.

Overall, the LP-IV results suggest that inflation-related *Tagesschau* coverage affects all three household outcomes. The responses are largest for perceptions, most persistent for attention, and more short-lived for expectations.

4.4 Effects at the micro level

We now turn to household-level evidence on the link between inflation-related TV news and household outcomes. Using the BOP-HH variables defined in Section 4, we ask whether households become more attentive to inflation, adjust their perceptions of current inflation, or revise their inflation expectations when inflation-related news appeared in the *Tagesschau* on the day before the household participated in the survey.

Our specification is

$$Y_{it} = \alpha_i + \beta \text{TVNews}_{t-1} + \delta_w + \eta_d + \varepsilon_{it}, \quad (16)$$

where Y_{it} is the household-level outcome of interest for household i on survey day t . Depending on the specification, Y_{it} is household attention, inflation perceptions, or inflation expectations. TVNews_{t-1} is a dummy equal to one if there was inflation-related *Tagesschau* coverage on the previous day. We estimate separate specifications for different types of inflation-related coverage. α_i denotes household fixed effects, δ_w week fixed effects, and η_d day-of-week fixed effects. Standard errors are clustered at the household level.

We consider three news indicators. The first two indicate whether any inflation-related *Tagesschau* coverage appeared on the previous day, based either on the archive dataset or on the video dataset. The third indicator equals one if the previous day included *Tagesschau* coverage of the preliminary inflation release. We focus on preliminary releases because they are the first official publication of monthly inflation and should therefore contain more new information than final releases. Final inflation releases usually occur in the first half of the month, when the BOP-HH is typically not yet in the field, leaving too few survey days for a separate analysis.

In the full BOP-HH sample up to wave 54 in June 2024, we observe 791 survey days. According to the archive dataset, 96 of these days followed an inflation-related *Tagesschau* report; according to the video dataset, 147 did. Eleven survey days followed *Tagesschau* coverage of the preliminary inflation release.¹⁹

¹⁹The estimation sample differs across outcomes because the inflation-perception question was introduced only in wave 8. Household attention and inflation perceptions are therefore estimated from wave 8 onward. In this sample, we observe 679 survey days. On 91 of these survey days, there was inflation-related *Tagesschau* coverage on the previous day according to the archive dataset. On 138 days, there was such coverage according to the video dataset. On 10 survey days, the previous day included *Tagesschau* coverage of the preliminary inflation release. Inflation expectations are available from wave 1 onward and are therefore estimated on the full BOP-HH sample. The expectation results are very similar when we restrict the sample to wave 8 onward.

Table 6: Household attention and *Tagesschau* coverage at the micro level

	Archive	Video	Preliminary release
TVNews _{t-1}	0.019 [0.007***]	0.021 [0.006***]	0.114 [0.037***]
Household FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
Day-of-week FE	Yes	Yes	Yes
# Obs.	115,701	115,701	115,701

Notes: The dependent variable in all specifications is the standardized household attention measure from the BOP-HH as described in the text. The first two columns use alternative dummy variables for inflation-related *Tagesschau* coverage on the previous day, based on the archive and video datasets, respectively. The third column uses a dummy variable indicating whether the previous day included a *Tagesschau* edition covering the preliminary inflation release. Standard errors clustered at the household level are reported in brackets. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table (6) reports the results for household attention. Inflation-related *Tagesschau* coverage is associated with significantly higher attention in both specifications using general inflation-related coverage: attention increases by 0.019 standard deviations using the archive dummy and by 0.021 standard deviations using the video dummy. Coverage of preliminary inflation releases is associated with a larger increase, with a coefficient of 0.114. This suggests that timely reports about newly released inflation figures may be especially relevant for household attention.

Table 7: Inflation perceptions and *Tagesschau* coverage at the micro level

	Archive	Video	Preliminary release
TVNews _{t-1}	0.001 [0.009]	-0.010 [0.008]	-0.114 [0.036***]
Household FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
Day-of-week FE	Yes	Yes	Yes
# Obs.	115,701	115,701	115,701

Notes: The dependent variable in all specifications is the standardized household inflation perception measure from the BOP-HH. The first two columns use alternative dummy variables for inflation-related *Tagesschau* coverage on the previous day, based on the archive and video datasets, respectively. The third column uses a dummy variable indicating whether the previous day included a *Tagesschau* edition covering the preliminary inflation release. Standard errors clustered at the household level are reported in brackets. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table (7) turns to inflation perceptions. The coefficients on the archive and video measures are close to zero and statistically insignificant. Coverage of preliminary

inflation releases, however, is associated with significantly lower perceived inflation, with a coefficient of -0.114 .

Table 8: Inflation expectations and *Tagesschau* coverage at the micro level

	Archive	Video	Preliminary release
TVNews $_{t-1}$	-0.008 [0.007]	-0.006 [0.006]	-0.030 [0.036]
Household FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
Day-of-week FE	Yes	Yes	Yes
# Obs.	187,953	187,953	187,953

Notes: The dependent variable in all specifications is the standardized household inflation expectation measure from the BOP-HH. The first two columns use alternative dummy variables for inflation-related *Tagesschau* coverage on the previous day, based on the archive and video datasets, respectively. The third column uses a dummy variable indicating whether the previous day included a *Tagesschau* edition covering the preliminary inflation release. Standard errors clustered at the household level are reported in brackets. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table (8) shows little evidence of an immediate effect on one-year-ahead inflation expectations. The coefficients are statistically insignificant across all three specifications. Thus, while inflation-related TV coverage is associated with higher attention, it does not translate into a robust contemporaneous adjustment of reported inflation expectations.

The weaker average response of perceptions and expectations may reflect two features of the micro setting. First, the daily regressions use individual broadcasts rather than the aggregate monthly news environment. Reports about rising and falling prices may therefore offset each other. Second, the BOP-HH sample covers a period of relatively high inflation. Average inflation between April 2020 and June 2024 was about 4.5%. Households may therefore already have been highly attentive to inflation and may have held elevated inflation priors. This is consistent with the simple model above: in Equation (7), belief updating depends on the difference between the signal and the prior. In case of an already elevated prior, additional reports about rising prices may contain little new information, while reports about falling prices may be more informative.

To examine this possibility, we estimate a second specification in which the single news indicator is replaced by two dummies, one for reports about rising prices and one for reports about falling prices. In the full sample, reports about rising prices are more frequent than reports about falling prices, appearing before 81 and 27 survey days, respectively.

Table 9: Inflation expectations perceptions, and price-direction news at the micro level

	Inflation expectations	Inflation perceptions
Rising prices	−0.000 [0.008]	−0.011 [0.009]
Falling prices	−0.022 [0.011**]	−0.032 [0.013**]
Household FE	Yes	Yes
Week FE	Yes	Yes
Day-of-week FE	Yes	Yes
# Obs.	187,953	115,701

Notes: The dependent variables are the standardized household inflation expectation measure from the BOP-HH in the first column and the standardized household inflation perception measure from the BOP-HH in the second column. The price-direction news dummies indicate whether the previous day included a *Tagesschau* edition covering rising or falling prices, respectively. Standard errors clustered at the household level are reported in brackets. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table (9) reports the corresponding results for inflation perceptions and expectations. We find no evidence that reports about rising prices affect either outcome: the coefficients are close to zero and statistically insignificant in both cases. By contrast, reports about falling prices are associated with significantly lower inflation perceptions and expectations. The estimated coefficients are -0.032 for perceptions and -0.022 for expectations.

These results suggest that household belief responses depend on both news content and the inflation environment. During the recent high-inflation episode, reports about falling prices appear to have contained more new information for households than additional reports about rising prices.

5 Conclusion

This paper provides new long-run evidence on television news as a channel through which inflation information reaches households. We study the *Tagesschau*, Germany’s main evening news program, and construct two datasets on its inflation coverage: a historical archive dataset from 1956 to 2024 and a video-based dataset from 1989 to 2024. These data allow us to measure not only how much inflation is covered, but also how prominently it is covered and how it is described.

We show that inflation coverage in the *Tagesschau* is highly state-dependent. Inflation receives little coverage in low-inflation periods, but appears more often, receives more

time, and is placed earlier in the broadcast when inflation is high. This relationship is nonlinear. The estimated threshold is around 5 to 6 percent in the full sample and in the Bundesbank period, but much lower in the ECB period, where it lies close to the ECB's inflation target. This points to greater sensitivity of inflation coverage in the more recent monetary policy regime.

We also show that variation in this coverage is reflected in household outcomes. Months with more inflation-related *Tagesschau* coverage are months with higher household attention and higher inflation perceptions. The IV-LP estimates suggest persistent and economically sizable effects on household attention, inflation perceptions, and inflation expectations. The household-level evidence points in the same direction: inflation-related news on the previous day raises household attention, and belief responses depend on whether the news reports rising or falling prices.

Our results make one important channel of household inflation attention observable. Households may become attentive to inflation through many routes, including their own shopping experiences, conversations, official statistics, or news coverage. We show that television news is one such channel and that it is linked to household attention and beliefs. The *Tagesschau* is therefore more than a record of inflation developments; it is part of the information environment households use to understand inflation. This provides one explanation for why attention, perceptions, and expectations change across inflation environments.

References

- Arndt, S. (2024). Different newspapers-Different inflation perceptions, *AWI Discussion Paper Series* (748).
- Berger, H., Ehrmann, M. and Fratzscher, M. (2011). Monetary policy in the media, *Journal of Money, Credit and Banking* **43**(4): 689–709.
- Binder, C., Frank, P. and Ryngaert, J. M. (2025). The causal effect of news on inflation expectations, *NBER Working paper series* (34088).
- Bracha, A. and Tang, J. (2025). Inflation levels and (in)attention, *Review of Economic Studies* **92**(3): 1564–1594.
- Buelens, C. (2023). Googling "inflation": What does internet search behaviour reveal about household (in)attention to inflation & monetary policy?, *Publications Office of the European Union*.
- Carroll, C. D. (2003). Macroeconomic expectations of households and professional forecasters, *the Quarterly Journal of economics* **118**(1): 269–298.
- Cavallo, A., Cruces, G. and Perez-Truglia, R. (2017). Inflation expectations, learning, and supermarket prices: Evidence from survey experiments, *American Economic Journal: Macroeconomics* **9**(3): 1–35.
- Chahrour, R., Shapiro, A. H. and Wilson, D. (2025). News selection and household inflation expectations, *NBER working paper series* (33837).
- Coibion, O. and Gorodnichenko, Y. (2012). What can survey forecasts tell us about information rigidities?, *Journal of Political Economy* **120**(1): 116–159.
- Coibion, O., Gorodnichenko, Y. and Kumar, S. (2018). How do firms form their expectations? New survey evidence, *American Economic Review* **108**(9): 2671–2713.
- Dao, T. (2023). Flashattention-2: Faster attention with better parallelism and work partitioning, *arXiv preprint arXiv:2307.08691*.
- Dräger, L. (2015). Inflation perceptions and expectations in Sweden – Are media reports the missing link?, *Oxford Bulletin of Economics and Statistics* **77**(5): 681–700.
- Heikkinen, J. and Heimonen, K. (2025). Media tone: The role of news and social media on heterogeneous inflation expectations, *Journal of Forecasting* **44**(3): 881–921.
- Hirsch, P., Feld, L. P. and Köhler, E. A. (2023). Breaking monetary policy news: The role of mass media coverage of ECB announcements for public inflation expectations, *CESifo Working Paper* (10285).
- Kim, C. H., Ranaldi, L. and Schularick, M. (2025). Inflation surprises and asset returns: a macrohistory perspective, *unpublished*.


- Korenok, O., Munro, D. and Chen, J. (2026). Inflation and attention thresholds, *The Review of Economics and Statistics* pp. 1–28.
- Lamla, M. J. and Lein, S. M. (2014). The role of media for consumers’ inflation expectation formation, *Journal of Economic Behavior & Organization* **106**: 62–77.
- Lamla, M. J. and Maag, T. (2012). The role of media for inflation forecast disagreement of households and professional forecasters, *Journal of Money, Credit and Banking* **44**(7): 1325–1350.
- Larsen, V. H., Thorsrud, L. A. and Zhulanova, J. (2021). News-driven inflation expectations and information rigidities, *Journal of Monetary Economics* **117**: 507–520.
- Le, Q. V. and Mikolov, T. (2014). Distributed Representations of Sentences and Documents, *International conference on machine learning*.
- Macaulay, A. and Song, W. (2023). News media, inflation, and sentiment, *AEA Papers and Proceedings* **113**: 172–176.
- Maćkowiak, B., Matějka, F. and Wiederholt, M. (2023). Rational inattention: A review, *Journal of Economic Literature* **61**(1): 226–273.
- Müller, H., Schmidt, T., Rieger, J., Hufnagel, L. M. and Hornig, N. (2022). A German inflation narrative. How the media frame price dynamics: Results from a RollingLDA analysis, *DoCMA Working Paper* (9).
- Nautz, D. and Scharff, J. (2012). Inflation and relative price variability in the euro area: evidence from a panel threshold model, *Applied Economics* **44**(4): 449–460.
- Newman, N., Ross Arguedas, A., Robertson, C. T., Nielsen, R. K. and Fletcher, R. (2025). Digital news report 2025, *Reuters Institute for the study of Journalism*.
- Pfäuti, O. (2023). The inflation attention threshold and inflation surges, *arXiv preprint arXiv:2308.09480*.
- Pfäuti, O. (2025). Inflation—who cares? monetary policy in times of low attention, *Journal of Money, Credit and Banking* **57**(5): 1211–1239.
- 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.
- Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C. and Sutskever, I. (2023). Robust speech recognition via large-scale weak supervision, *International conference on machine learning*, PMLR, pp. 28492–28518.
- Weber, M., Candia, B., Afrouzi, H., Ropele, T., Lluberas, R., Frache, S., Meyer, B. H., Kumar, S., Gorodnichenko, Y. and Georgarakos, D. (2025). Tell me something i don’t already know: Learning in low and high-inflation settings, *Econometrica* **93**(1).

- Xu, Y., Liu, Z., Chen, J. and Salem, S. (2024). How official TV news affect public inflation expectations? Evidence from the Chinese national broadcaster China Central Television, *International Journal of Finance & Economics* **29**(1): 819–831.
- Çekin, S. E. and Polattimur, H. (2025). Televised inflation: Measuring tv news coverage and its effect on household expectations, *ZEW Discussion Paper* (25-051).
- Řehůřek, R. and Sojka, P. (2010). Software framework for topic modelling with large corpora, *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, ELRA, Valletta, Malta, pp. 45–50.

Appendices

Appendix A Additional Tables and Figures

Figure A.1: Reach of the *Tagesschau* since 1992



Notes: This figure shows the average daily audience of the *Tagesschau* in millions of viewers by year since 1992. The time-series data are available from the corresponding Statista dataset.

Table A.1: Keywords to measure inflation-related TV reports

German	Translation
Inflation*	Inflation*
Deflation*	Deflation*
Disinflation*	Disinflation*
Geldwert	Value of money
Geldentwertung	Currency depreciation
Geldabwertung	Monetary devaluation
Kostenschub	Cost push
Lebenshaltungskosten	Cost of living
Preise	Prices
Preisstabilität	Price stability
Preisniveaustabilität	Stability of the price level
Verbraucherpreise	Consumer prices
Konsumentenpreise	Consumer prices
Ölpreise	Oil prices
Vermögenspreise	Asset prices
Energiepreise	Energy prices
Preisentwicklung	Price development
Preisniveau	Price level
Teuerung	Inflation (historical term)
Kaufkraft	Purchasing power
Kaufkraftverlust	Loss of purchasing power
Währungsstabilität	Currency stability

Notes: The table lists the German keywords used to identify inflation-related TV reports and their English translations. An asterisk indicates that all word forms starting with the respective stem are included. Searches are case-insensitive.

Figure A.2: Example page of the archive output

Titel

Serien-/Reihentitel Tagesschau F: 3254
 Sende-/Haupttitel Tagesschau vom 22.01.1965
 Beitragstitel 02 Berlin: Bundesbankpräsident Blessing zur Inflation|

Plan- und Sendedaten

Start-TC: 10:00:45.09

Ereignistyp	Btr.	Datum/Uhrzeit	Ereignisdauer	publiziert von/in
ESD		22.01.1965 20:00	1'28" in 14'40"	NDR @ ARD-1
ESD	02	22.01.1965	1'27"	NDR @ ARD-1

Urheberchaft/Produktion/Mitwirkung

Redaktion (Bereich) Tagesschau
 Programmbereich ARD-Aktuell
 Redaktion (Person) Ising
 Rohleder
 Sprechstimme Köpcke
 Berichterstattung 02 N.N.(SFB)

Sachinhalt

02 **Start-TC Beitrag:** 10:00:45.09 **Dauer:** 1'28"
 Veranstaltung mit rede Karl Blessing, Präs. Bundesbank, zur Diskonterhöhung und Währungsstabilität (Die Preissteigerung der letzten Jahre habe nichts mit einer Inflation im Stil der beiden Weltkriegsinflationen zu tun).

Schlagworte


IND > Standard 02 BRD; Zinserhöhung; Geldstabilität; Währungspolitik; Inflation
 IND > Sache 02 Währungspolitik; Inflation

Kategorie/Sparte

Sachgebiet Politik
 Präsentationsform Berichterstattung


Notes: The figure shows an example page from the archive output provided by the NDR archive.

Figure A.3: *Tagesschau* attention and content: original German terms




Notes: The figure shows the ten most frequent words in the titles and short descriptions of inflation-related *Tagesschau* reports by decade. Technical archive terms, city names, and the words “Inflation” and “Preise” are excluded. The words are shown in the original German. The line shows the historical *Tagesschau* attention index.

Figure A.4: Attention threshold full sample (January 1960-July 2024)



Notes: This figure plots regression results from a threshold model without control variables for the full period. The monthly inflation rate is the threshold variable. The model also includes a regime-independent constant. Confidence intervals for the threshold estimates are derived using a bootstrap procedure.

Figure A.5: Attention Thresholds for the Bundesbank (January 1960–December 1998) and the ECB period (January 1999–July 2024)



Notes: This figure plots regression results from threshold models without control variables for both the Bundesbank period and the ECB period. The monthly inflation rate is the threshold variable. The model also includes a regime-independent constant. Confidence intervals for the threshold estimates are derived using a bootstrap procedure.

Appendix B Further robustness checks

B.1 A segmented threshold model

As a robustness check, we estimate a segmented threshold model. In contrast to the baseline threshold specification in Section 3, which allows for different regression lines across inflation regimes, the segmented model keeps the relationship continuous at the threshold and allows only the slope to change.

The model is given by

$$TA_t = \alpha + \delta\pi_t + \beta(\pi_t - \gamma)^+ + \phi X_t + \varepsilon_t, \quad (\text{B.1})$$

where TA_t is the standardized *Tagesschau* attention index, π_t is the inflation rate, and

$$(\pi_t - \gamma)^+ = \max(\pi_t - \gamma, 0).$$

The coefficient δ measures the slope below the threshold. Above the threshold, the slope is $\delta + \beta$. As in the baseline specification, X_t includes unemployment, inflation volatility, log industrial production, and the interest rate.

Unlike in the baseline specification, where we allow for up to three thresholds and let the model select the number of regimes, the segmented specification imposes a single threshold. We estimate this threshold γ by exact grid search over observed inflation values in the 10% trimmed search region. For each candidate threshold, we construct the hinge term $(\pi_t - \gamma)^+$ and choose the threshold that minimizes the sum of squared residuals. We estimate the model for the same three samples as in the baseline analysis: the full sample, the Bundesbank period, and the ECB period, each with and without controls.

Table (B.2) reports the results. The estimated thresholds are broadly consistent with the baseline specification. In the full sample, the segmented model estimates a threshold of about 5.6%, with bootstrap confidence intervals that closely overlap with the baseline threshold estimates. In the Bundesbank period, the threshold is about 5.9%, again close to the baseline estimates. In the ECB period, the point estimate is lower, at about 1.7%, although the confidence intervals are wider. Overall, the segmented model confirms the main pattern from the baseline analysis: the estimated threshold is high in the full sample and the Bundesbank period, but lower in the ECB period.

The slope estimates also support the main pattern from the baseline model. Below the threshold, the relationship between inflation and *Tagesschau* attention is small and often statistically insignificant. Above the threshold, the implied slope $\delta + \beta$

Table B.2: Segmented threshold estimates

	Full Sample		Bundesbank		ECB	
	January 1960–July 2024		January 1960–December 1998		January 1999–July 2024	
Threshold (γ)	5.567 (4.888–5.877)	5.567 (4.888–5.882)	5.908 (4.645–6.117)	5.908 (4.667–6.117)	1.678 (0.713–3.711)	1.683 (0.536–3.811)
Region 1 ($\pi_t \leq \gamma$)						
Inflation rate	0.037 [0.024]	0.087 [0.033***]	0.068 [0.022***]	0.050 [0.030*]	0.013 [0.081]	0.070 [0.106]
Region 2 ($\gamma < \pi_t$)						
Inflation rate	1.157 [0.234***]	1.037 [0.240***]	1.397 [0.327***]	1.438 [0.324***]	0.506 [0.107***]	0.471 [0.143***]
Inflation volatility		0.224 [0.165]		−0.101 [0.146]		0.242 [0.296]
Unemployment rate		−0.037 [0.012***]		−0.089 [0.017***]		0.016 [0.042]
Log industrial production		0.006 [0.001***]		0.010 [0.002***]		−0.004 [0.007]
Interest rate		−0.042 [0.022**]		−0.049 [0.022**]		−0.015 [0.053]
#Obs.	775	771	468	468	307	303

Notes: The table reports estimates of Equation (B.1) for the full sample, the Bundesbank period, and the ECB period. The dependent variable is the standardized *Tagesschau* attention index. The threshold is estimated by exact grid search over observed inflation values using a trimming parameter of 10%. Region 1 reports the slope below the threshold, δ . Region 2 reports the implied slope above the threshold, $\delta + \beta$. The 95% confidence intervals for the estimated thresholds are obtained from a bootstrap procedure with 2,000 replications and are reported in parentheses. Robust standard errors are reported in brackets. All specifications include a constant. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

is much larger and statistically significant in all specifications. In the full sample, the above-threshold slope is about 1.16 without controls and 1.04 with controls. In the Bundesbank period, it is about 1.40 to 1.44. In the ECB period, it is lower in magnitude, around 0.47 to 0.51, but remains statistically significant.

The segmented specification therefore supports the main robustness conclusion: inflation coverage in the *Tagesschau* reacts nonlinearly to inflation. Even when the regression line is constrained to be continuous at the threshold, media attention becomes substantially more responsive once inflation exceeds the estimated threshold. The confidence intervals indicate some uncertainty around the exact threshold values, especially in the shorter ECB sample, but the qualitative pattern remains similar to the baseline results.

B.2 Positive and negative inflation surprises

As an additional robustness check, we examine whether the LP-IV results differ between positive and negative inflation surprises. In the baseline specification, the

first stage in Equation (14) uses the signed inflation surprise Z_t and its interaction with the release-coverage indicator R_t . This allows the direction of the surprise to matter, but it does not separately show whether the identifying variation comes primarily from positive or negative surprises.

We therefore replace Z_t with three alternative surprise measures. First, we use the absolute size of the surprise, $|Z_t|$. Second, we use the positive component,


$$Z_t^+ = \max(Z_t, 0).$$

Third, we use the absolute size of the negative component,

$$Z_t^- = |\min(Z_t, 0)|.$$

In each case, we interact the respective surprise measure with the release-coverage indicator R_t . In the positive-surprise specification, we also control for Z_t^- and $Z_t^- \times R_t$. Conversely, in the negative-surprise specification, we control for Z_t^+ and $Z_t^+ \times R_t$. The results suggest that the baseline effects are mainly driven by positive inflation surprises. Absolute surprises generate very weak first stages, with Kleibergen–Paap rk Wald F -statistics below 1 across specifications, and the corresponding impulse responses are not statistically significant. This suggests that *Tagesschau* coverage does not respond symmetrically to the size of inflation surprises. Negative surprises also generate weak first stages, again with F -statistics close to or below 1, and do not produce robust responses of household outcomes. By contrast, positive surprises provide the strongest first-stage relationship among the split specifications, with F -statistics between about 2 and 3. The corresponding impulse responses are similar in sign and shape to the baseline estimates, although somewhat smaller and less persistent.


Figure B.6: LP-IV robustness: absolute inflation surprises



Notes: This figure shows local-projection IV estimates of the responses of household attention, inflation perceptions, and inflation expectations to a one-standard-deviation shock in *Tagesschau* attention to inflation. The specification uses the absolute size of the inflation surprise and its interaction with *Tagesschau* coverage of the preliminary inflation release as instruments. Panel (a) uses the standardized *Tagesschau* attention index, while panel (b) uses the standardized number of inflation-related *Tagesschau* videos from our video dataset. Shaded areas indicate 68% and 90% confidence intervals based on heteroskedasticity- and autocorrelation-robust standard errors.


Since the sign-specific first stages are weaker than in the baseline specification, where the interaction instrument yields a Kleibergen–Paap rk Wald F -statistic of 11.9, we interpret these results as suggestive evidence on the direction of the identifying variation rather than as separate causal estimates. Overall, the pattern is consistent with asymmetric inflation news: inflation-related television coverage and household responses appear to be triggered mainly when inflation turns out higher than expected, rather than by inflation surprises in either direction. The weak results for absolute surprises suggest that pooling positive and negative surprises washes out this asymmetry.

Figure B.7: LP-IV robustness: positive inflation surprises



Notes: This figure shows local-projection IV estimates of the responses of household attention, inflation perceptions, and inflation expectations to a one-standard-deviation shock in *Tagesschau* attention to inflation. The specification uses the positive component of the inflation surprise and its interaction with *Tagesschau* coverage of the preliminary inflation release as instruments. Panel (a) uses the standardized *Tagesschau* attention index, while panel (b) uses the standardized number of inflation-related *Tagesschau* videos from our video dataset. Shaded areas indicate 68% and 90% confidence intervals based on heteroskedasticity- and autocorrelation-robust standard errors.

Figure B.8: LP-IV robustness: negative inflation surprises



Notes: This figure shows local-projection IV estimates of the responses of household attention, inflation perceptions, and inflation expectations to a one-standard-deviation shock in *Tagesschau* attention to inflation. The specification uses the absolute size of the negative inflation surprise component and its interaction with *Tagesschau* coverage of the preliminary inflation release as instruments. Panel (a) uses the standardized *Tagesschau* attention index, while panel (b) uses the standardized number of inflation-related *Tagesschau* videos from our video dataset. Shaded areas indicate 68% and 90% confidence intervals based on heteroskedasticity- and autocorrelation-robust standard errors.

Appendix C Construction of inflation surprises

Our instrumental-variable strategy uses inflation surprises, defined as the difference between realized year-on-year inflation and expected inflation formed before the release:

$$Z_t = \pi_t - \mathbb{E}_{t-1}[\pi_t], \quad (\text{C.2})$$

where π_t denotes realized inflation in month t . Since long monthly survey-based measures of inflation expectations are not available for our full sample, we construct $\mathbb{E}_{t-1}[\pi_t]$ using model-based forecasts. This follows Kim et al. (2025), who use time-series forecasts as proxies for expectations when direct expectation measures are unavailable. Similarly, we combine several forecasting models rather than relying on a single specification.

We compute one-month-ahead forecasts of year-on-year inflation from a set of univariate and multivariate time-series models. The univariate models include ARMA(1,1) specifications estimated both over the available sample and over rolling 36-month windows, a random-walk forecast, and an unobserved-components model with a local-level trend and irregular component. We also estimate ARMA(1,1) models for transitory inflation. To do so, we first estimate trend inflation $\bar{\pi}_t$ using the local-level unobserved-components model and define transitory inflation as

$$\tilde{\pi}_t = \bar{\pi}_t - \pi_t.$$

Forecasts of transitory inflation are transformed back into level forecasts using the lagged trend component.

We further estimate a backward-looking Phillips-curve specification for transitory inflation, using lagged transitory inflation and the lagged output gap as predictors. The output gap is obtained from an HP filter applied to log industrial production. Finally, we estimate VAR(2) models with inflation, the output gap, and the short-term interest rate. The VAR models are estimated both for inflation in levels and for transitory inflation, using rolling 36-month windows and available-sample in-sample forecasts.


Let $\hat{\pi}_{t|t-1}^j$ denote the forecast of inflation in month t from model j . We combine the individual model forecasts into a single model-based expectation measure using the median across models:

$$\mathbb{E}_{t-1}[\pi_t] = \text{median}_j \left(\hat{\pi}_{t|t-1}^j \right).$$

Using the median reduces the influence of outlier forecasts from individual models. The inflation surprise is then the difference between realized inflation and this

model-based expectation. Figure (C.9) plots the resulting inflation surprise.

Figure C.9: Model-based inflation surprises



Notes: The figure shows monthly inflation surprises, defined as realized year-on-year inflation minus model-based expected inflation formed before the release.

Affidavit

Ich erkläre hiermit, dass ich die vorgelegten und nachfolgend aufgelisteten Aufsätze selbstständig und nur mit den Hilfen angefertigt habe, die im jeweiligen Aufsatz angegeben oder zusätzlich in der nachfolgenden Liste aufgeführt sind. In der Zusammenarbeit mit den angeführten Koautoren war ich mindestens anteilig beteiligt. Bei den von mir durchgeführten und in den Aufsätzen erwähnten Untersuchungen habe ich die Grundsätze guter wissenschaftlicher Praxis, wie sie in der Satzung der Justus-Liebig-Universität Gießen zur Sicherung guter wissenschaftlicher Praxis niedergelegt sind, eingehalten.

Moritz Grebe

18.05.2026

I hereby declare that I completed the papers submitted and listed hereafter independently and only with those forms of support mentioned in the relevant paper. When working with the authors listed, I contributed no less than a proportionate share of the work. In the analyses that I have conducted and to which I refer in the papers, I have followed the principles of good academic practice, as stated in the Statute of Justus Liebig University Giessen for ensuring good scientific practice.

Moritz Grebe

May 18, 2026

Submitted Papers

- I. Grebe, M., Kandemir, S. and Tillmann, P. (2024). Uncertainty about the war in Ukraine: Measurement and effects on the German economy, *Journal of Economic Behavior and Organization* **217**: 493–506.
- II. Grebe, M., Kandemir, S. and Tillmann, P. (2024). Geopolitics in the boardroom: How German managers respond to the war in Ukraine, available at SSRN 5098484.
- III. Grebe, M. and Hassanin, S. (2026). Monetary metanomics: The effect of monetary policy on virtual real estate prices.
- IV. Grebe, M. and Tillmann, P. (2025). Household expectations and dissent among policymakers, *European Journal of Political Economy* **86**: 102638.
- V. Grebe, M. and Tillmann, P. (2026). Inflation expectations and a conservative central banker: A natural experiment, IMFS Working Paper Series No. 234.
- VI. Bruns, B. J. and Grebe, M. (2026). Prices at prime time: How TV news shape inflation attention in Germany.