Year : 2023

Three Essays in Macroeconomics

Bolliger Elio Rico

UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES

DÉPARTEMENT D'ÉCONOMIE

Three Essays in Macroeconomics

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales
de l'Université de Lausanne

pour l'obtention du grade de

Doctorat en Économie

par

Elio Rico BOLLIGER

Directrice de thèse
Prof. Kenza BENHIMA

Co-Directeur de thèse
Prof. Adrian BRUHIN

Jury

Prof. Boris Nikolov, président
Prof. Jean-Paul Renne, expert interne
Prof. Lorenz Küng, expert externe

Lausanne
2023

UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES

DÉPARTEMENT D'ÉCONOMIE

Three Essays in Macroeconomics

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales
de l'Université de Lausanne

pour l'obtention du grade de

Doctorat en Économie

par

Elio Rico BOLLIGER

Directrice de thèse
Prof. Kenza BENHIMA

Co-Directeur de thèse
Prof. Adrian BRUHIN

Jury

Prof. Boris Nikolov, président
Prof. Jean-Paul Renne, expert interne
Prof. Lorenz Küng, expert externe

Lausanne
2023

UNIL | Université de Lausanne
Faculté des hautes études commerciales
Internef - 261
CH-1015 Lausanne

IMPRIMATUR

La Faculté des hautes études commerciales de l'Université de Lausanne autorise l'impression de la thèse de doctorat rédigée par

Elio Bolliger

intitulée

Three Essays in Macroeconomics

sans se prononcer sur les opinions exprimées dans cette thèse.

Lausanne, le 22.09.2023

Professeure Marianne Schmid Mast, Doyenne

Members of the thesis committee

Prof. Kenza Benhima

Full Professor, University of Lausanne

Thesis Supervisor

Prof. Adrian Bruhin

Full Professor, University of Lausanne

Thesis Co-Supervisor

Prof. Jean-Paul Renne

Full Professor, University of Lausanne

Internal member of the thesis committee

Prof. Lorenz Küng

Associate Professor, Università della Svizzera italiana (USI)

External member of the thesis committee

University of Lausanne
Faculty of Business and Economics

PhD in Economics

I hereby certify that I have examined the doctoral thesis of

Elio Rico BOLLIGER

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: _______________ Date: 31/07/2023

Prof. Kenza BENHIMA
Thesis supervisor

University of Lausanne
Faculty of Business and Economics

PhD in Economics

I hereby certify that I have examined the doctoral thesis of

Elio Rico BOLLIGER

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: ______________________ Date: 31.7.2023

Prof. Adrian BRUHIN
Thesis co-supervisor

University of Lausanne
Faculty of Business and Economics

PhD in Economics

I hereby certify that I have examined the doctoral thesis of

Elio Rico BOLLIGER

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: ______________________. Date: 31.07.2023

Prof. Jean-Paul RENNE
Internal member of the doctoral committee

University of Lausanne
Faculty of Business and Economics

PhD in Economics

I hereby certify that I have examined the doctoral thesis of

Elio Rico BOLLIGER

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: Date: 31.7.2023

Prof. Lorenz KUENG
External member of the doctoral committee

Three Essays in Macroeconomics

Elio Bolliger
University of Lausanne

2023

Acknowledgments

My gratitude goes first and foremost to my supervisor, Kenza Benhima, and co-supervisor, Adrian Bruhin.

During the courses taught by Kenza at the University of Lausanne, I was quickly impressed by her methodological approach and how she allowed students to approach sometimes complex topics in macroeconomics not only theoretically but also in an applied context with a strong focus on policy implications. This greatly shaped my fascination with macroeconomics.

Furthermore, Kenza provided me with unique opportunities to enhance my professional experience, including teaching, additional research projects, and hands-on experience at the Centre of Applied Economic Research, CREA. I enormously benefited from her knowledge in those matters and will surely continue to do so.

Most importantly, Kenza consistently demonstrated herself to be a highly supportive mentor who cares about the well-being and success of all her students.

I had the great pleasure of following one of the courses taught by Adrian Bruhin during my Master's programme in Lausanne. The way Adrian has taught topics related to financial institutions left an imprint on my way of thinking about the financial crisis. Furthermore, he smartly combined perspectives from models with his extensive experience in central banking, coupled with insights from research.

Adrian also gave me the unique opportunity to conduct my thesis chapter in collaboration with the Swiss National Bank - an experience for which I am very thankful. I benefited immensely from this project, as Adrian's guidance provided me with crucial skills in writing, drafting, and developing research ideas. Paired with his continuing support and availability at any time, he laid the foundation for a successful start to my PhD.

Special thanks also go to the members of my thesis committee.

I am very grateful to Prof. Jean-Paul Renne, internal expert, for his helpful guidance throughout my studies at the University of Lausanne. He has always been very approachable, paired with an intrinsic motivation to help others benefit from his expertise in econometrics. Throughout my PhD, he provided continuous support and help, which was instrumental in the development of the chapters.

I would like to express my sincere gratitude to Associate Prof. Lorenz Küng, external expert, for the fruitful discussions as well as his invaluable comments and recommendations, which significantly enhanced the quality and content of my thesis.

I also want to warmly thank Associate Prof. Andreas Fuster, my co-author on the first chapter.

His established expertise and insightful feedback undoubtedly had a significant impact on the improvement of our work. I am also very grateful for his support during the job market process.
I would also like to thank all members of the academic community who participated in improving the content of this thesis.
I am deeply grateful to Jeremy, Brendan, Amalia, Fléance, and Adrien. As I embarked on my journey of studying economics in Lausanne, I consider myself incredibly fortunate to have been surrounded by this remarkable group of scholars. Their invaluable support and inspiration played a significant role in shaping my decision to pursue a PhD. I am truly honored to not only have them as friends but also to have been able to share many journeys and festivities together.
Kudos also go to Matthias, Anis, and Maxime, with whom I enjoyed many excellent and (really) tense table football games, classy football action on the pitch, and many more memorable discussions not very much related to economics.
Surely, a big thank you also goes to Noëmi, Andrea, Paola, and Pauline, all of whom contributed to a perfect working and social environment during and after work in our shared offices and teaching activities.
I am also very grateful to all my friends and colleagues from the iconic fifth floor at UNIL, with whom I have shared countless discussions and good memories, whether around a coffee or a beer: Giacomo, Kevin, Guillaume, Andrea, Laia, Jeremias, Tobias, Fabrizio, Sofia, Pascal, Alice, Bernhard, Christoph, Katja, Resuf, Marius, and all others I may have forgotten.
Many thanks also to the Gerzensee crew 2018-2019, who made my time at the doctoral school memorable: Evert, Florian, Hyacinthe Georges, Juliette, Flavia, Tabea, Anja, Pia, Sarah.
I am profoundly thankful to my family.
My mom and dad, Ines and Jean-Pierre, for their unwavering encouragement and support throughout my life. They granted me the freedom to pursue any aspirations I set my mind to, and I will forever appreciate their contributions.
Thanks also to my brother Renato, who sparked my interest in economics at an early stage of my studies. Our shared Futsal experience laid the foundation for a healthy study-life balance, providing me with the necessary motivation to pursue my academic goals.
Finally, I would like to thank Monica for accompanying me on this challenging but fulfilling journey. Not only did she exhibit extraordinary patience and support as I discussed macroeconomics and time series analysis or coding on the weekend, but Monica also imparted invaluable life lessons, teaching me the importance of living in the present, finding moments of relaxation, and embracing joy to its fullest.
Meeting her throughout my PhD journey fills me with deep gratitude and has enriched every moment since.

Contents

Introduction

In economic research, frictions are omnipresent. These frictions can be understood as various barriers, inefficiencies, and imperfections that impede the smooth functioning of markets and economic activities. They arise in different parts of the economy and have a significant impact on resource allocation, decision-making, and overall economic outcomes.

Given the various impacts of frictions on the economy and individuals, identifying and studying them is essential for economists and policymakers.

The aim of this thesis is to investigate such frictions, emphasizing their identification and analyzing their effects on macroeconomic outcomes, in order to provide new insights for policymakers and academics.

Specifically, the three chapters contribute to the bodies of research on financial frictions related to macroprudential regulation as well as informational frictions, studying how professional economists and households use and are affected by information, with the ultimate goal of learning more about expectations and perceptions of important macroeconomic variables.

In the first chapter, co-authored with Andreas Fuster, Adrian Bruhin, and Maja Ganarin, we examine the impact of macroprudential policies on homeownership and their effect on borrowing constraints in Switzerland. This analysis is crucial for understanding how these policies influence homeownership, particularly considering the institutional goal of promoting homeownership in Switzerland and the documented social and financial benefits of owning a house.

To investigate the effects of macroprudential policies on homeownership, we utilize a comprehensive and granular dataset of tax reports from the canton of Bern. What sets our study apart is the ability to differentiate between inheritances and predeath bequests in the tax reports. While predeath bequests are more targeted transfers, inheritances are considered relatively unpredictable. We use these transfers to identify borrowing constraints of households and elucidate the underlying mechanism using a simplified theoretical framework.

Our primary analysis employs an event study to compare the effects of macroprudential regulation on the probability of transitioning into homeownership (extensive margin). Specifically, we examine how transfers influence the likelihood of transitioning into homeownership and assess whether they have become more significant following the introduction of macroprudential policies. Additionally, we use a related approach to investigate whether similar effects are observed regarding the purchase price of a home (intensive margin).

On average, we find a moderate decrease from 3.4% to 3% in the probability to transition into homeownership after 2012 which is likely due to the introduction of the macroprudential policies. This decrease is particularly pronounced for younger households and households with low wealth. Transfers are in general important to overcome borrowing constraints. They increase the probability to transition into homeownership from 2.4 up to 12 percentage points (p-values$<$0.01) for inheritances and predeath bequests, respectively. Moreover, after the introduction of macroprudential policies, predeath bequests have become more important and increase the probability to transition into homeownership by an additional 0.8 percentage points (p-value$<$0.05).

Regarding the intensive margin, we find some evidence that since 2012, predeath bequest and wealth also have had a higher effect on the purchase price of a home, reflecting the results we observe for the extensive margin. A household that receives a predeath bequest or who has above median wealth increases the purchase price of a home by 3% and 5%, respectively. Our robustness checks suggest that the stronger effect of predeath bequests after 2012 cannot solely be explained by rising house prices. Furthermore, evidence from homeowners who acquire an additional home shows that they are not confronted with higher borrowing constraints. Predeath bequests have no different effects before and after 2012 for homeowners acquiring an additional home.

Our results have the following implications. Macroprudential policies aim at reducing imbalances on the housing and mortgage market. Hence, a tightening of borrowing constraints has to be expected. However, our results show that younger households, who had less time to accumulate wealth, are particularly affected by tighter borrowing constraints. Via predeath bequests, they rely more heavily on their family wealth to transition into homeownership. As family wealth is heterogeneous, the introduction of macroprudential policies may entail distributional consequences.

In the second chapter, in collaboration with Kenza Benhima, we study the information set of professional forecasters to identify information asymmetries between local and foreign forecasters. Our analysis focuses on updating behavior and forecasting precision for two crucial macroeconomic fundamentals: inflation and real GDP growth. The results obtained serve as a fundamental basis for refining international finance and trade models with heterogeneous information.

To achieve this, we leverage a unique dataset containing yearly inflation and GDP growth forecasts, gathered at a monthly frequency. By differentiating between local and foreign forecasters based on the countries they forecast for, we enhance the dataset's granularity and control for a wide range of fixed effects.

Our analysis reveals that foreign forecasters update their predictions about 10% less frequently than local forecasters and tend to make more mistakes, with their excess absolute error reaching up to 9%, depending on the forecast horizon and variable. The local advantage is especially large when predicting inflation as opposed to GDP and it is stronger for shorter forecasting horizons.

We then investigate the role of information frictions and behavioral biases in explaining our results about forecasting precision. We do this in two steps. First, we rule out behavioral biases such as over-reaction to new information as explanations of the foreigners' excess mistakes, by showing that the local and foreign behavioral biases do not differ systematically. Second, we test for the relative precision of local and foreign forecasters' private information, and find that local forecasters have more precise private information. Our approach builds upon the growing literature that uses model-based tests to identify frictions in survey respondents' expectation formation, providing robust results even in the presence of public signals.

Surprisingly, we observe that the information asymmetry does not diminish with reduced forecasting uncertainty; in fact, it tends to increase. The local advantage is most prominent in short forecasting horizons, inflation predictions, and in larger countries, despite these variables affecting average forecasting uncertainty. However, the difference in forecast errors between local and foreign forecasters does not seem to correlate with a country's development status, institutional quality, or the volatility of business cycles or financial markets.

Our findings suggest that information tends to flow more readily to local forecasters, likely due to their better access to locally-produced information and their understanding of relevant information release schedules. This information asymmetry is stronger for nowcasting (when forecasting the current year's GDP growth or inflation), and it increases in the course of a year (the asymmetry is higher in December than in January). This is consistent with the idea that local forecasters are exposed to the regular releases of partial GDP growth and inflation figures and integrate this information faster.

While we do not directly measure the incentives for information acquisition at the forecaster level, we find that forecasts issued by the financial industry are generally more precise, likely driven by their incentives to allocate portfolios effectively. However, there is no significant difference between the local advantage of the financial sector and that of the non-financial sector.

In conclusion, our study offers direct evidence of information asymmetries between local and foreign forecasters. Furthermore, we provide new model-based tests to identify such asymmetries that are robust to public signals. Finally, our results provide a basis for disciplining international finance and trade models with heterogeneous information.

In the third chapter, single-authored, we investigate how information in the form of news articles affects households' inflation expectations and perceptions. Since the onset of the Covid-19 crisis, inflation rates have spiked to historically high levels, presenting a significant challenge for policymakers and households alike. Consequently, it is crucial to better understand the drivers of expectations and perceptions.

To study this question, we use a novel newspaper database from Switzerland. We construct two indices from this data that measure inflation news content in newspapers. Specifically, a quantitative news measure, which represents the difference between articles writing about an inflation increase versus an inflation decrease, and a novel qualitative inflation sentiment measure, that shows how newspapers assess the inflation rate (positively or negatively).

To investigate how these news measures affect inflation expectations and perceptions, we proceed in two steps. First, we analyze and compare inflation news reporting in French and German articles. Second, we study how news reporting affects inflation expectations and perceptions, exploiting the language barrier of French and German-speaking households in Switzerland.

The results from our first step challenge conflicting evidence of the so-called negativity bias in inflation news reporting. This bias refers to the over-reporting of bad news compared to good news. We show that newspapers in Switzerland do not suffer from such a bias. Despite time-varying differences in inflation-reporting of newspapers written in French compared to those written in German, we find no systematic difference for both the quantitative and qualitative news measure. This is important for interpreting our results from the second step.

In the second step of our study, we leverage the language barrier in Switzerland to adopt a more rigorous econometric approach compared to existing literature. We assume that households living in the French-speaking part consume more French-written news, and vice-versa for households living in the German-speaking part with respect to German-written news. This allows us to control for time fixed-effects, that capture any confounds affecting expectations, perceptions and news reporting at the same time.

Our results provide new evidence that both the quantitative and qualitative news measures have a significant effect on expectations and perceptions. For the quantitative news measure, more articles written about an inflation increase compared to a decrease augment inflation expectations, especially in times of a high inflation environment. On the other hand, we find that the qualitative inflation news measure has a countercyclical effect. Namely, a more positive inflation assessment of newspapers can have a dampening effect on expectations and perceptions when current inflation is increasing and it can have the opposite effect when inflation is decreasing. The results of the qualitative news measure are therefore similar to evidence showing that general economic sentiment can also Furthermore, we show that the effects of news are increasing in age and stronger for households living in the German-speaking part. This is consistent with the fact that elderly households spend more time reading the news and are more affected by inflation, given they are mostly net savers. Also, households in the German-speaking part of Switzerland have been shown to be more inflation averse than those in the French-speaking part.

To conclude, we provide new evidence of how news affects inflation expectations and perceptions by exploiting the language barrier in Switzerland. These results are important from a policy perspective for several reasons. First, inflation expectations can be self-fulfilling and well-anchored inflation expectations improve the effectiveness of monetary policy. As newspapers frequently report about the communication of central banks, policymakers could use the news channel to better anchor inflation expectations through inflation sentiment. Finally, our novel inflation sentiment measure can be used as a timely policy indicator that shows how newspapers assess the inflation environment.

Chapter 1

The Effect of Macroprudential Policies on Homeownership: Evidence from Switzerland

Elio Bolliger
University of Lausanne

Adrian Bruhin
University of Lausanne

Andreas Fuster
EPFL, Swiss Finance Institute

Maja Ganarin
Swiss National Bank

Abstract

This paper analyzes how the introduction of macroprudential policies in the Swiss residential mortgage market affected the propensity of households to become homeowners. We exploit a unique administrative data set of individual tax records containing detailed financial and socio-demographic information. We show that the mean share of renter households transitioning into homeownership decreased from 3.4% per year in the five years prior to the introduction of macroprudential policies to 3.0% per year in the five years afterward. This decrease is more pronounced for young and middle-aged households with relatively low income and wealth, suggesting that it is at least partly due to a tightening in borrowing constraints. Moreover, intergenerational transfers in the form of predeath bequests have become more important for homebuying both at the extensive and intensive margin.

Keywords: Homeownership, Macroprudential Policy, Borrowing Constraints, Intrafamily Transfers, Wealth Inequality.

JEL: E5, D14, D31, G18.

1. Introduction

Since the global financial crisis, the housing market and its participants have attracted growing attention from policy makers. The risk of systemic crises arising from the vulnerability of highly-leveraged households and banks has strengthened the case for policy intervention. One increasingly common set of interventions, such as countercyclical buffers (CCyB) on bank capital or leverage restrictions on households, is referred to as macroprudential policies. These policies aim to strengthen banks' and borrowers' resilience during a housing market downturn and to mitigate the build up of systemic risk in the first place. As such, macroprudential policies can be beneficial from an aggregate perspective. At the same time, to be effective, they tighten borrowing constraints and could therefore make homeownership more difficult to attain for some households.

This paper focuses on the latter aspect. It studies how the propensity of renter households to transition into homeownership changed after the introduction of macroprudential policies in Switzerland, and how these policies have affected borrowing constraints. The paper exploits a comprehensive administrative data set on individual tax reports from one of the largest Swiss cantons, Bern, for the years 2007–2016.[1] This data set contains precise information on taxpayers' income and wealth. Furthermore, similarly to Blickle and Brown (2019), it allows us to identify borrowing constraints by exploiting intergenerational wealth transfers, such as inheritances and predeath bequests.

After the financial crisis, the first macroprudential policy in Switzerland was introduced in 2012. It imposed, among other things, stricter requirements on down-payments. As a consequence, home buyers need to finance at least 10 percent of the housing value with own equity capital, without drawing from their mandatory pension savings. The policy aims to enhance financial stability (in an aging population) primarily by ensuring that homeowners have sufficient pension savings to finance retirement and are not forced to sell their homes or take out home equity loans once they retire. In addition, it may also enhance financial stability by reducing the overall financial burden on households from servicing their debt and rebuilding their pension savings as well as by screening out financially literate households who built up savings besides their mandatory pension funds (Swiss Financial Market Supervisory Authority FINMA, 2012; Seiler Zimmermann, 2012; International Monetary Fund, 2014).

The policy has a direct impact on borrowing constraints – particularly for households from low-wealth families who cannot use predeath bequests to finance the necessary down-payment. This policy was soon followed by additional measures, for example, a CCyB, which aimed at increasing the resilience of the banking sector and, along with other measures, dampening

[1]We also confirm some of our main results in a second tax data set, for the canton of Lucerne, as well as in nationwide survey data.

credit growth in the mortgage market (Danthine, 2013).

We start by discussing descriptive evidence. The mean share of renter households transitioning into homeownership decreased by 11.8% from 3.4% per year in the five years prior to the introduction of macroprudential policies to 3% per year in the five years afterward. This decrease is primarily driven by young and middle-aged households with relatively low income and wealth, who already had a lower probability of becoming homeowners before the introduction of the macroprudential policies than their more affluent counterparts. Moreover, the total amount withdrawn from mandatory pension savings dropped. At the same time, after the measures were introduced, a larger share of young and middle-aged households that do become homeowners received an intrafamily transfer (namely, a predeath bequest), presumably to help with the down-payment. These patterns suggest that households tap more into family wealth to overcome the tighter borrowing constraints and finance the transition into homeownership.

Next, we estimate the effects of households receiving a transfer on the extensive and intensive margins of homeownership (where the intensive margin refers to the value of the home purchased by a newly owning household). Motivated by a simple theoretical framework, we use intergenerational wealth transfers to identify borrowing constraints.

At the extensive margin, our estimates reveal that predeath bequests became more important after the introduction of the macroprudential policies. The estimated effect of receiving a predeath bequest on the probability of becoming a homeowner is 12 percentage points prior to 2012, and increases by approximately another 0.8 percentage points after 2012 (controlling for many other observable characteristics). This implies that receiving a predeath bequest more than offsets the overall decrease in renters' probability of becoming homeowners, given that after 2012, the annual probability of transitioning into homeownership decreased by 0.45 percentage points for households that receive no predeath bequest.[2]

At the intensive margin, we find that, after the introduction of the macroprudential policies, being above median wealth has a stronger positive effect on the purchase price of the new home. Receiving a predeath bequest has also a stronger positive effect after 2012. We estimate that the average initial effect, a 9 percent higher purchase price for households receiving a predeath bequest, increased by an additional 3 percentage points, although this estimate is not very precise. While predeath bequests have become more important after 2012, especially at the extensive margin, we find no such evidence for inheritances.[3]

---

[2]The point estimate in fact implies a 0.35 percentage points *higher* propensity of becoming homeowners after 2012 for the group of households that receive transfers, but this effect is not statistically significant (in contrast to the decrease in the propensity for those renters who do not receive predeath bequests, which is highly significant).

[3]A possible reason for the different effects of inheritances and predeath bequests may be that these transfers occur, on average, at a different stage in a household's life-cycle. In fact, the average age of households (across the two spouses) receiving an inheritance is 58 years, whereas for predeath bequests the average age is 47 years. As

A potential concern for our interpretation is that the change in the probability of transitioning into homeownership and the stronger effect of predeath bequests are not driven by macroprudential policies, but simply reflect steadily increasing house prices over time. To test this alternative channel, we examine whether predeath bequests have stronger effects in regions with higher price-to-rent ratios (or just price levels). However, we find no evidence of such differential effects.

Another interesting feature of our data set is that we are able to identify the number of properties a household owns. This allows us to analyze the effects of macroprudential policies on households with potentially different characteristics than first-time homeowners. As real estate is expensive, households with more than one property are likely less credit constrained and, therefore, less affected by the introduction of the macroprudential policies. Indeed, we find that, for such households, predeath bequests have the same effect on the propensity to acquire an additional property before and after the introduction of the macroprudential policies.

Note that, from a theoretical perspective, tighter borrowing constraints do not necessarily affect homeownership. Property prices may adjust downward, making homeownership more affordable and, thus, leaving the allocation of real estate unchanged. However, according to our results, this is not the equilibrium outcome in the Swiss setting, perhaps because a substantial fraction of the real estate is owned by investors who are less affected by macroprudential policies.

Our results have several implications. First, they suggest that macroprudential policies succeeded at fostering financial stability, partly because they prevented households with insufficient hard equity from transitioning into homeownership and financing the necessary downpayment by tapping into their pension savings.

Consequently, these policies likely lower the likelihood and depth of a potential property price correction and ensuing recession. Second, if some households transition into homeownership later or forgo homeownership entirely due to these policies, there may be distributional consequences—akin to those prominently discussed in the context of the prolonged low interest rates (Coibion et al., 2017; Saiki and Frost, 2014).[4] However, the extent of such potential distributional consequences is unclear, as macroprudential policies not only tighten borrow-

entry into homeownership is most common between ages 35 to 50 years (see Appendix Figure A1), an inheritance might occur when most households already transitioned into homeownership or decided to stay renters.

[4]The literature documents several potential wealth benefits of homeownership. Homeowners exhibit higher savings rates, leading to higher net wealth compared to renters (Di et al., 2007; Turner and Luea, 2009). Other benefits are the higher internal rate of return and the favorable tax treatment of housing compared to alternative investments – which are less relevant in Switzerland, where the imputed rent of housing costs counts as taxable income. Sodini et al. (2016) find that homeownership boosts consumption. In general, the financial benefits of homeownership depend on its duration, and whether households can maintain homeownership also during economic downturns (Goodman and Mayer, 2018).

ing constraints of certain households but probably also dampen future house price growth. This, in turn, makes homeownership more accessible again to constrained households.

Our analysis contributes to the small but growing literature about the effects of macroprudential policies on homeownership. Several studies in this literature use loan-level data and focus exclusively on new mortgage originations. For Ireland, Kinghan et al. (2019) document a decrease in loan-to-value (LTV) ratios and report that high-income households increased their down-payments to keep the house price constant, while low-income households purchased a cheaper home to keep the down-payment constant. Using similar data, Acharya et al. (2021) show evidence that banks reallocate mortgage loans away from low-income households and urban regions toward high-income households and more rural regions, resulting in a dampening effect on house prices. Peydró et al. (2020) find similar effects on low-income borrowers and house prices in the United Kingdom. Tzur-Ilan (2019) studies the effects of LTV limits in Israel on housing choices on the intensive margin, and finds that tighter borrowing constraints can lead to higher commuting costs. In Switzerland, the introduction of macroprudential policies has been found to reduce high-LTV mortgages (Behncke, 2022), shift lending from residential mortgages to commercial loans (Auer et al., 2022), and reallocate mortgage lending from less to more resilient banks (Basten, 2020).

In contrast to the above studies, our paper uses administrative data comprising the universe of households including renters. Only a few other studies use similar data. For the Netherlands, Van Bekkum et al. (2020) find that macroprudential policies reduced the share of households transitioning into homeownership, especially among liquidity constrained households. In related work, Aastveit et al. (2022) find similar effects for the extensive margin in Norway, and further show that those households that do still buy subsequently have less liquid wealth and more volatile consumption. We extend this literature and show that households seeking homeownership react to the introduction of macroprudential policies by relying more on intergenerational wealth transfers to overcome the tighter borrowing constraints. Moreover, young and middle-aged households with relatively low income and wealth are particularly affected by the stricter requirements on down-payments.

The paper also contributes to the strand of literature on borrowing constraints and homeownership. In early work, Linneman and Wachter (1989) and Haurin et al. (1997) document the importance of income and wealth constraints for the households' propensity to transition into homeownership. More recently, Fuster and Zafar (2016, 2021) use strategic surveys to highlight the relevance of down-payment requirements for homebuying, especially for households who are more liquidity constrained. Benetton et al. (2019) and Tracey and van Horen (2021) find that a large-scale UK policy initiative, which relaxed down-payment constraints, increased access to homeownership, especially for young households. For Switzerland, Bütler and Stadelmann (2020) use administrative data from a pension provider to analyze the change in pension withdrawals for funding homeownership after the introduction of the stricter requirements on down-payments. In line with our findings, they document a sizeable decrease in the probability of households withdrawing pension savings and present suggestive evidence of a decrease in aggregate home purchase activity.

Intergenerational wealth transfers within families are one way to overcome borrowing constraints. Recent work by Bond and Eriksen (2021) confirms the significant role of such transfers and points to the potential of family wealth to explain differences in homeownership between white and non-white households in the U.S. In particular, they find that differences in parental wealth explain the largest share of the white versus non-white gap in the probability of becoming a homeowner. In related work, Brandsaas (2021) shows that wealth transfers from parents are instrumental for young US households to transition into homeownership. Similarly, in Swiss data, Blickle and Brown (2019) find that intergenerational wealth transfers increase the probability of transitioning into homeownership by 6 to 8 percentage points. They rely on the nationwide Swiss Household Panel Data Survey (SHP), which we also exploit for robustness checks.

We extend this strand of literature by analyzing whether the importance of wealth transfers has changed in response to the introduction of macroprudential policies. Moreover, in contrast to the SHP, our data set allows us to discriminate between the effects of inheritances and predeath bequests.

The rest of the paper is organized as follows. Section 2 provides some background about the housing market in Switzerland and describes the macroprudential policies. Section 2 outlines the data. Section 4 provides a descriptive analysis of the importance of wealth and income to transition into homeownership and analyzes the use of wealth transfers. Section 5 presents the theoretical framework. Section 6 describes our empirical strategy to estimate the effect of the introduction of the macroprudential policies on borrowing constraints. Section 7 shows our main results both at the extensive and the intensive margin. Section 8 provides further evidence that the tightening of borrowing constraints is likely due to the introduction of the macroprudential policies and features some robustness checks. Finally, Section 8 concludes.

2. Background

2.1 Housing Market in Switzerland

Even though the promotion of homeownership is a constitutional goal in Switzerland, the share of homeowners is markedly lower than in neighboring countries. In 2019, 41.6% of Swiss households owned either a house or an apartment, compared to 51.1% in Germany, 55.2% in Austria, 64.1% in France, and 72.4% in Italy (Eurostat, 2021).

In the literature, the scarcity of land as well as the well-developed and regulated rental market are often named among the main reasons for the difference in the rate of homeownership between Switzerland and its neighbors (Kuhn and Grabka, 2018; Schneider and Wagner, 2015). Moreover, the purchase of individual units in apartment buildings was not allowed in Switzerland before 1965 (Wehrmüller, 2014).

Relatively high down-payment requirements pose an additional barrier for transitioning into homeownership in Switzerland. Households typically have to finance at least 20% of the acquired home's value with equity. However, to facilitate meeting these down-payment requirements, households are allowed to withdraw and use part of their tax-privileged pension savings.[5]

The Swiss pension system is based on three pillars.[6] The first pillar comprises a pay-as-you-go insurance. Its goal is to guarantee a living standard after retirement at the subsistence level. The second pillar comprises an individual pension fund. Its goal is to guarantee the continuation of the current living standard. The first and the second pillar are mandatory for all employees. The third pillar is a voluntary and tax-privileged private supplement comprising additional pension accounts and funds.

Importantly for us, Swiss households can withdraw a limited amount ahead of retirement from their second and third pillar savings to provide the required equity for acquiring their principal residence. Such withdrawals are common (e.g., Seiler Zimmermann, 2013), but as discussed below, a policy change introduced in 2012 restricted their use.

Since the year 2000, Swiss home prices have grown much faster than household incomes. Figure 2.1 shows the evolution of apartment and house prices, average household income, and mortgage volume, all in nominal terms and indexed to the year 2000. Between 2000 and 2018, apartment and house prices increased by 114% and 74%, respectively, while average household income grew by only 20%. Over the same period, mortgage volume increased by 124%.

The widening gap between house prices and household income as well as the trend towards more mortgage debt has led to concerns among regulators and policy makers. For instance, the Swiss National Bank noted in its 2012 Financial Stability Report that rising debt relative to GDP, reflected by the increase in mortgage volume in the household sector, makes households vulnerable to potential macroeconomic shocks (SNB, 2012). These concerns are compounded by the fact that mortgages are the most important asset of Swiss banks, accounting, on average, for 70% of the domestically focused banks' total assets (Behncke, 2022).

[5]Besides down-payment requirements that restrict the LTV ratio of a household, banks usually do not grant loans to households whose monthly installments are higher than a certain threshold of the household's income. In Switzerland, the threshold is usually at 33%.

[6]See e.g. https://www.ch.ch/en/manage-retirement-provision/.

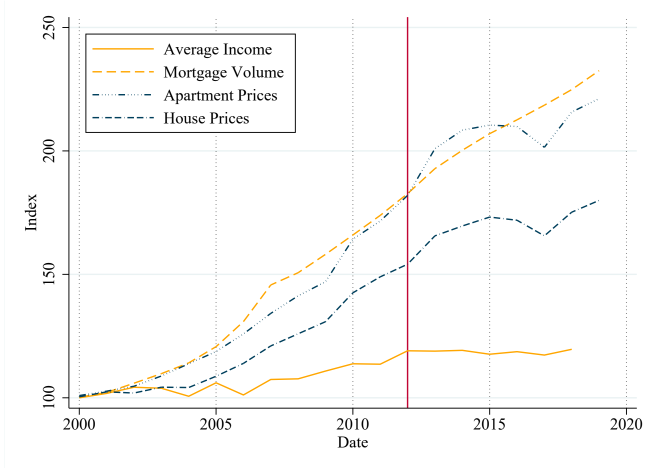

Figure 2.1: Income, Mortgages, and Prices in Switzerland

Note: This figure shows the time series of the average nominal gross income per household, total mortgages outstanding (gross claim) by banks in Switzerland measured in million CHF, as well as the price indices for privately owned apartments and single-family houses (Sources: FSO, 2020b; FSO, 2020a; SNB, 2020, and Wüest & Partner retrieved from SNB, 2020). The vertical line indicates the year 2012 when macroprudential policies were introduced in Switzerland. All series are indexed to a base of 100 in year 2000.

2.2 Macroprudential Policies in Switzerland between 2012 and 2016

In light of these concerns, Switzerland implemented several macroprudential policies with the goal of countering potentially damaging developments in the mortgage and real estate markets, and strengthening the resilience of the banking system. There are three relevant policies.

First, in June 2012, the Swiss Bankers Association tightened the down-payment requirements. Under the new requirements, home buyers need to finance at least 10% of the purchase price with "hard" equity capital, without drawing from second-pillar pension savings. This policy might have considerable effects on borrowing constraints, particularly for households with only little wealth outside their mandatory pension savings. Figure 2.2 shows the share of households in our data making such a withdrawal when transitioning into homeownership. The drop in the share of households withdrawing in 2012 illustrates that we can expect first effects at that time. While the average share was 22.4% before 2012, it decreased to 16.4% thereafter.[7]

Second, in June 2012, the Swiss Federal Council (the executive branch of the Swiss government) raised banks' capital requirements for originated mortgage loans with high LTV ratios: by January 2013, the risk-weights for the loan tranche exceeding an LTV ratio of 80% in-

[7] In addition to the restriction on the second-pillar withdrawals, the change in self-regulation included a maximal duration for the repayment of the loan. However, for most banks, this policy change likely had minor effects as their own requirements were already more restrictive (Behncke, 2022). This first self-regulation became effective in July 2012 with a transition period of 5 months. In June 2014, the self-regulation was revised and the maximal duration for repayment of the loan was shortened further.

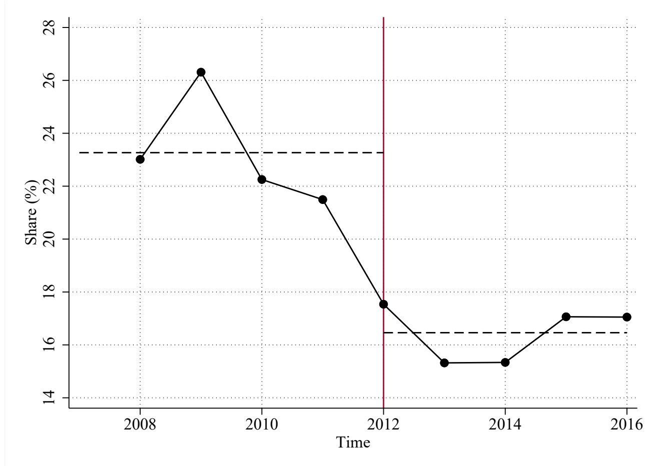

Figure 2.2: Annual Share of Households Withdrawing Second-Pillar Pension Savings

Note: The figure shows the annual share of households withdrawing from the second-pillar pension fund to finance their transition into homeownership. The share is calculated for all households renting in the previous year $t - 1$ and transitioning into homeownership in the current year t.

creased from 75% to 100%.

Third, in February 2013, the Federal Council activated the sectoral CCyB, requiring banks to hold additional common equity Tier 1 (CET1) capital on domestic residential mortgage loans. The CCyB initially amounted to 1% of a bank's relevant risk-weighted assets and was subsequently increased to 2% in January 2014. Table D1 provides a detailed timeline of all macroprudential policies from 2012 to 2016.

The tightening of the down-payment restrictions has the potential to directly impact borrowing constraints, particularly for households from low-wealth families that cannot access predeath bequests. Such a policy might enhance financial stability through three channels.

First, borrowers will have a stronger balance sheet. Only households with enough family-wide equity will be able to transition into homeownership. Households from richer families can tap into predeath bequests to finance homeownership. On the other hand, households from poorer families with weaker balance sheets will be credit constrained. Thus, only households from families with a relatively strong balance sheet can finance homeownership.

Second, borrowers face a lower financial burden resulting from debt service and rebuilding the second pillar pension fund, if households finance the necessary down-payment with predeath bequests instead of withdrawals from the 2nd pillar pension fund.

Third, borrowers get screened for financial responsibility. In particular, households with savings besides their second pillar pension fund are arguably financially more responsible and still able to transition into homeownership. Such a mechanism is consistent with evidence from the United States where Federal Housing Administration loans that received down-payment assistance are more likely to default (Lam et al., 2013).

3. Data

To analyze the effects of macroprudential policies on homeownership in more detail, we turn to administrative tax data containing information about the tenure status, intrafamily wealth transfers and other household characteristics. We exploit a unique administrative data set on the universe of individual tax records in the canton of Bern from 2007 to 2016.[8] Bern is Switzerland's second-largest canton, accounting for 12% of the total population (Federal Statistical Office, 2019a). It features both rural and urbanized areas. The data set contains information on 723,273 individual taxpayers, resulting in 5.7 million observations.

Every tax record includes detailed information on the taxpayer's income and wealth. It also comprises the taxpayer's marital status and age, tax deductions for childcare, and second- and third-pillar withdrawals for financing the down-payment necessary to acquire the principal residence.

An important feature of the data set is that it allows us to differentiate between two types of intergenerational wealth transfers, predeath bequests and inheritances. While predeath bequests can be planned for, the exact timing of inheritances is unpredictable in most cases. In particular, bequests can be timed to serve as additional equity in a planned purchase of a home.[9]

In Switzerland, married couples are required to file taxes jointly. So they are recorded as a single taxpayer, although we observe the income and age of each spouse. Our main analysis focuses on households of married couples.[10]

We observe the tax-assessed value of each household's real estate holdings, which allows us to follow the household's tenure status and identify a potential transition into homeownership.[11] We consider a household to be a homeowner if the tax-assessed value of one of its properties exceeds CHF 100,000. This threshold ensures that non-habitable properties, such as garages or small plots of land, are excluded.

In addition to the tax-assessed value, 39.4% of households transitioning into homeownership also report the purchase price of the acquired property. However, reporting is voluntary.

[8]Other authors use the same data but in different contexts. Galli and Rosenblatt-Wisch (2022) analyze the consumption and saving pattern of households, while Brülhart et al. (2021) study the effects of wealth taxation on reported wealth.

[9]In the canton of Bern, predeath bequests and inheritances to descendants, stepchildren or foster children, as well as predeath bequests between spouses or people in a registered partnership are tax-free (Grand Council of the Canton of Bern, 2014).

[10]We cannot identify individual taxpayers living in cohabitation and identify their homeownership status. However, in Section 8.1, we present a robustness check that relies on an alternative data set without such a restriction on the civil status.

[11]The tax-assessed value of a property is periodically updated by the tax authority and corresponds to approximately 70% of the market value (Steuerverwaltung Kanton Bern, 2020).

Table B1 in the Appendix compares the characteristics of reporting and non-reporting households in the year they acquire the property. While for many characteristics, the differences in means are statistically significant due to the large sample size, most of them are economically modest.

We observe each household's place of residence as a so-called MS-region. MS-regions are small labour market areas with a functional orientation towards centres. They are constructed by the Federal Statistical Office and feature a high degree of spatial homogeneity (FSO, 2019). There are 16 MS-regions in the canton of Bern. We match each household's MS-region with a price index for single family houses and a rent index for apartments, which allows us to construct local price-to-rent indices.[12]

Besides the administrative data from Bern, we use the nationwide SHP data and administrative tax data from the canton of Lucerne to assess the external validity of our results. We discuss these additional data sets and the results we derive from them in Section 8.2.3.

4. Descriptive Analysis

In this section, we provide descriptive evidence suggesting that borrowing constraints for households tightened with the introduction of macroprudential policies in 2012. We show that the share of households transitioning into homeownership decreased, starting in 2012, and that wealth and income became more important for making such a transition. Moreover, the evidence suggests that, since 2012, young and middle-aged households that transitioned into homeownership have relied more on predeath bequests compared to their peers who stayed renters or already owned a home.

4.1 Renter Households

To analyze the effects of macroprudential policies on homeownership, we are interested in the propensity of a household to transition into homeownership. Consequently, we restrict our focus to households of married couples and their potential children who initially rent and follow their tenure status over the subsequent years. This leaves us with 126,708 households and 780,955 observations. We identify 25,704 households who transitioned from renting to owning over our sample period from 2007 to 2016. A household is removed from the sample the year after it transitioned into homeownership.

[12]Figure C1 in the appendix shows all MS-regions in Switzerland. The indices are calculated based on transaction data for both rental units as well as residential property prices. We are grateful to Fahrländer Partner for providing price indices at the MS-region level.

Table 4.1: Summary Statistics

	(1) Mean	(2) Std. Dev.	(3) N
ΔHO	3.3	17.8	780,955
Received a Predeath Bequest (0/100)	4.9	21.5	780,955
Received a Predeath Bequest kCHF 1 to 10 (0/100)	0.8	8.7	780,955
kCHF 10 to 25	1.3	11.4	780,955
kCHF 25 to 50	0.8	8.9	780,955
kCHF 50 to 100	0.9	9.4	780,955
kCHF 100 or more	1.1	10.3	780,955
Received an Inheritance (0/100)	4.7	21.2	780,955
Received an Inheritance kCHF 1 to 10 (0/100)	1.0	10.1	780,955
kCHF 10 to 25	1.1	10.4	780,955
kCHF 25 to 50	0.9	9.3	780,955
kCHF 50 to 100	0.8	8.8	780,955
kCHF 100 or more	0.9	9.6	780,955
Purchase Price (in kCHF)	534.4	284.1	10,365
Age	51.6	15.9	780,955
Share of people with age ≤ 35 (0/100)	17.9	38.3	780,955
$35 < \text{Age} \leq 50$	34.1	47.4	780,955
$50 < \text{Age} \leq 100$	48.1	50.0	780,955
Lag Income (in kCHF)	85.6	44.8	780,955
Lag Wealth (in kCHF)	107.2	389.5	780,955
Has Children (0/100)	45.4	49.8	780,955
Second Pillar Withdrawal (0/100)	0.9	9.5	780,955
Third Pillar Withdrawal (0/100)	1.0	9.8	780,955
Price-to-Rent ratio (100 = 2007)	108.1	9.2	780,955

Note:The table shows the summary statistics of all variables for households renting in the previous year $t-1$. Variables with (0/100) in parentheses are dummy variables scaled from 0 to 100 to indicate percentages. ΔHO refers to the share of households who rented in year $t-1$ and transitioned into homeownership in year t. For both predeath bequests and inheritances, we consider transfers bigger or equal to CHF 1,000 in order to eliminate small transfers and reporting errors. Age refers to the mean age of the main taxpayer and the spouse. The base year of the price-to-rent ratio index is 2007.

Table 4.1 shows summary statistics for these households. On average, 3.3% of them transition into homeownership per year. The average purchase price of the acquired property is CHF 534,400. 4.9% of the households who are renting in the previous year receive a predeath bequest and 4.7% an inheritance. On average, the mean age of the main taxpayer and the spouse is 51.6 years. The average joint income is CHF 85,600 and the average financial wealth amounts to CHF 107,200.

4.2 Share of Households Transitioning into Homeownership

Figure 4.1 shows how the share of households transitioning into homeownership evolved over the sample period. It depicts the share of households who transitioned into homeownership in a given year t conditional on being renters in the previous year $t-1$ and remaining in the data set in t. The share decreased around the introduction of the macroprudential policies: it dropped by 11.8 percent from an average of 3.4% per year from 2007-2011 to an average of 3.0% per year from 2012-2016.[13]

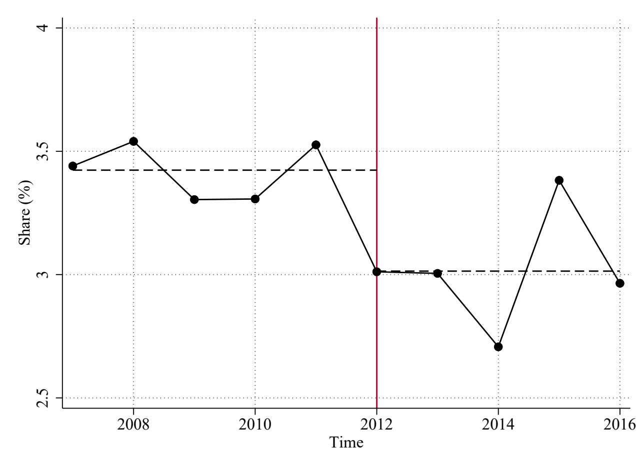

Figure 4.1: Annual Share of Households Transitioning into Homeownership in Bern

Note: The figure shows the share of households transitioning into homeownership for each year in the canton of Bern. The dashed lines indicate the mean before and after 2012. The vertical line shows the timing of the introduction of macroprudential policies in Switzerland. The sharp drop in 2014 and the following spike in 2015 are likely due to a cantonal tax reform in 2015 that gave an incentive to postpone property transactions from 2014 to 2015. The analogous figure for the canton of Lucerne can be found in Appendix Figure E1.

Next, we look at different age groups, as the propensity to transition into homeownership varies throughout a household's life-cycle. For instance, the median age for households transitioning into homeownership is 40 years, while the first and third quartiles are 34 and 58 years, respectively.[14]

Variation in homeownership timing may have various reasons. Besides career and family planning, financial considerations probably play an important role as accumulating the required savings for making a down-payment takes time. Thus, borrowing constraints likely vary with age.

To explore this point, we split the households into three different age categories according to the mean age between the main taxpayer and the spouse: up to 35 years (category 1), 36-50 years (cat. 2), or above 50 years (cat. 3).

Figure 4.2 shows the share of households transitioning into homeownership for each of the

[13]The low share in 2014, followed by a rebound in 2015, is likely due to a reform in tax law in the canton of Bern, which was passed in 2014 and came into force in 2015. This reform reduced the taxes on property purchases (Kanton Bern, 2020). Households, therefore, had an incentive to postpone property purchases from 2014 to 2015. Note that the general pattern and the pronounced drop around 2012 in the rate of transition into homeownership also occur in the canton of Lucerne, as shown in the Appendix Figure E1; in fact, the drop there is even larger (in both absolute and relative terms).

[14]For details, we show the age histogram for households transitioning into homeownership and for all households in Appendix Figure A1 for Bern and in Appendix Figure E2 for Lucerne.

three age categories. For the youngest households in age category 1, the mean share of households transitioning into homeownership decreases slightly from 5.2% before 2012 to 4.9% thereafter. This corresponds to a decrease of 5.8% in the probability of transitioning into homeownership. The decrease is much more pronounced in the second age category, where the mean share of households transitioning into homeownership falls from 5% before 2012 to 4.5% thereafter, which translates to a decrease of 10%. In age category 3, the mean share of households transitioning into homeownership before 2012 is 1.6% compared to 1.4% thereafter. This corresponds to a decrease of 12.5%. Note that the share of households transitioning into homeownership is more volatile in the younger age categories as there are fewer observations.[15]

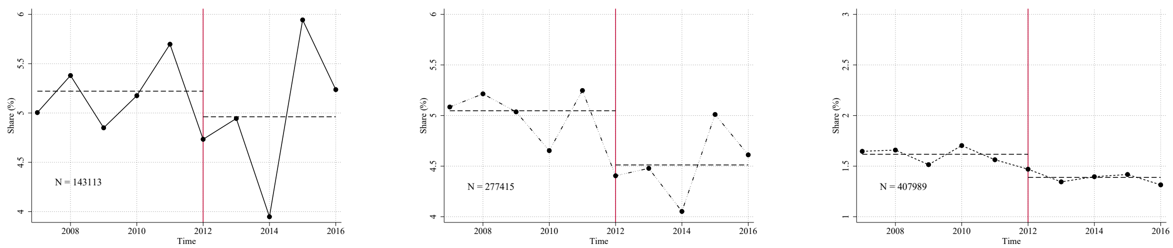

(A) Age ≤ 35 (B) $35 <$ Age ≤ 50 (C) $50 <$ Age ≤ 100

Figure 4.2: Share of Households Transitioning into Homeownership by Age Category

Note: The figure shows the share of households transitioning into homeownership in the canton of Bern, separately for the three age categories. Dashed horizontal lines without symbols represent the average before and after 2012. "N" indicates the total number of households renting in the previous year $t-1$ per age category.

4.3 The Role of Income and Wealth

Next, we look at how income and wealth relate to the propensity to transition into homeownership. Figure 4.3 illustrates that, after 2012, income and especially wealth became more important. It shows the mean share of households transitioning into homeownership before and after 2012, conditional on their lagged income quintile (Panel a) and lagged wealth quintile (Panel b). The numbers above the bars indicate relative changes in percent.

For lagged income quintiles, there are two main observations. First, high-income households are more likely to transition into homeownership, both before and after the introduction of the macroprudential policies in 2012. Second, the relative change in the share of households transitioning into homeownership before versus after 2012 is generally more pronounced for low- and middle-income households than for high-income households. For example, the share in the second income quintile was 1.51% before 2012 and 1.21% thereafter, corre-

[15]Similar figures for the canton of Lucerne can be found in Appendix Figure E3. There as well, the decrease is most pronounced for the middle age category.

sponding to a relative change of -20.2%. In contrast, the share of households transitioning into homeownership in the fifth income quintile was 7.81% before 2012 and 6.96% thereafter, corresponding to a smaller relative change of -10.9%.

For lagged wealth quintiles, we observe a similar pattern. First, wealthy households are more likely to transition into homeownership. Second, the relative change in the share of households transitioning into homeownership before versus after 2012 is stronger for the low wealth quintiles. For instance, in the bottom quintile, the relative change in the share of households transitioning into homeownership is -37.2%, while in the top quintile, it is just -0.3%, and the pattern is monotonic in between.

In summary, these results suggest that income and wealth have become more important in enabling households to transition into homeownership.

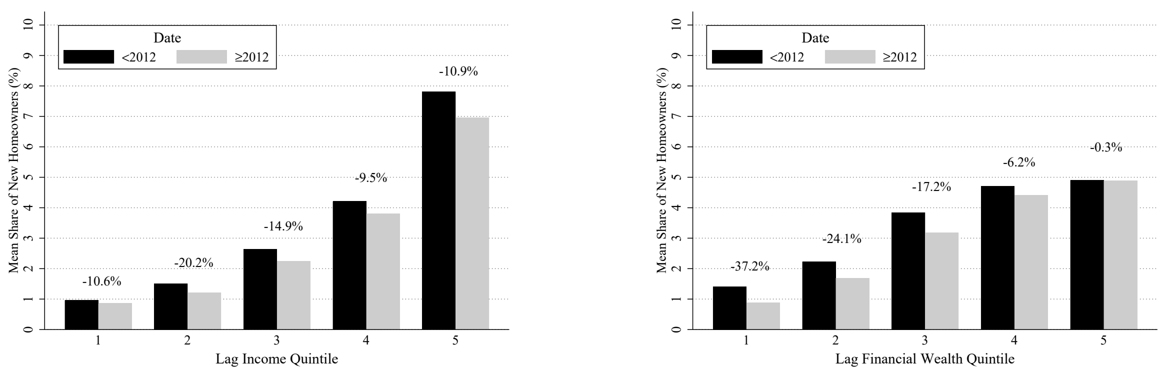

(A) Conditional on Lag Income

(B) Conditional on Lag Financial Wealth

Figure 4.3: Share of Households Transitioning into Homeownership by Income and Wealth

Note: The figure shows the share of households transitioning into homeownership conditional on the position in the lag income and wealth distributions for the periods before and after 2012. The percentage at the top of each bar indicates the relative change in the share before versus after 2012. All quantiles are calculated over all households renting in the previous year $t-1$.

Figure 4.4 focuses on second-pillar withdrawals conditional on the household's position in the wealth distribution. It suggests that the restriction on withdrawals from pension savings introduced in 2012 primarily affects low-wealth households. It reveals that low-wealth households who transitioned into homeownership strongly relied on withdrawals from their pension savings before 2012 and that the propensity to use such withdrawals fell significantly thereafter.

4.4 The Role of Intergenerational Wealth Transfers

Households with insufficient wealth may rely on intergenerational wealth transfers to transition into homeownership. Predeath bequests might be particularly effective for that purpose as they can be timed, in contrast to inheritances that are relatively unpredictable.

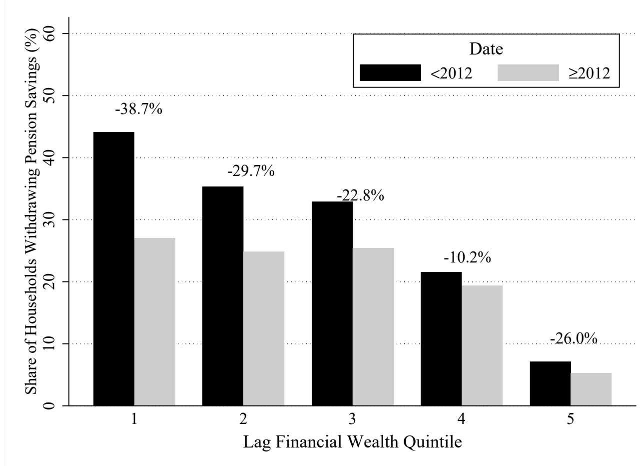

Figure 4.4: Share of Households Transitioning into Homeownership using Pension Saving Withdrawals by Wealth Quintile

Note: The figure shows the share of households withdrawing vs. not withdrawing second-pillar pension savings to finance the transition into homeownership, conditional on their position in the wealth distribution. The sample conditions on all households renting in the previous year $t-1$ and transitioning into homeownership in t. The percentage at the top of each bar indicates the relative change in the share before versus after 2012. The wealth quintiles are calculated over all households renting in $t-1$.

Figure 4.5 displays the share of households who receive a wealth transfer, conditional on their age category and tenure status.[16] A wealth transfer refers either to a predeath bequest or an inheritance received during the current year t or the previous year $t-1$. We consider three different tenure statuses. The first tenure status, "ΔHO", represents households who were renting in the previous year $t-1$ and transitioned into homeownership in the current year t. The second, "Staying Renter", refers to households who were renting in both years $t-1$ and t. The third, "Staying Owner", refers to households who owned a home in both years $t-1$ and t.

The upper panels show the share of households receiving a predeath bequest. There are three key observations. First, across all age categories, households who transition into homeownership are more likely to receive a predeath bequest than households who stay renters or already own a home. Thus, predeath bequests are commonly used to finance homeownership. Second, young and middle-aged households who transition into homeownership receive a predeath bequest more often than their older counterparts. This is probably because they have had less time to accumulate the wealth necessary to finance the purchase of a home. Third, the share of young and middle-aged households who transition into homeownership and receive a predeath bequest increases sharply after 2012. This is particularly the case for households in the youngest age category 1. The average share for them increases from 28% before 2012 to 36% thereafter. This observation suggests that borrowing constraints

[16] For more details about the distribution of the transfers over time, see Appendix Figure F1.

got tighter after 2012 and, in response, households started relying more on intergenerational wealth transfers.[17]

The lower panels show the share of households receiving an inheritance. They reveal three noteworthy observations. First, in age categories 1 and 2, the share of households receiving an inheritance does not vary with tenure status. Second, in the older age category 3, the share is higher for households who transition into homeownership than for those who stay renters or already own a property. Third, there is no systematic increase in the share of households receiving an inheritance around 2012. Overall, these three observations suggest that, due to their unpredictable nature, inheritances are less well suited than bequests to overcome the tighter borrowing constraints after 2012.

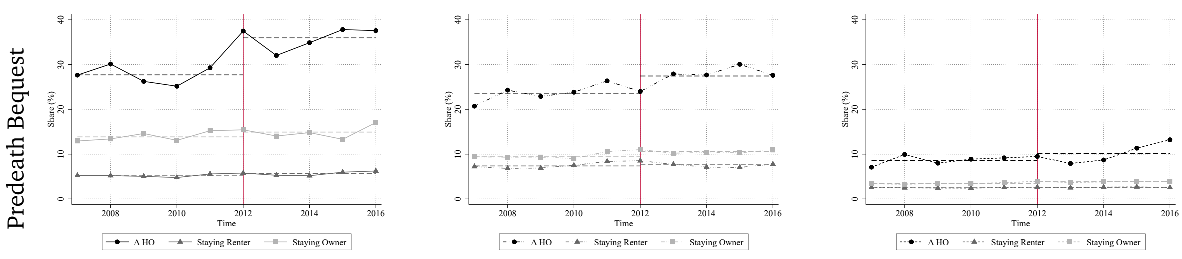

(A) Age ≤ 35 (B) $35 < \text{Age} \leq 50$ (C) $50 < \text{Age} \leq 100$

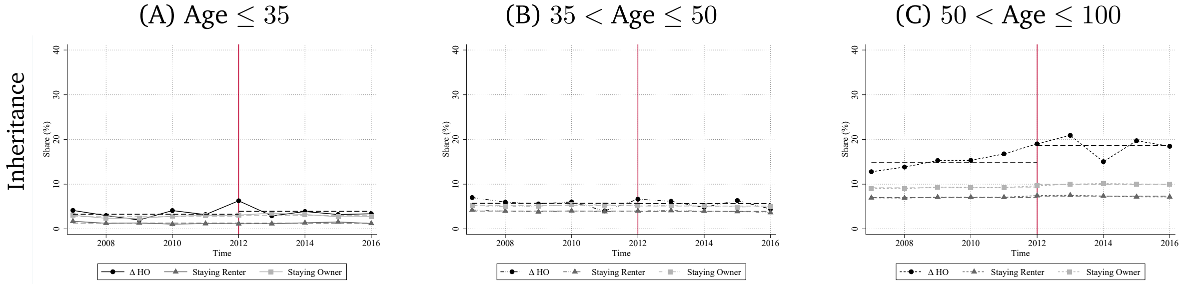

(D) Age ≤ 35 (E) $35 < \text{Age} \leq 50$ (F) $50 < \text{Age} \leq 100$

Figure 4.5: Share of Households Receiving a Transfer by Tenure Status and Age Group

Note: The figure shows the mean share of households receiving a transfer (predeath bequest or inheritance) conditional on their tenure status and age category. The dashed lines without symbols indicate the means before and after 2012.

To summarize, we observe a pronounced drop in the share of households transitioning into homeownership around 2012—especially for young and middle-aged households. In addition, wealth and income have become more important after 2012. Young and middle-aged households who accumulated less wealth often receive predeath bequests. Around 2012, the share of households transitioning into homeownership that recieved bequests increased and has stayed at a higher level since then. Thus, it seems that, following the introduction of

[17]Note that these results are robust to a different lag structure where the transfer dummy is equal to one if the household has received any transfer (CHF $> 1,000$) in the last two years and zero otherwise.

the macroprudential policies, young and middle-aged households have had to rely more on family wealth to overcome the tighter borrowing constraints.

Next, we provide a simple theoretical framework to illustrate that, as down-payment requirements increase, we expect the propensity of low-income households to transition into homeownership to decrease, but this decrease should be weaker for households that are able to access their family wealth via a predeath bequest. The framework will motivate our main regressions in Section 6.

5. Theoretical Framework

In this section, we provide a simple theoretical framework, based on a modified version of **?**. The goal is to explain how tighter borrowing constraints may affect the incentives for family wealth transfers to finance homeownership.

5.1 Setup and Calibration

There are many households. Each household lives for two periods, $t = 1$ and $t = 2$, and earns a constant income y in every period. However, income is heterogeneous across households. We assume that income in the second period also incorporates family wealth. In the first period, households choose their consumption. Additionally, they have the possibility to save using a risk-free bond b. In this model, we abstract from lending and assume that access to homeownership is based on a household's amount of savings. In the second period, households use their second-period income and savings $(1 + r)b$ to finance their consumption c_2.

The model features a down-payment restriction b^*. If a household saves more than b^*, it is considered to be a homeowner and gets a homeownership utility bonus Φ.[18] If the household saves less than b^*, the household is a renter. This status is fixed across both periods.

Additionally, households can access family wealth from the second period via transfers TR (e.g., predeath bequests). However, transfers have additional utility costs, such as liquidation costs when family assets are illiquid, transaction costs when legal documents are required, and psychological costs related to the discomfort of asking the family for money. These utility costs are assumed proportional to the transfer size and represented by the parameter λ.

The household maximizes its lifetime utility, which is the sum over the per-period utilities

[18]The assumption of a utility gain for homeownership is common in the literature. The higher utility may represent that housing serves as an important savings instrument (Goodman and Mayer, 2018) and provides consumption insurance (Lustig and Van Nieuwerburgh, 2005). Other studies find that homeownership is associated with an increase in personal well-being (White and Schollaert, 1993).

plus the eventual utility gain of homeownership net of the costs of transfers,

$$\sum_{t=1}^{2} u(c_t) + \mathbb{I}_{\Phi}\Phi - \mathbb{I}_{TR}\lambda TR, \text{ with } t = 1,2, u' > 0, \text{ and } u'' < 0, \text{ s.t.}$$

$$c_1 = y_1 - b + TR$$

$$c_2 = y_2 + b - TR$$

$$b \geq 0$$

$$0 \leq TR < y_2$$

$$\mathbb{I}_{\Phi} = \begin{cases} 1 & \text{if } b \geq b^* \\ 0 & \text{else} \end{cases}$$

$$\mathbb{I}_{\text{TR}} = \begin{cases} 1 & \text{if } TR > 0 \\ 0 & \text{else} \end{cases}$$

Notice that households face a trade off between consumption smoothing, the utility benefit of housing, and the disutility from the potential use of transfers.

We use a simple calibration of the model to illustrate how households solve this trade off in two different scenarios. In both scenarios, we study the behavior of a continuum of households that vary in their income $y_1 = y_2 \in (0.1, 1)$. We assume the following parameters: $b^* = 0.2$, $\Phi = 0.3$ and $\lambda = 0.8$. We abstract from discounting of future utility and assume that interest rates are zero. We specify the utility function to be the log of consumption.

5.2 Scenario 1: No transfers versus transfers

In the first scenario, we compare the optimal decisions of households that have access to family wealth via transfers to those of households that do not. Panel (a) of Figure 5.1 shows the results.

If households have no access to transfers, only households with income equal or above to 0.39 are homeowners. For the ones with lower income, the benefit of homeownership is lower than the utility cost of unequal consumption.

In contrast, when households have access to transfers, households are homeowners if their total income is larger than 0.14. Thus, the number of households that finance the down-payment and are homeowners increases relative to the situation without transfers. Households use transfers to finance homeownership as long as the benefit of homeownership is bigger than the disutility of unequal consumption plus the disutility of transfers. Note that the disutility of unequal consumption is partially offset by transfers.

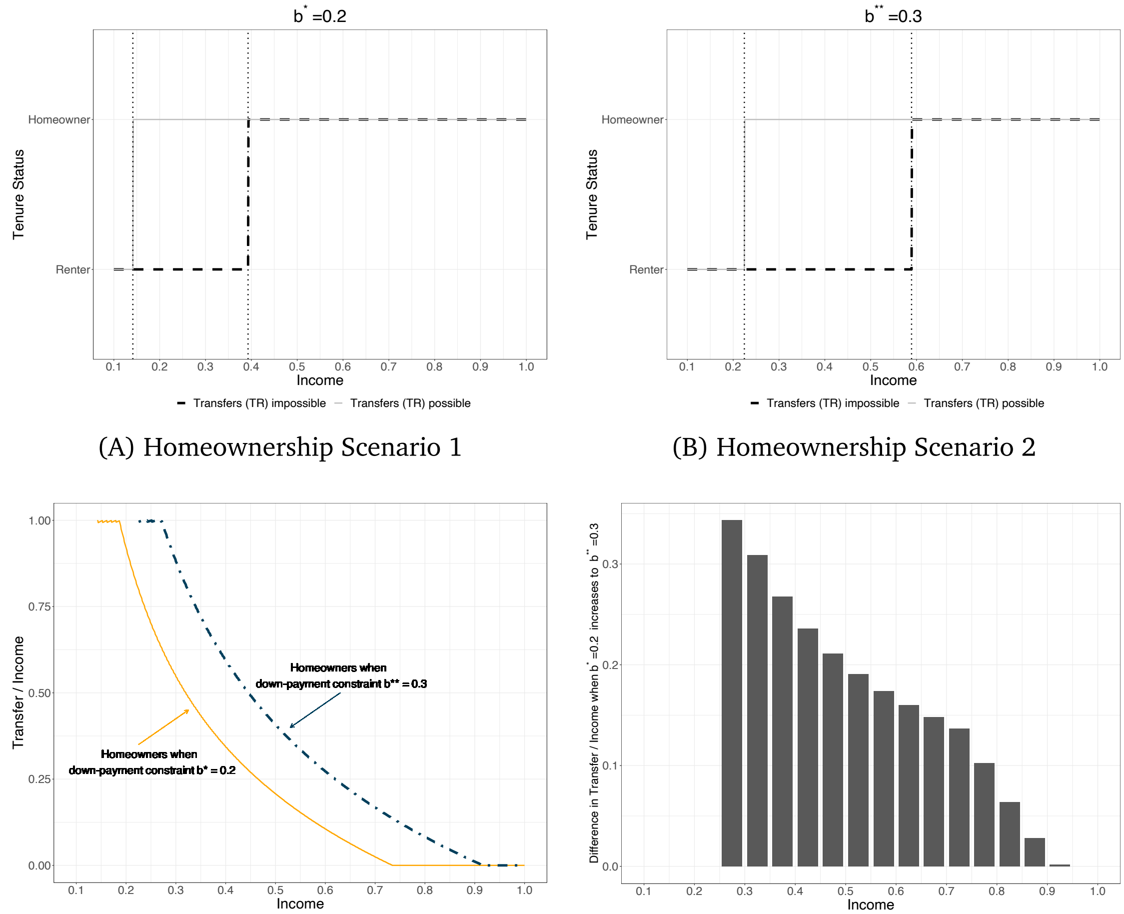

(A) Homeownership Scenario 1

(B) Homeownership Scenario 2

(C) Transfers Relative to Income

(D) Change in Transfers Relative to Income when Down-Payment Restrictions Tighten

Figure 5.1: Results from Illustrative Calibration

Note: The figure shows the results of the numerical calibration for the two scenarios described in the text for households with different levels of income. Panel (a) displays whether a household optimally chooses to be a renter or homeowner, conditional on their income. To be a homeowner, the household must finance the down-payment restriction with savings. Panel (b) is the same graph but households are faced with a higher down-payment constraint ($b^* < b^{**}$). Panel (c) plots the size of used transfers relative to the households' income (y_2) for both scenarios. Panel (d) plots the average differences in the share of transfers after the increase in the down-payment restriction for each income bin of size 0.05.

5.3 Scenario 2: Increase down-payment restriction

In the second scenario, we increase the down-payment restriction from $b^* = 0.2$ to $b^{**} = 0.3$. Such an increase could result from the introduction of macroprudential policies. Again, we compare households that have access to transfers to those that do not.

The resulting decisions in this scenario are shown in panel (b) of Figure 5.1. Now, only households with an income above or equal to 0.59 choose to save enough for the down-payment if they do not have access to transfers. Instead, if households have access to transfers, they choose homeownership if their income is higher or equal to 0.22. The width of the income range in which households only become homeowners when they have access to transfers has

widened compared to the scenario in panel (a).
Panel (c) shows the size of the used transfers as a share of the household's income in the two scenarios. If households stay renters, they do not use transfers as they can perfectly smooth consumption. If households become homeowners, low-income households use larger transfers relative to their income than high-income households. Moreover, when down-payment restrictions tighten, transfers increase for a given income level.
Panel (d) illustrates that, when down-payment restrictions get tighter, transfers become more important for low-income households. It shows how the ratio of transfers relative to income changes for different income levels when down-payment restrictions tighten from $b^* = 0.2$ to $b^{**} = 0.3$. Those households with lower income level increase their transfers more to finance homeownership after down-payment restrictions tighten.
In sum, the simple calibrated model illustrates how transfers become more important for enabling homeownership when down-payment requirements are higher. Moreover, households that finance the down-payment also increase the amount of transfers used, and this increase is stronger for low-income households.

6. Empirical strategy

In an ideal experimental set-up, we could measure the causal effect of the introduction of macroprudential policies on homeownership and borrowing constraints directly, as households would be randomly assigned to a treatment group subject to tighter borrowing constraints and a control group facing no policy change. However, in our set-up, the introduction of macroprudential policies affects all households at the same time.
Thus, we exploit the equilibrium relationship outlined in the theoretical framework and use intergenerational wealth transfers to identify a potential change in borrowing constraints. More precisely, we compare the impact of intergenerational wealth transfers on the probability of transitioning into homeownership (extensive margin) and the price of the acquired home (intensive margin) before and after 2012. This strategy is similar to Blickle and Brown (2019), who also rely on intergenerational wealth transfers to identify borrowing constraints.
Our main analysis focuses on the extensive margin. It exploits the panel structure of the data and controls for various potential confounds by applying the following linear probability model, conditional on all households renting in the previous year $t-1$:

$$\Delta\text{HO}_{i,t} = \alpha_{m(,t)} + \beta_1\text{TR}_{i,t} + \beta_2 MP_t + \beta_3\text{TR}_{i,t} \times \text{MP}_t + \beta_4\text{HH}_{i,t} + \epsilon_{i,t}\,. \tag{6.1}$$

The dependent variable, $\Delta\text{HO}_{i,t}$, is a dummy indicating whether household i transitions into homeownership. It takes the value 100 if the household buys a home in year t and is 0 otherwise. Once a household has transitioned into homeownership, it exits the sample in the following period.

The main independent variable of interest is $\mathrm{TR}_{i,t}$, a dummy indicating whether household i receives a wealth transfer. This dummy takes the value of 1 if the household receives a transfer in the current year t or the previous year $t-1$. We include transfers in the previous year to take into account that the decision to buy a home might not be immediate, or that the transfer occurs just before the end of the reporting period for taxes. As the theoretical framework illustrates, borrowing constrained households tend to rely on intergenerational wealth transfers to meet the down-payment requirement when transitioning into homeownership. Hence, we interpret β_1 as a measure of how tight borrowing constraints are.

Another important independent variable is the dummy MP_t, which takes the value of 1 between 2012 and 2016 when the macroprudential policies are in force, and 0 otherwise. Its coefficient, β_2, indicates how the introduction of the macroprudential policies affects the probability of transitioning into homeownership of households that do not rely on transfers.

We interact the indicator for the transfer with the dummy MP_t. Again, as the theoretical framework suggests, households' reliance on wealth transfers increases when borrowing constraints tighten. Thus, a positive β_3 signals that borrowing constraints became tighter after 2012.

We use two specifications of the model to discriminate between the effects of inheritances and bequests. In the first specification, $\mathrm{TR}_{i,t}$ indicates that the household received a predeath bequest, while in the other, it indicates that it received an inheritance.

To estimate the effects of different transfer sizes, we also use alternative versions of the above specifications and replace the dummy $\mathrm{TR}_{i,t}$ with six categorical variables, indicating the following transfer sizes: CHF 1,000–9,999; CHF 10,000–24,999; CHF 25,000–49,999; CHF 50,000–99,999; and CHF 100,000 or higher. Receiving no transfer is always the base category. This categorical variable also takes into account potential non-linear effects of transfers.

$\mathrm{HH}_{i,t}$ is a vector of control variables at the household level. It includes three measures for the household's financial strength as well as demographic variables. The first two measures for financial strength are the household's position in the income and wealth distribution at first observation in ventiles, i.e., 20 categorical dummies each for income and for wealth. Another measure is the household's log income, which we lag by one period to avoid endogeneity bias. Demographic control variables include the mean age of the main taxpayer and the spouse (rounded to the nearest integer to allow for the use of age fixed effects), and a lagged dummy indicating whether any children live in the household at $t-1$.

α_m represent MS-region fixed effects. They control for heterogeneous local housing market conditions as well as for other local characteristics, such as the structure of the local banking market. We also estimate a version of the model with year $\times$ MS-region fixed effects, $\alpha_{m,t}$, which further capture changes in market conditions and local characteristics over time but also absorb the MP_t-dummy.

Adding these fixed effects along with the household characteristics captures various potential confounds. For instance, it captures that older households had more time to build up wealth and, therefore, are less likely to rely on transfers; or that households with children might have different housing preferences than households without children.

We cluster standard errors at the household level. While filed taxes of married couples are recorded as single taxpayers, we observe the income and age of each spouse separately. Therefore, we allow errors to be correlated at the household level.

Besides the extensive margin, we also study the intensive margin. That is, we estimate an analogous model but use as the dependent variable the log purchase price of the property to estimate how the introduction of the macroprudential policies affected the average purchase price. In this specification, we cluster the standard errors at the region $\times$ year level to account for potential correlation at the region$\times$year level.

7. Results

In this section, we present and interpret the estimation results. At the extensive margin, we show that households have, on average, a lower probability to transition into homeownership after the introduction of macroprudential policies. However, this is not the case for households that could tap into family wealth via predeath bequest. Their probability of transitioning into homeownership stays roughly the same. At the intensive margin, the results are similar. After 2012, receiving a predeath bequest as well as having more wealth increases the purchase price of the acquired property.

7.1 Extensive margin

Table 7.1 shows the effects of predeath bequests and inheritances on a household's probability of transitioning into homeownership. In Columns (1), (2), (4) and (5), we provide the results for the model specified in Equation 6.1 using an indicator for wealth transfers and including either MS-region fixed effects or year $\times$ MS-region fixed effects. In Columns (3) and (6), we focus on categorical transfer sizes.

Column (1) displays the main result. Before 2012, receiving a predeath bequest increases the probability of transitioning into homeownership by 11.99 percentage points. After 2012, the probability of transitioning into homeownership decreases by 0.45 percentage points but only for households that receive no predeath bequests. Households that could draw on predeath bequests have a 0.35 (= 0.8 - 0.45) percentage point higher transition probability after the introduction of the macroprudential policies, although this combined effect is not statistically significantly different from zero (p =0.35).

Column (2) shows the same regression with year $\times$ MS-region fixed effects, which leaves the estimated effect of bequests and the change in this effect after 2012 essentially unchanged. Thus, after the introduction of the macroprudential policies, predeath bequests have become more important for households to transition into homeownership.

Table 7.1: Effect of Transfer on the Probability of Transitioning into Homeownership

	Predeath Bequest			Inheritance		
	(1)	(2)	(3)	(4)	(5)	(6)
	ΔHO	ΔHO	ΔHO	ΔHO	ΔHO	ΔHO
Transfer=1	11.99***	12.00***		2.40***	2.41***	
	(0.27)	(0.27)		(0.17)	(0.17)	
$\text{MP}_{t,12}$=1	-0.45***			-0.36***		
	(0.04)			(0.04)		
$\text{MP}_{t,12}$=1 × Transfer=1	0.80**	0.77**		0.23	0.20	
	(0.38)	(0.38)		(0.24)	(0.24)	
kCHF 1 to 10			0.30			-0.00
			(0.35)			(0.26)
kCHF 10 to 25			2.77***			-0.10
			(0.35)			(0.26)
kCHF 25 to 50			8.58***			1.09***
			(0.59)			(0.34)
kCHF 50 to 100			16.82***			3.12***
			(0.69)			(0.44)
kCHF 100 or more			30.35***			9.15***
			(0.77)			(0.56)
$\text{MP}_{t,12}$=1 × kCHF 1 to 10			-0.38			-0.04
			(0.47)			(0.37)
$\text{MP}_{t,12}$=1 × kCHF 10 to 25			0.62			0.29
			(0.50)			(0.37)
$\text{MP}_{t,12}$=1 × kCHF 25 to 50			2.10**			-0.15
			(0.87)			(0.48)
$\text{MP}_{t,12}$=1 × kCHF 50 to 100			2.15**			0.08
			(1.00)			(0.63)
$\text{MP}_{t,12}$=1 × kCHF 100 or more			0.35			-0.05
			(1.06)			(0.78)
Year FE × MS Region FE	No	Yes	Yes	No	Yes	Yes
Main Controls	Yes	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	780,955	780,955	780,955	780,955	780,955	780,955
$\bar{y}$	3.28	3.28	3.28	3.28	3.28	3.28

Note:The table shows the effect of a transfer (predeath bequest or inheritance) on the probability of transitioning into homeownership. In Columns (1), (2), (4) and (5), "Transfer" is a dummy equal to one if a household received a transfer of at least CHF 1,000 in year t or $t-1$, and zero otherwise. Columns (3) and (6) use categorical variables for different transfer sizes (omitted category: households who receive no transfer). Columns (1) and (4) are estimated without year but MS region fixed effects while columns (2), (3), (5) and (6) include year × MS Region fixed effects. Main controls include lag income, having children, financial wealth and income ventiles at first observation. $MP_{t,12}$ is a dummy indicating when macroprudential policies are active during our sample period (2012 to 2016). ΔHO indicates whether a household transitioned into homeownership. Regressions are calculated for households renting in the previous year $t-1$. $\bar{y}$ is the mean of the dependent variable. * p<0.10, ** p<0.05, *** p<0.010. Standard errors are clustered at the household level. In the Appendix Table I1 we provide results for different specifications of the model.

The estimates in Column (3) show that the effect of predeath bequests increases with their size. Receiving under CHF 10,000 has no significant effect on the probability of transitioning into homeownership, while a predeath bequest of CHF 100,000 or more increases this probability by 30.4 percentage points (relative to not receiving a transfer). After 2012, predeath bequests between CHF 25,000-50,000 or CHF 50,000-100,000 have a significantly stronger effect on the probability of transitioning into homeownership. Notice that predeath bequests in that range are most likely to alleviate the additional borrowing constraints for households on the margin of being able to make a down-payment after the introduction of the restriction on pension savings withdrawals. Overall, the results suggest that, after 2012, borrowing constraints became tighter.

Next, we turn to inheritances, which are less easily planned for than predeath bequests. In Column (4), the estimates reveal that receiving an inheritance increases the probability of transitioning into homeownership by 2.4 percentage points. This effect is smaller than the one of receiving a predeath bequest but still significant. The effect is not significantly changed after the introduction of the macroprudential policies, although the point estimate is also positive, like for predeath bequests. Adding year $\times$ MS-region fixed effects does not alter these coefficients (Column (5)). In Column (6), we see that larger inheritances increase the probability of transitioning into homeownership but these effects are not significantly different in the post-2012 period.

As discussed, inheritances are not only less easily planned for than predeath bequests but also tend to occur later in the life-cycle of a household. The results confirm that these features make them less suitable for overcoming tighter borrowing constraints—particularly for young and middle-aged households who are most affected by the tighter borrowing constraints after 2012. Thus, from now on, we focus exclusively on the effects of predeath bequests.

In Table 7.2, we present several modifications to our baseline model. For better comparability, Column (1) shows our baseline estimates with year $\times$ MS-region fixed effects. Column (2) interacts predeath bequests with the age categories. Columns (3) and (4) control for real estate prices to avoid a potential confound.

Column (2) reveals that receiving a predeath bequest has a stronger effect on the probability of transitioning into homeownership for households in the younger two age categories than for those where the mean age between the main taxpayer and the spouse is 50 and older (the omitted category for the interaction terms). This holds over the sample period as a whole, as indicated by the economically and statistically highly significant interaction terms of the dummy for receiving a predeath bequest with the dummies for being in the youngest age category 1 or the middle-aged category 2. The effect sizes are +15.3 percentage points and +7.4 percentage points, respectively. After 2012, the effect of receiving a predeath bequest increases by an additional 3.0 percentage points (or 19%) for households in age category 1, and by 2.1 percentage points (or 29%) for households in age category 2 relative to households where the mean age is 50 and older.[19]

[19] The table also shows that the simple interactions of being in the post-2012 period with the two age category dummies are negative, meaning that for households with mean age below 50 years and without a predeath bequest, the probability of becoming a homeowner decreases post-2012 relative to older households.

A potential concern for our analysis is that the increased importance of predeath bequest could reflect not only the tightening of borrowing constraints due to macroprudential policy but also the general increase in Swiss real estate prices (see Figure 2.1 above). In this case, we would expect that predeath bequests have a stronger effect in regions with a higher price-to-rent ratio or price index. To test whether this is the case, in Columns (3) and (4), we interact the dummy for receiving a predeath bequest with the lag price-to-rent ratio and the lag of the price index, respectively.

These interaction terms are small and insignificant for the lag price-to-rent ratio or significantly negative (at the 10% level). More importantly, the estimate of receiving a predeath bequest interacted with the MP dummy remains similar to our baseline results in Column (1). Consequently, the increased importance of predeath bequests after 2012 appears to be due to tighter borrowing constraints rather than to the upward trend in real estate prices.

In Table G1 in the Appendix, we use an alternative approach to control for the effect of real estate price dynamics. We split the sample in half based on either the level or the growth rate of the regional price-to-rent ratio. If the stronger effect of predeath bequests was solely due to an increase in real estate prices, we would expect our interaction coefficient of interest to be larger in regions with an above-median price-to-rent ratio. However, we find it to be larger in regions with relatively low price-to-rent ratios.

7.2 Intensive Margin

Next, we turn to the intensive margin. Evidence from the United States and Italy shows that households tend to buy larger homes after they received a wealth transfer (Engelhardt and Mayer, 1998; Guiso and Jappelli, 2002). We use the purchase price reported by the households transitioning into homeownership as a proxy for the size and quality of a property.[20]

Table 7.3 analyzes how receiving a predeath bequest and a household's position in the wealth and income distributions affect the price of the purchased property, and how this changed after 2012. Our results suggest that a positive effect of wealth transfers on the intensive margin also exists in Switzerland, and that this effect got stronger after 2012.

All models are estimated conditional on the household reporting the purchase price. The dependent variable, the purchase price of the property, is in logs. As independent variables, we use a dummy for having received a predeath bequest, and dummies indicating whether the household's income or wealth is above the median. We use lagged income and wealth, as

[20]For those not reporting the purchase price, one could consider using the tax-assessed value available in the tax data of all households. As noted earlier, this is supposed to correspond to about 70% of the property's market value. However, anecdotally, these tax-assessed values contain quite a bit of variation around this target, so that we are reluctant to use them for the intenive-margin analysis.

Table 7.2: Heterogeneity across Age Categories and Effect of Real Estate Prices

	Baseline	Age Categories	Lag Price-to-Rent Ratio	Lag Price
	(1)	(2)	(3)	(4)
	ΔHO	ΔHO	ΔHO	ΔHO
Predeath Bequest=1	12.00***	4.44***	12.51***	14.81***
	(0.27)	(0.39)	(1.95)	(1.62)
$\text{MP}_{t,12}$=1 × Predeath Bequest=1	0.77**	-0.88*	0.82*	1.29***
	(0.38)	(0.52)	(0.43)	(0.49)
Predeath Bequest=1 × Age Category=1		15.29***		
		(0.76)		
Predeath Bequest=1 × Age Category=2		7.37***		
		(0.54)		
$\text{MP}_{t,12}$=1 × Age Category=1		-0.41***		
		(0.11)		
$\text{MP}_{t,12}$=1 × Age Category=2		-0.52***		
		(0.09)		
$\text{MP}_{t,12}$=1 × Predeath Bequest=1 × Age Category=1		2.97***		
		(1.07)		
$\text{MP}_{t,12}$=1 × Predeath Bequest=1 × Age Category=2		2.12***		
		(0.75)		
Predeath Bequest=1 × Lag Price-to-Rent Ratio			-0.00	
			(0.02)	
Predeath Bequest=1 × Lag Price				-0.02*
				(0.01)
Year FE × MS Region FE	Yes	Yes	Yes	Yes
Main Controls	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Observations	780,955	780,955	780,955	780,955
$\bar{y}$	3.28	3.28	3.28	3.28

Note: The table shows the effect of receiving a predeath bequest on the probability of transitioning into homeownership. Column (1) shows the baseline results. We interact the dummy for receiving a predeath bequest with the age category in Column (2). Columns (3) and (4) show the effect of receiving a predeath bequest when controlling for the interaction of the predeath bequest with the lag price-to-rent ratio as well as the lag price index. Main controls include lag income, having children, financial wealth and income ventiles at first observation. $MP_{t,12}$ is a dummy indicating when the macroprudential policies are active during our sample period (2012 to 2016). ΔHO indicates whether a household transitioned into homeownership. Regressions calculated for all households renting in the previous year $t-1$. $\bar{y}$ is the mean of the dependent variable. * p<0.10, ** p<0.05, *** p<0.010. Standard errors are clustered at the household level.

the current levels might be affected by the home purchase. We include the same controls and fixed effects as in the previous extensive margin analysis.[21]

Column (1) reveals that receiving a predeath bequest has a positive effect on the purchase price of the new property. Households who received a predeath bequest spend, on average, 9% more on their new home than households who buy without having received a predeath bequest. Moreover, the difference increases by another 3 percentage points after the introduction of the macroprudential policies in 2012. However, the effect is less precisely estimated than the one at the extensive margin, due to the relatively small sample size. In Appendix Table H1, we show that the effect after 2012 is significant and positive mostly for transfers between CHF 50,000 and 100,000, which is also the category for which extensive-margin effects were the largest.

Column (2) reveals that households whose wealth is above the median acquire homes which are, on average, 14% more expensive than those of households whose wealth is below the median. This effect significantly increases by an additional 5 percentage points after the

[21] An exception is that we do not show a model estimated without year fixed effect. As prices of real estate increase over time, in the model without year fixed effects the time dummy for macroprudential policies is necessarily positive.

introduction of the macroprudential policies.

Column (3) shows that households whose income is above the median buy properties which are, on average, 6% more expensive than those of households whose income is below the median. This effect does not change significantly after the introduction of the macroprudential policies. Column (4) jointly controls for above-median wealth and income, and their interactions with the post-2012 indicator, and shows that the results from the previous two columns remain unchanged when doing so.

In sum, we find that receiving a predeath bequest as well as having more wealth and income is relevant for the purchase price of a property. Moreover, wealth has become more important after the introduction of the macroprudential policies. We interpret this finding as additional evidence of a tightening in borrowing constraints. Households receiving a predeath bequest after 2012 are likely to be able to afford a similar quality house despite the stricter down-payment requirements, whereas those that do not might be forced to downsize.

Table 7.3: Effect of Transfer, Wealth, and Income on the Log Purchase Price

	Transfer	Position in the wealth/income distribution		
	(1)	(2)	(3)	(4)
	ln PP	ln PP	ln PP	ln PP
Predeath Bequest=1	0.09***			
	(0.02)			
$\text{MP}_{t,12}$=1 × Predeath Bequest=1	0.03			
	(0.03)			
Above median Wealth		0.14***		0.13***
		(0.02)		(0.02)
$\text{MP}_{t,12}$=1 × Above median Wealth		0.05**		0.05**
		(0.02)		(0.02)
Above median Income			0.06**	0.06**
			(0.02)	(0.02)
$\text{MP}_{t,12}$=1 × Above median Income			0.01	0.00
			(0.03)	(0.03)
Year FE × MS Region FE	Yes	Yes	Yes	Yes
Main Controls	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Observations	10,365	10,365	10,365	10,365
R^2	0.20	0.20	0.19	0.20
$\bar{y}$	13.02	13.02	13.02	13.02

Note: Column (1) shows the effect of receiving a transfer on the log purchase price of the new property. Columns (2) and (3) separately estimate the effects of having above-median wealth and income on the log purchase price of the new property. Column (4) estimates these effects jointly. The medians were calculated for the households reporting the purchase price of the new property. Main controls include lag income, having children, financial wealth, and income ventiles at first observation. $MP_{t,12}$ is a dummy indicating when the macroprudential policies are active during our sample period (2012 to 2016). $\bar{y}$ and $\sigma(y)$ indicate the mean and standard deviation of the dependent variable. * p<0.10, ** p<0.05, *** p<0.010. Standard errors are clustered at the year × MS-Region level. In Table I2 in the Appendix, we provide results for both types of wealth transfers and different specifications of the model.

8. Additional Evidence

In this section, we present additional evidence that borrowing constraints got tighter after the introduction of the macroprudential policies. In particular, we show that the stronger effect of predeath bequests on the probability of acquiring a home is only present for first-time buyers. For supposedly less credit-constrained households who already own a home and could acquire an additional property, predeath bequests have the same effect before and after the introduction of the macroprudential policies. Moreover, we conduct several robustness checks and discuss the external validity of our results.

8.1 Existing Homeowners Buying Additional Properties

Withdrawals of second-pillar pension savings to finance the down-payment for buying a property are only allowed for the principal residence. For this reason, the macroprudential policy requiring households to finance at least 10% of the housing value without second-pillar savings only constrains households who have been renting and transition into homeownership but not those who already own a home and buy an additional property. Moreover, as homeownership is costly, we can expect households with multiple properties to be less credit-constrained in general. Accordingly, predeath bequests may be less important for these households.

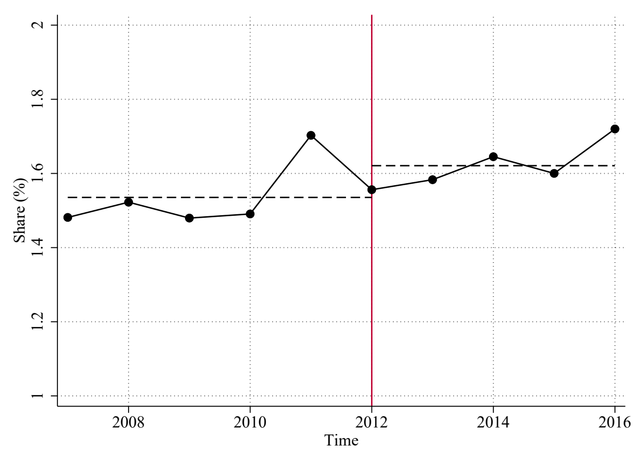

Figure 8.1: Share of Households Buying an Additional Property

Note: This figure shows the annual share of households who already own at least one home and acquire an additional property. The vertical line indicates when macroprudential policies were introduced in Switzerland.

Figure 8.1 shows the share of households who already own a home in the previous year $t-1$ and acquire an additional property in the current year t. The share varies little over the sample period. In contrast to the share of households transitioning into homeownership for the first time, there is no abrupt change around 2012 when the macroprudential policies were introduced.

Figure 8.2 exhibits the share of households receiving a predeath bequest, conditional on their age category and tenure status. The tenure status refers to two different categories of households. The first tenure status, "ΔMRE", refers to households who already own at least one property in $t-1$ and acquire an additional one in t. The second tenure status, "Staying Owner", refers to households who own at least one property in $t-1$ and keep their real estate holdings constant in t.

There are two noteworthy observations. First, households who acquire an additional property receive a predeath bequest more often than households who keep their real estate holdings constant. Second, the share of households receiving a predeath bequest varies only slightly around the introduction of the macroprudential policies. For the youngest age category, the share is highly volatile due to the few observations of young households who acquire an additional property. For the older two age categories, the share is constant.

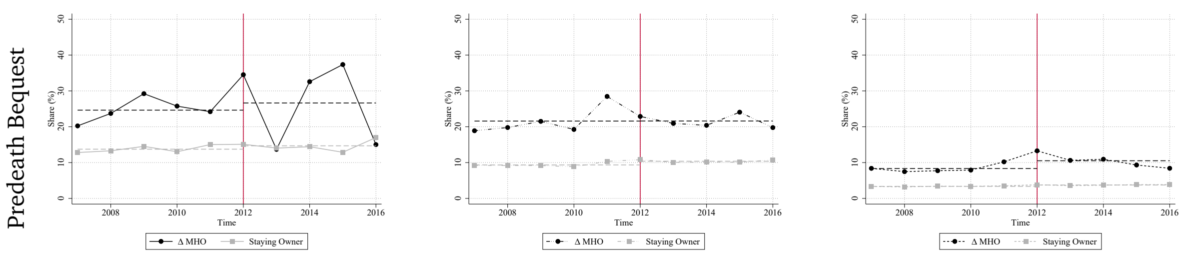

(A) Age ≤ 35 (B) $35 <$ Age ≤ 50 (C) $50 <$ Age ≤ 100

Figure 8.2: Share of Households Receiving a Predeath Bequest by Age Group and Tenure Status

Note: The figure shows the share of households who receive a predeath bequest, conditional on their tenure status and age group. The tenure status refers to two different categories of households. The first tenure status, "ΔMRE", refers to households who own already at least one home in the previous year $t-1$ and acquire additional real estate in year t. The second tenure status, "Staying Owner", refers to households who own at least one home in $t-1$ and keep their real estate holdings constant in t.

In Table 8.1, we use a similar regression model as in Section 7.1 for estimating the effects at the extensive margin. However, the dependent variable ΔMRE indicates households already owning a property in $t-1$ and acquiring an additional one in t. The fixed effects and main control variables are the same as in the baseline specification, except for the income and wealth ventiles at first observation, which are calculated conditional on all households already owning at least one property.

In Column (1), we observe that predeath bequests have a significant effect for households who acquire an additional property. However, the effect is much smaller than for first-time home buyers. Moreover, the effect of receiving a predeath bequest on the probability of acquiring an additional property does not increase after the introduction of the macroprudential policies.

In Column (2), we observe similar results for the different transfer sizes. Larger predeath bequests have a stronger effect on the probability of acquiring an additional property. Yet, across almost all transfer sizes, the effect remains unchanged after the introduction of the macroprudential policies. An exception are transfers between CHF 25,000 and 50,000 which have a stronger effect after 2012. However, the effect is only marginally significant and smaller compared to transfers of the same size for households who were renting before.

Table 8.1: Effect of Receiving a Predeath Bequest for Households Buying an Additional Property

	(1) ΔMRE	(2) ΔMRE
Transfer=1	1.98***	
	(0.11)	
$\text{MP}_{t,12}$=1 × Transfer=1	-0.02	
	(0.15)	
kCHF 1 to 10		0.21
		(0.25)
kCHF 10 to 25		0.54***
		(0.17)
kCHF 25 to 50		0.34*
		(0.19)
kCHF 50 to 100		1.71***
		(0.23)
kCHF 100 or more		4.84***
		(0.27)
$\text{MP}_{t,12}$=1 × kCHF 1 to 10		-0.28
		(0.32)
$\text{MP}_{t,12}$=1 × kCHF 10 to 25		-0.11
		(0.23)
$\text{MP}_{t,12}$=1 × kCHF 25 to 50		0.48*
		(0.29)
$\text{MP}_{t,12}$=1 × kCHF 50 to 100		-0.24
		(0.31)
$\text{MP}_{t,12}$=1 × kCHF 100 or more		0.38
		(0.38)
Year FE × MS Region FE	Yes	Yes
Main Controls	Yes	Yes
Age FE	Yes	Yes
Observations	1,049,114	1,049,114
$\bar{y}$	1.57	1.57

Note: ΔMRE is a dummy indicating households who already own at least one home in the previous year $t-1$ and buy an additional property in t. Transfer is a dummy for a predeath bequest that takes the value of one if the transfer is $\geq$ CHF $1,000$ and zero otherwise. Main controls include lag income, having children, financial wealth and income ventiles at first observation. $MP_{t,12}$ is a dummy indicating when the macroprudential policies are active during our sample period (2012 to 2016). $\bar{y}$ is the mean of the dependent variable. * p<0.10, ** p<0.05, *** p<0.010. Standard errors are clustered at the household level.

8.2 Robustness Checks

8.2.1 Timing of the Introduction of the Macroprudential Policies

As described in Section 2, the introduction of the macroprudential policies occured gradually. While the first policy change was announced in June 2012, it became effective in July 2012 and allowed for a transition period of 5 months. For this reason, we verify whether our results still hold when we use two alternative definitions for the introduction of the macroprudential policies. First, we treat the policies as active from 2013 to 2016. Second, we drop all observations from 2012.

We find that our results are robust to both definitions. Table 8.2 shows the effect of predeath bequests on the probability of transitioning into homeownership for the two alternative definitions. In Column (1), we set the dummy for active macroprudential policies equal to one from 2013 to 2016. In Column (2), we estimate the model without the observations from 2012. The effect of receiving a predeath bequest on the probability of transitioning into homeownership is very close to our baseline estimate in Table 7.1. The coefficient on the interaction of the predeath bequest with the macroprudential policy dummy in Columns (1) and (2) is larger. Thus, if anything, our main result is stronger under these alternative definitions.

Table 8.2: Robustness Checks for Different Definitions of the Macroprudential Policy Dummy

	(1)	(2)
△ HO	Dummy 2013	without 2012
Transfer=1	11.85***	11.99***
	(0.25)	(0.27)
MP_t=1	1.37***	1.22***
× Transfer=1	(0.39)	(0.41)
Year FE × MS Region FE	Yes	Yes
Main Controls	Yes	Yes
Age FE	Yes	Yes
Observations	780,955	703,906
$\bar{y}$	3.28	3.30

Note: The table shows the effect of receiving a predeath bequest on the probability of transitioning into homeownership using different definitions of the dummy for the introduction of macroprudential policies. In Column (1), we set the macroprudential policy dummy equal to 1 for 2013 to 2016. In Column (2), we drop observations from 2012 completely. Transfer is a dummy for a predeath bequest that takes the value of one if the transfer is $\geq$ CHF $1,000$ and zero otherwise. Main controls include lag income, having children, financial wealth, and income ventiles at first observation. △HO indicates whether a household transitioned into homeownership. Regressions are calculated for households renting in the previous year $t-1$. MP_t is a dummy indicating when the macroprudential policies are active; in these specifications, this dummy is = 1 for 2013 to 2016. $\bar{y}$ is the mean of the dependent variable. * p<0.10, ** p<0.05, *** p<0.010. Standard errors are clustered at the household level.

8.2.2 Change in Preferences among Age Groups

Preferences over tenure choices among age groups might change over time. For example, after 2012, homeownership might have become more desirable for households where the mean age of the main taxpayer and the spouse is between 30 and 35. Separate age and year fixed effects do not absorb such a potential change in preferences. For this reason, we include an additional control variable and interact the year of an observation with a variable that groups households according to their mean age in 5 year bins.

Table 8.3 provides evidence that our results are robust to the inclusion of this additional control variable. Column (1) shows the effect of receiving a predeath bequest on the probability of transitioning into homeownership without the additional control variable. Column (2) adds the additional control variable. Columns (3) and (4) show the analogous regressions replacing the dummy for receiving a predeath bequest with the categorical variables for the size of the predeath bequest. The estimates remain robust across all specifications. Thus, we conclude that the additional effect of predeath bequests after 2012 cannot be explained by a change in preferences over tenure choices among different age groups.

Table 8.3: Additional Fixed Effects Absorbing Shifts in Preferences for Housing Tenure Choice

	(1) ΔHO	(2) ΔHO	(3) ΔHO	(4) ΔHO
Transfer=1	12.00***	11.93***		
	(0.27)	(0.27)		
MP_t=1 × Transfer=1	0.77**	0.92**		
	(0.38)	(0.38)		
kCHF 1 to 10			0.30	0.25
			(0.35)	(0.35)
kCHF 10 to 25			2.77***	2.71***
			(0.35)	(0.35)
kCHF 25 to 50			8.58***	8.50***
			(0.59)	(0.59)
kCHF 50 to 100			16.82***	16.72***
			(0.69)	(0.69)
kCHF 100 or more			30.35***	30.25***
			(0.77)	(0.77)
MP_t=1 × kCHF 1 to 10			-0.38	-0.27
			(0.47)	(0.47)
MP_t=1 × kCHF 10 to 25			0.62	0.74
			(0.50)	(0.50)
MP_t=1 × kCHF 25 to 50			2.10**	2.25***
			(0.87)	(0.87)
MP_t=1 × kCHF 50 to 100			2.15**	2.32**
			(1.00)	(1.00)
MP_t=1 × kCHF 100 or more			0.35	0.54
			(1.06)	(1.06)
Year FE × MS Region FE	Yes	Yes	Yes	Yes
Main Controls	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Year FE × 5 Year Age Groups	No	Yes	No	Yes
Observations	780,955	780,947	780,955	780,947
$\bar{y}$	3.28	3.28	3.28	3.28

Note: The table shows the effect of receiving a predeath bequest on the probability of transition into homeownership. It compares two different specifications. In Column (1), we show our baseline model. In Column (2), we add an additional fixed effect from the interaction of the year and five year age group. In Column (3), we have the baseline model specification using a categorical variable for the transfer size. In Column (4), we add to this specification the fixed effect of the interaction of the year and five year age group. The dummy of a transfer is equal to one if the household receives a predeath bequest $\geq$ CHF $1,000$. The additional fixed effect absorbs potential shifts in preferences for housing tenure choice across age groups. Main controls include lag income, having children, financial wealth, and income ventiles at first observation. $MP_{t,12}$ is a dummy indicating when the macroprudential policies are active during our sample period (2012 to 2016). ΔHO indicates whether a household transitioned into homeownership. Regressions are calculated for households renting in the previous year $t-1$. $\bar{y}$ is the mean of the dependent variable. * p<0.10, ** p<0.05, *** p<0.010. Standard errors are clustered at the household level.

8.2.3 External Validity

Besides the administrative data from Bern, we use two additional data sets to assess the external validity of our results.

First, we exploit the nationwide SHP data. The SHP project is based at the Swiss Centre of Expertise in the Social Sciences (FORS) and financed by the Swiss National Science Foundation. The advantage of this data set is that it is representative for the entire country, and that we can include households without restrictions on their civil status. However, compared to the administrative data set from the canton of Bern, the SHP data is only a sample and does not contain the universe of taxpayers. Moreover, it provides just one transfer variable and does not allow to distinguish between inheritances and predeath bequests.

Overall, our main results also hold in the SHP data. Wealth transfers have a significant and positive effect on the probability of transitioning into homeownership. After 2012, we observe a stronger effect of wealth transfers, even though the effect is estimated imprecisely due to the limited number of observations and therefore not statistically significant. Reassuringly and similar to the results for the canton of Bern, wealth transfers with a size from CHF 50'000-100'000 have the strongest effect on the probability of transitioning into homeownership after 2012. Details can be found in Appendix J.

Second, we use a similar administrative data set of individual tax records from the canton of Lucerne. Lucerne provides a particularly useful check on external validity, as it experienced stronger real estate price growth over the sample period than Bern. Transaction prices of single-family houses in Lucerne increased by about 37% from 2007 to 2016, versus about 27% in Bern (SNB, 2020). As tax reports share a common structure between cantons, observable characteristics are similar in both administrative data sets (see Table E1 in the Appendix for summary statistics for Lucerne). However, the data set from Lucerne has two limitations. First, it does not allow us to discriminate between predeath bequests and inheritances. Second, it does not allow us to identify the individual properties owned by a taxpayer. Due to these limitations, we use this data set primarily to compare the descriptive results between Bern and Lucerne. Conditional on households renting in the previous year, the administrative data from Lucerne comprises 334,014 observations.[22]

The descriptive patterns in the data from Lucerne are similar to those from Bern. In Figure E3 in the appendix, we find a similar but more pronounced drop in the share of new homeowners after 2012 across all age categories. Additionally, the patterns in Figure E4 in the appendix depicting the share of homeowners transitioning into homeownership conditional on wealth and income quintiles in Lucerne are comparable to those in the analogous Figure 4.3 for Bern. First, households with higher income and wealth are more likely to transition into homeownership. Second, the drop in the share of households transitioning into homeownership after 2012 is more pronounced when income and wealth are low.

[22]For the canton of Lucerne, tax reports show no information about the age of the spouse. Hence, we use the age of the main taxpayer as a measure of the household's age. However, results from Bern are similar when using the mean age or only the main taxpayer's age.

9. Conclusion

Using administrative tax data from Switzerland, we study how the introduction of macroprudential policies affects the propensity of households to become homeowners and the borrowing constraints they face. We identify borrowing constraints by analyzing the effect of receiving a predeath bequest on renter households' probability of transitioning into homeownership. We find that the yearly share of renter households transitioning into homeownership decreased from an average of 3.4% in the four years prior to the introduction of the macroprudential policies to 3.0% in the four years afterward. However, this decrease is not present for households that could draw from family wealth via predeath bequests. As borrowing constraints tightened, predeath bequests became more important for financing the transition into homeownership, especially for young and middle-aged households. We also find similar evidence at the intensive margin. Predeath bequests and wealth have stronger effects on the purchase price of homes bought after the introduction of macroprudential policies.

The results are robust to different model specifications. In particular, while predeath bequests became more important for first-time home buyers, this is not the case for households who already own at least one home and acquire additional property. These households are presumably less borrowing-constrained and are not affected by the policy limiting withdrawals of pension savings. In addition, we use a representative nationwide sample and similar tax data from a second canton to check for external validity. Qualitatively, our results hold in these samples too.

The stronger effect of predeath bequests on access to homeownership after the introduction of macroprudential policies is probably not unique to Switzerland. Research conducted in both the United States (Bond and Eriksen, 2021; Brandsaas, 2021) and the Euro area Spilerman and Wolff (2012); Mathä et al. (2017) provides similar evidence when it comes to recognizing the significance of intergenerational wealth transfers in facilitating homeownership. Moreover, these jurisdictions have also implemented comparable macroprudential measures (e.g. higher LTV ratios or qualified mortgages), which have increased down-payment requirements and consequently imposed higher borrowing restrictions on households.

Our results have implications for the discussion surrounding macroprudential policies. In Switzerland, these policies have aimed at countering potentially damaging developments in the mortgage and real estate markets, and at strengthening the resilience of the banking system. If effective, macroprudential policies reduce the likelihood and depth of a housing market downturn. Among other things, this happens by tightening borrowing constraints with the aim of preventing households from taking on excessive debt. The reduced ownership propensities of low-wealth and low-income households suggest that the macroprudential policies in Switzerland achieve this aim. At the same time, such policies could have distributional consequences to the extent that homeownership has potential long-run wealth benefits. Some households are able to overcome the tighter constraints via predeath bequests, but this is certainly not an option for all households, as family wealth is highly heterogeneous. To the extent that homeownership has potential long-run wealth benefits, such policies therefore likely also have distributional consequences. A full evaluation of this aspect should, however, also take into account the effects of the policies on home prices, which we have not attempted to study in this work.

Bibliography

Aastveit, K. A., Juelsrud, R., and Getz Wold, E. (2022). The leverage-liquidity trade-off of mortgage regulation. Norges Bank Working Paper 6/2022.

Acharya, V. V., Bergant, K., Crosignani, M., Eisert, T., and McCann, F. J. (2021). The Anatomy of the Transmission of Macroprudential Policies. *Journal of Finance*, forthcoming.

Auer, R., Matyunina, A., and Ongena, S. (2022). The Countercyclical Capital Buffer and the Composition of Bank Lending. *Journal of Financial Intermediation*, forthcoming.

Basten, C. (2020). Higher Bank Capital Requirements and Mortgage Pricing: Evidence from the Counter-Cyclical Capital Buffer. *Review of Finance*, 24(2):453–495.

Behncke, S. (2022). Effects of macroprudential policies on bank lending and credit risks. *Journal of Financial Services Research*.

Benetton, M., Bracke, P., Cocco, J. F., and Garbarino, N. (2019). Housing consumption and investment: evidence from shared equity mortgages. Bank of England Working Paper 790.

Blickle, K. and Brown, M. (2019). Borrowing Constraints, Home Ownership and Housing Choice: Evidence from Intra-Family Wealth Transfers. *Journal of Money, Credit and Banking*, 51(2-3):539–580.

Bond, S. A. and Eriksen, M. D. (2021). The role of parents on the home ownership experience of their children: Evidence from the health and retirement study. *Real Estate Economics*, 49(2):433–458.

Brandsaas, E. E. (2021). Illiquid Homeownership and the Bank of Mom and Dad. Working paper, University of Wisconsin-Madison.

Brülhart, M., Gruber, J., Krapf, M., and Schmidheiny, K. (2021). Behavioral responses to wealth taxes: Evidence from switzerland. *American Economic Journal: Economic Policy*, forthcoming.

Bütler, M. and Stadelmann, S. (2020). Building on a pension: Second pillar wealth as a way to finance real estate? *The Journal of the Economics of Ageing*, 17:100261.

Coibion, O., Gorodnichenko, Y., Kueng, L., and Silvia, J. (2017). Innocent bystanders? monetary policy and inequality. *Journal of Monetary Economics*, 88:70–89.

Danthine, J.-P. (2013). A macroprudential progress report. Remarks by Vice Chairman of the Governing Board of the Swiss National Bank Jean-Pierre Danthine at the Society for Financial Econometrics (SoFiE) Conference, Lugano [Accessed: 2021 07 06].

Di, Z. X., Belsky, E., and Liu, X. (2007). Do homeowners achieve more household wealth in the long run? *Journal of Housing Economics*, 16(3-4):274–290.

Engelhardt, G. V. and Mayer, C. J. (1998). Intergenerational transfers, borrowing constraints, and saving behavior: Evidence from the housing market. *Journal of Urban Economics*, 44(1):135–157.

Eurostat (2021). Distribution of population by tenure status, type of household and income group - EU-SILC survey. Retrieved from https://ec.europa.eu/eurostat/databrowser/view/ILC_LVHO02__custom_1122592/default/table?lang=en.

Federal Statistical Office (2019a). Bilanz der ständigen wohnbevölkerung, nach kanton und stadt, 1999-2019. Retrieved from https://www.bfs.admin.ch/bfs/de/home/statistiken/bevoelkerung.assetdetail.13707271.html.

Federal Statistical Office (2019b). Ms-regionen. Retrieved from https://www.bfs.admin.ch/bfs/de/home/statistiken/raum-umwelt/nomenklaturen/msreg.assetdetail.415729.html.

Federal Statistical Office (2020a). Entwicklung der nominallöhne, der konsumentenpreise und der reallöhne. Retrieved from https://www.bfs.admin.ch/bfs/en/home/statistics/work-income/surveys/sli.assetdetail.13067305.html.

Federal Statistical Office (2020b). Household income and expenditure. Retrieved from https://www.bfs.admin.ch/bfs/en/home/statistics/economic-social-situation-population/income-consumption-wealth/household-budget.html.

Federal Statistical Office (2020c). ThemaKart map boundaries - Set 2020. Retrieved from https://www.bfs.admin.ch/bfs/en/home/statistics/regional-statistics/base-maps/cartographic-bases.assetdetail.11927607.html.

Fuster, A. and Zafar, B. (2016). To Buy or Not to Buy: Consumer Constraints in the Housing Market. *American Economic Review*, 106(5):636–40.

Fuster, A. and Zafar, B. (2021). The Sensitivity of Housing Demand to Financing Conditions: Evidence from a Survey. *American Economic Journal: Economic Policy*, 13(1):231–65.

Galli, A. and Rosenblatt-Wisch, R. (2022). Analysing households' consumption and saving patterns using tax data. Working Paper 3/2022, Swiss National Bank.

Goodman, L. S. and Mayer, C. (2018). Homeownership and the American Dream. *Journal of Economic Perspectives*, 32(1):31–58.

Grand Council of the Canton of Bern (2014). Gesetz über die erbschafts- und schenkungssteuer (eschg), artikel 9.
hhttps://www.belex.sites.be.ch/frontend/versions/195?locale=de.

Guiso, L. and Jappelli, T. (2002). Private transfers, borrowing constraints and the timing of homeownership. *Journal of Money, Credit, and Banking*, 34(2):315–339.

Haurin, D. R., Hendershott, P. H., and Wachter, S. M. (1997). Borrowing constraints and the tenure choice of young households. *Journal of Housing Research*, 8(2):137–154.

International Monetary Fund (2014). Switzerland: Technical Note - Macroprudential Institutional Arrangements and Policies. Technical report, International Monetary Fund, Washington, D.C.

Jann, B. and Fluder, R. (2015). Erbschaften und schenkungen im kanton bern, steuerjahre 2002 bis 2012.

Kanton Bern (2020). Gesetz betreffend die Handänderungssteuer. Retrieved from https://www.belex.sites.be.ch/frontend/versions/1969?locale=de.

Kinghan, C., McCarthy, Y., and O'Toole, C. (2019). How do macroprudential loan-to-value restrictions impact first time home buyers? A quasi-experimental approach. *Journal of Banking & Finance*, page 105678.

Kuhn, U. and Grabka, M. (2018). Homeownership and Wealth in Switzerland and Germany. In *Social Dynamics in Swiss Society*, pages 175–185. Springer, Cham.

Lam, K., Dunsky, R. M., and Kelly, A. (2013). Impacts of down payment underwriting standards on loan performance—evidence from the gses and fha portfolios. *Federal Housing Finance Agency Working Paper*, 12.

Linneman, P. and Wachter, S. (1989). The Impacts of Borrowing Constraints on Homeownership. *Real Estate Economics*, 17(4):389–402.

Lustig, H. N. and Van Nieuwerburgh, S. G. (2005). Housing collateral, consumption insurance, and risk premia: An empirical perspective. *The Journal of Finance*, 60(3):1167–1219.

Mathä, T. Y., Porpiglia, A., and Ziegelmeyer, M. (2017). Household wealth in the euro area: The importance of intergenerational transfers, homeownership and house price dynamics. *Journal of Housing Economics*, 35:1–12.

Peydró, J.-L., Rodriguez Tous, F., Tripathy, J., and Uluc, A. (2020). Macroprudential policy, mortgage cycles and distributional effects: Evidence from the UK. Bank of England Working Paper 866.

Saiki, A. and Frost, J. (2014). Does unconventional monetary policy affect inequality? evidence from japan. *Applied Economics*, 46(36):4445–4454.

Schneider, M. and Wagner, K. (2015). Housing Markets in Austria, Germany and Switzerland. *Monetary Policy & the Economy*, 1:42–58.

Seiler Zimmermann, Y. (2012). Auswirkungen der neuen Eigenmittel-Mindestanforderung auf den Hypothekarmarkt. *Die Volkswirtschaft*.

Seiler Zimmermann, Y. (2013). Nutzung von Vorsorgegeldern zur Finanzierung von selbstgenutztem Wohneigentum. Eine deskriptive Analyse. *IFZ Hochschule Luzern*.

Sodini, P., Van Nieuwerburgh, S., Vestman, R., and von Lilienfeld-Toal, U. (2016). Identifying the Benefits from Homeownership: A Swedish Experiment. Swedish House of Finance Working Paper 16-11.

Spilerman, S. and Wolff, F.-C. (2012). Parental wealth and resource transfers: How they matter in france for home ownership and living standards. *Social Science Research*, 41(2):207–223.

Steuerverwaltung Kanton Bern (2020). Erläuterungen zum steuerlichen Bewertungssystem von Grundstücken und Liegenschaften. *Wegleitungen*.

Swiss Financial Market Supervisory Authority FINMA (2012). Mortgage financing: Finma recognises new minimum standards. Accessed July 2023, retrieved from https://www.finma.ch/en/news/2012/06/mm-hypo-richtlinien-20120601/.

Swiss Financial Market Supervisory Authority FINMA (2014). Mortgage financing: Finma approves amended self-regulation. Accessed June 2022, retrieved from https://www.finma.ch/en/news/2014/07/mm-selbstregulierung-hypothekarbereich-20140702/.

Swiss National Bank (2012). Swiss National Bank Financial Stability Report. *Financial Stability Report 2012*.

Swiss National Bank (2020). SNB data portal. Retrieved from https://www.snb.ch/en/iabout/stat/id/statpub_dataportal.

Tracey, B. and van Horen, N. (2021). The consumption response to borrowing constraints in the mortgage market. Bank of England Working Paper 919.

Turner, T. M. and Luea, H. (2009). Homeownership, wealth accumulation and income status. *Journal of Housing Economics*, 18(2):104–114.

Tzur-Ilan, N. (2019). The Real Consequences of LTV Limits on Housing Choices. Working paper.

Van Bekkum, S., Gabarro, M., Irani, R. M., and Peydró, J.-L. (2020). Take It to the Limit? The Effects of Household Leverage Caps. Working paper.

Wehrmüller, A. (2014). Tenlaw: Tenancy law and housing policy in multi-level europe – national report for switzerland. Retrieved on July 6, 2021, from http://www.iut.nu/wp-content/uploads/2017/03/National-Report-for-Switzerland.pdf.

White, G. F. and Schollaert, P. T. (1993). Home Ownership and Well-Being. *Housing and Society*, 20(1):31–40.

Appendix

A. Age Histogram

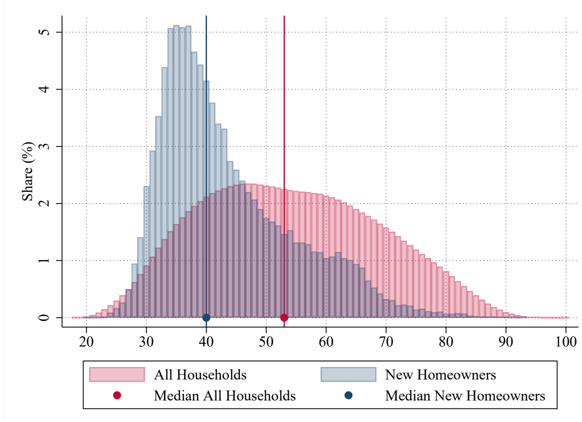

Figure A1: Age Histogram for Households in Bern

Note: The figure shows age distribution separately for all households in the sample and households that rented in year $t-1$ and then transitioned into homeownership in year t for the canton of Bern. The vertical lines indicate the median age for each group of households.

B. Intensive Margin T-tests

Table B1: Intensive Margin - Tests across Groups

	(1) Mean Price reported	(2) Mean No Price reported	(3) Difference	(4) Std.	(5) N Price reported	(6) N No Price reported
Age	44.18	41	3.18***	.14	15,339	10,365
Received an Inheritance (0/100	9.69	5.91	3.77***	.35	15,339	10,365
Received a Predeath Bequest (0/100)	23.9	26.56	–2.66***	.55	15,339	10,365
Lag Income (in kCHF)	115.53	118.36	–2.83**	1.06	15,339	10,365
Lag Wealth (in kCHF)	190.6	160.68	29.92***	6.3	15,339	10,365
Has Children (0/100)	66.56	74.02	–7.46***	.58	15,339	10,365
Year	2011.19	2011.32	–0.12***	.04	15,339	10,365
Price-to-Rent ratio (100 = 2007	119.55	120.52	–0.97***	.16	15,339	10,365
Max. Number of Obs. in sample	11.23	10.79	0.44***	.04	15,339	10,365

Note: The table shows the characteristics of households that are not reporting the purchase price and those households reporting the purchase price in the tax report. Reporting the purchase price in the tax report is voluntary. Difference shows the difference in mean for the characteristics of the households. The stars indicate the significance of the difference in mean among the households reporting the purchase price and those that do not using a t-test. Std. indicates the standard deviation and N indicate the number of observations for those who report the purchase price and those who do not, respectively. Note that the variable Max. Number of obs. in sample represents the maximum number of observations a household was observed in the Sample.

C. Map of Switzerland

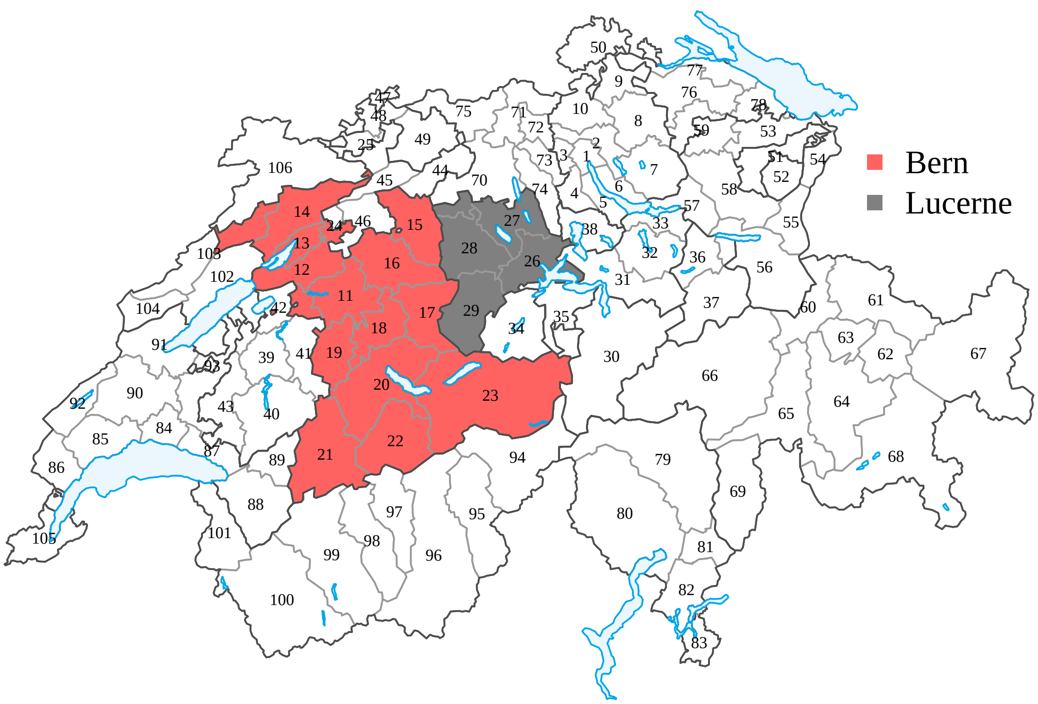

Figure C1: MS-regions of Switzerland

Note: The figure shows all 106 MS-regions of Switzerland (FSO, 2020c). The MS-regions are small labour market areas with a functional orientation towards centres and are characterized by a certain spatial homogeneity. Highlighted are the canton of Bern and the canton of Lucerne. Dark grey lines show cantonal borders whereas light grey lines indicate different MS-regions.

D. Macroprudential Policies in Switzerland - Timeline

Table D1: Timeline Introduction Macroprudential Policies

(1) Announced	(2) Effective	(3) Policy
June 2012	July 2012, with a 5 month transition period	The Swiss Bankers Association tightened the down-payment requirements in its catalogue of self-regulations. Under the new requirements, home buyers need to finance at least 10% of the purchase price with "hard" equity capital, without drawing from second-pillar pension savings.
June 2012	January 2013	The Swiss Federal Council (the executive branch of the Swiss government) raised banks' capital requirements for originated mortgage loans with high LTV ratios: by January 2013, the risk-weights for the loan tranche exceeding an LTV ratio of 80% increased from 75% to 100%.
February 2013	September 2013	The Swiss Federal Council activated the sectoral countercyclical capital buffer (CCyB), requiring banks to hold additional common equity Tier 1 (CET1) capital on domestic residential mortgage loans. The CCyB was set to 1% of a bank's relevant risk-weighted assets.
January 2014	June 2014	The Swiss Federal Council increased the sectoral CCyB to 2% of a bank's relevant risk-weighted assets.
June 2014	September 2014, with a 5 month transition period	The Swiss Bankers Association tightened the amortisation structure in its catalogue of self-regulations. New mortgages must be amortised to a LTV of two-thirds within 15 years, subject to linear repayment.

Note: The table shows the timeline of the introduction of Macroprudential Policies in Switzerland from 2012 to 2016. Sources: Behncke (2022); Swiss Financial Market Supervisory Authority FINMA (2014)

E. Lucerne Tax Data

E.1 Summary Statistics from Lucerne

Table E1: Summary Statistics

	(1) Mean	(2) Std. Dev.	(3) N
ΔHO (0/100)	3.1	17.4	334,014
Age	53.2	16.6	334,014
Share of people with age ≤ 35 (0/100)	16.4	37.0	334,014
$35 < \text{Age} \leq 50$	34.0	47.4	334,014
$50 < \text{Age} \leq 100$	53.9	49.9	334,014
Lag Income (in kCHF)	84.5	49.4	334,014
Lag Wealth (in kCHF)	129.3	506.8	334,014
Has Children (0/100)	45.7	49.8	334,014
Price-to-Rent ratio (100 = 2007)	126.1	14.3	334,014

Note: The table shows the summary statistics of all variables for households renting in the previous year $t-1$. Variables with (0/100) in parentheses are dummy variables scaled from 0 to 100 to indicate percentages. ΔHO refers to the share of households who rented in year $t-1$ and transitioned into homeownership in year t. Age refers to the mean age of the main taxpayer. The base year of the price-to-rent ratio index is 2007.

E.2 Descriptive Evidence from Lucerne

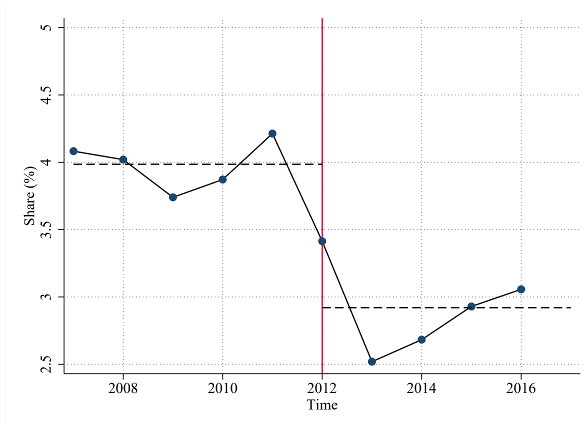

Figure E1: Share of Households transitioning into Homeownership in Lucerne

Note: The figure shows the share of households transitioning into homeownership for the canton of Lucerne. It is calculated conditional on all households renting at time $t-1$. Horizontal lines indicate means before and after 2012.

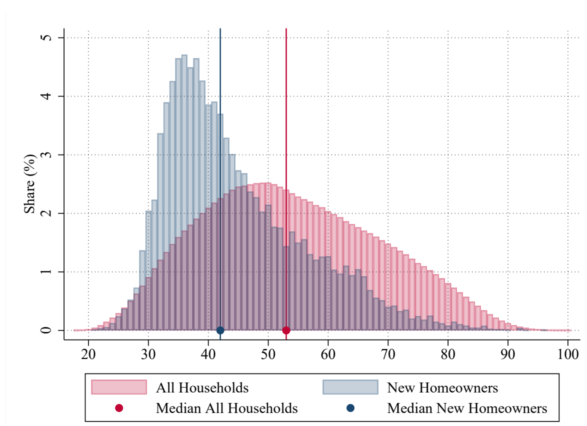

Figure E2: Age Histogram for Households in Lucerne

Note: The figure shows age distribution for households who were renting in year $t-1$ and transitioned into homeownership in year t as well as the overall age distribution of all households in the sample.

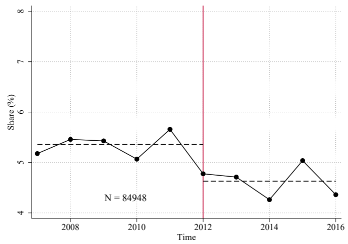

(A) Age category 1

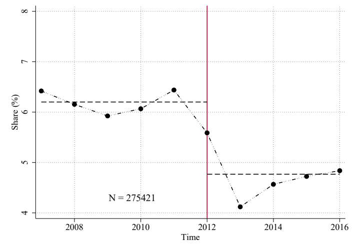

(B) Age category 2

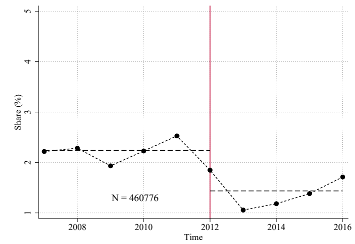

(C) Age category 3

Figure E3: Share of New Homeowners conditional on Age in Lucerne

Note: The figure shows the mean share of new homeowners for the canton of Lucerne conditional on the age category. "N" is the number of all renters in year $t-1$ in each age category. Horizontal lines indicate means before and after 2012.

E.3 Changing Effects of Wealth and Income

The pattern in Lucerne is qualitatively similar to the one in Bern: after 2012, the propensity to enter homeownership decreased relatively more for renters with low income and/or low wealth.

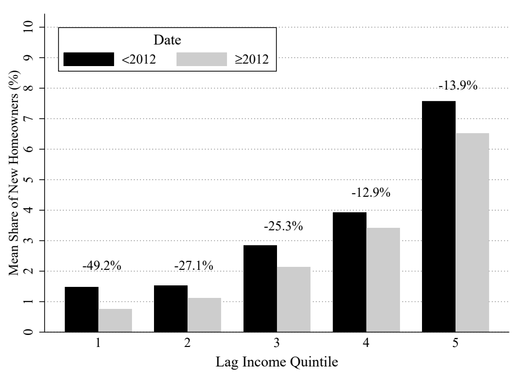

(A) Conditional on Lag Income

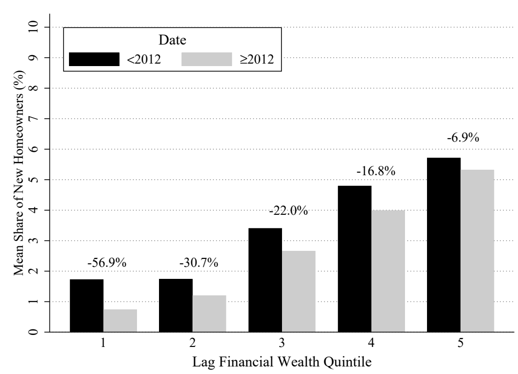

(B) Conditional on Lag Financial Wealth

Figure E4: Share to transition into Homeownership conditional on Wealth and Income Quintile in Lucerne

Note: The figure shows the share of households transitioning into homeownership conditional on their lag income and wealth quintile for periods before 2012 and after, respectively for Lucerne. The quintiles are calculated conditional on being a renter in year $t-1$. The numbers above the bars indicate the percentage change in the share before and after 2012.

F. Distribution of Transfers over Time

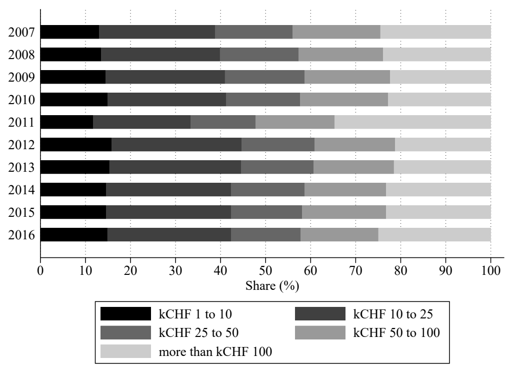

(A) Predeath Bequest

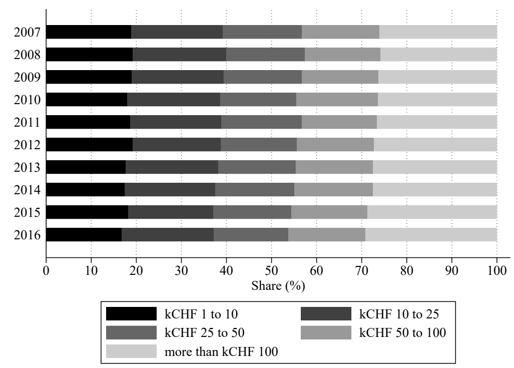

(B) Inheritances

Figure F1: Share of transfers conditional on having received a positive transfer

Note: The figure shows the distribution of predeath bequests and inheritances over time, conditional on households who received a positive transfer of CHF > 1000. Note that the higher share of predeath bequests of more than 100,000 CHF in 2011 is mainly due to the anticipation of a possible retroactive introduction of inheritance tax from January 1, 2012. We record a larger number of predeath bequests above 500,000 CHF in 2011. This is most likely due to a precautionary measure taken by households in order to avoid the impending inheritance tax by passing on large portions of their assets to their future heirs by way of a gift in 2011 (Jann and Fluder, 2015).

G. Regions with Different Price Dynamics

Table G1: Effect of a Predeath Bequest for Regions with Different Price Dynamics

	Growth: Below Median	Growth: Above Median	Level: Below Median	Level: Above Median
	(1)	(2)	(3)	(4)
	ΔHO	ΔHO	ΔHO	ΔHO
Transfer=1	11.98***	12.03***	11.15***	12.87***
	(0.40)	(0.38)	(0.38)	(0.39)
MP_t=1	1.06*	0.55	2.22***	-0.65
$\times$ Transfer=1	(0.55)	(0.53)	(0.54)	(0.54)
Year FE $\times$ MS Region FE	Yes	Yes	Yes	Yes
Main Controls	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Observations	369,443	411,512	372,105	408,850
$\bar{y}$	3.52	3.06	3.72	2.89

Note: The table shows the effect of a predeath bequest on the probability of transition into homeownership. The columns refer to different samples. In Columns (1) and (2), we split the sample across households that live in regions with below or above median price-to-rent ratio growth, respectively. In Columns (3) and (4), we split the sample across households that live in regions with below or above median price-to-rent ratio, respectively. "Transfer" is a dummy equal to one if a household received a transfer of at least CHF 1,000 in year t or $t-1$, and zero otherwise. Main controls include lag income, having children, financial wealth and income ventiles at first observation. $MP_{t,12}$ is a dummy indicating when the macroprudential policies are active during our sample period (2012 to 2016). ΔHO indicates whether a household transitioned into homeownership. Regressions calculated conditional on households renting year $t-1$. $\bar{y}$ is mean of the dependent variable. * p<0.10, ** p<0.05, *** p<0.010. Standard errors are clustered at the household level.

H. Intensive Margin and Different Transfer Sizes

Table H1: Effect of different Transfer Sizes on the Log Purchase Price

	Predeath Bequest	
	(1)	(2)
	ln PP	ln PP
Transfer=1	0.09***	
	(0.02)	
$\text{MP}_{t,12}$=1 × Transfer=1	0.03	
	(0.03)	
kCHF 1 to 10		-0.03
		(0.12)
kCHF 10 to 25		-0.03
		(0.05)
kCHF 25 to 50		0.03
		(0.04)
kCHF 50 to 100		0.04
		(0.03)
kCHF 100 or more		0.18***
		(0.03)
$\text{MP}_{t,12}$=1 × kCHF 1 to 10		0.01
		(0.16)
$\text{MP}_{t,12}$=1 × kCHF 10 to 25		0.02
		(0.07)
$\text{MP}_{t,12}$=1 × kCHF 25 to 50		0.03
		(0.06)
$\text{MP}_{t,12}$=1 × kCHF 50 to 100		0.11**
		(0.05)
$\text{MP}_{t,12}$=1 × kCHF 100 or more		-0.01
		(0.04)
Year FE × MS Region FE	Yes	Yes
Main Controls	Yes	Yes
Age FE	Yes	Yes
Observations	10,365	10,365
R^2	0.20	0.20
$\bar{y}$	13.02	13.02

Note: Column (1) shows the effect of receiving a transfer on the log purchase price of the new property. Column (2) uses categorical variables for different transfer sizes (omitted cateogry: households who receive no transfer). Main controls include lag income, having children, financial wealth, and income ventiles at first observation. $MP_{t,12}$ is a dummy indicating when the macroprudential policies are active during our sample period (2012 to 2016). $\bar{y}$ and $\sigma(y)$ indicate the mean and standard deviation of the dependent variable. * p<0.10, ** p<0.05, *** p<0.010. Standard errors are clustered at the year × MS-Region level.

I. Different Model Specifications for the Extensive and Intensive Margin

Table I1: Different Model Specifications for the Effect of a Predeath Bequest on Homeownership

	(1) ΔHO	(2) ΔHO	(3) ΔHO	(4) ΔHO	(5) ΔHO	(6) ΔHO
Transfer=1	13.59***	12.45***	12.00***			
	(0.27)	(0.28)	(0.27)			
MP_t=1	0.86**	0.81**	0.77**			
× Transfer=1	(0.38)	(0.38)	(0.38)			
kCHF 1 to 10				0.86**	0.59*	0.30
				(0.35)	(0.35)	(0.35)
kCHF 10 to 25				3.73***	3.05***	2.77***
				(0.34)	(0.35)	(0.35)
kCHF 25 to 50				9.97***	8.97***	8.58***
				(0.58)	(0.59)	(0.59)
kCHF 50 to 100				18.65***	17.32***	16.82***
				(0.69)	(0.70)	(0.69)
kCHF 100 or more				32.95***	31.04***	30.35***
				(0.76)	(0.77)	(0.77)
MP_t=1				-0.26	-0.40	-0.38
× kCHF 1 to 10				(0.46)	(0.47)	(0.47)
MP_t=1				0.85*	0.72	0.62
× kCHF 10 to 25				(0.49)	(0.50)	(0.50)
MP_t=1				2.25***	2.11**	2.10**
× kCHF 25 to 50				(0.86)	(0.87)	(0.87)
MP_t=1				2.29**	2.19**	2.15**
× kCHF 50 to 100				(1.00)	(1.01)	(1.00)
MP_t=1				0.11	0.38	0.35
× kCHF 100 or more				(1.05)	(1.06)	(1.06)
Year FE × MS Region FE	Yes	Yes	Yes	Yes	Yes	Yes
Main Controls	No	Yes	Yes	No	Yes	Yes
Age	No	No	Yes	No	No	Yes
Observations	828,517	780,955	780,955	828,517	780,955	780,955
$\bar{y}$	3.22	3.28	3.28	3.22	3.28	3.28

Note:The table shows the effect of a predeath bequest on the probability of a household transitioning into homeownership for different model specifications. In Columns (1) to (3), "Transfer" is a dummy equal to one if a household received a transfer of at least CHF 1,000 in year t or $t-1$, and zero otherwise. Columns (4) to (6) use categorical variables for different transfer sizes (omitted category: households who receive no transfer). Main controls include lag income, having children, financial wealth and income ventiles at first observation. $MP_{t,12}$ is a dummy indicating when the macroprudential policies are active during our sample period (2012 to 2016). ΔHO indicates whether a household transitioned into homeownership. Regressions calculated conditional on households renting in the previous year $t-1$. $\bar{y}$ is the mean of the dependent variable. * p<0.10, ** p<0.05, *** p<0.010. Standard errors are clustered at the household level.

Table I2: Different Model Specifications for the Effect of a Transfer on the Log Purchase Price

	Predeath Bequest			Inheritance		
	(1) ln PP	(2) ln PP	(3) ln PP	(4) ln PP	(5) ln PP	(6) ln PP
Transfer=1	0.11***	0.10***	0.09***	-0.03	-0.02	0.01
	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
MP_t=1 × Transfer=1	0.02	0.03	0.03	0.04	0.06	0.06
	(0.03)	(0.03)	(0.03)	(0.05)	(0.05)	(0.05)
Year FE × MS Region FE	Yes	Yes	Yes	Yes	Yes	Yes
Main Controls	No	Yes	Yes	No	Yes	Yes
Age	No	No	Yes	No	No	Yes
Observations	11,002	10,368	10,365	11,002	10,368	10,365
$\bar{y}$	13.02	13.02	13.02	13.02	13.02	13.02
$\sigma(y)$	0.65	0.64	0.64	0.65	0.64	0.64

Note: The table shows the effect of a transfer (predeath bequest or inheritance) on the log purchase price of the new property for different model specifications. Main controls include lag income, having children, financial wealth and income ventiles at first observation. $MP_{t,12}$ is a dummy indicating when the macroprudential policies are active during our sample period (2012 to 2016). $\bar{y}$ and $\sigma(y)$ are mean and st. dev. of the dependent variable. * p<0.10, ** p<0.05, *** p<0.010. Standard errors are clustered at the household level.

J. Results for Nationwide SHP Data

In Tables J1 and J2, we compare results from the SHP panel data set to the tax data from Bern. The SHP data also encompasses non-married households and observations from all cantons. However, it only contains data for a general wealth transfer and does not allow to distinguish between a predeath bequest and inheritance. This wealth transfer could include predeath bequests and/or inheritances.

To compare the results of the SHP data set to the tax data, we use information about inheritances and predeath bequests in the tax data to generate a similar variable of wealth transfer. In more detail, the dummy for a wealth transfer in the tax data is equal to one if a household has received a predeath bequest or an inheritance in year t or $t-1$.

Then, we estimate similar regressions for the extensive margin using both the SHP and the tax data set from Bern over the same sample periods. We use the same control variables in both data sets with the exception of initial wealth ventiles. The SHP data set does not include information about the wealth of the households. For this reason, we only control for initial income ventiles in the regressions with the SHP data.

Table J1 shows the results of the extensive margin regressions for both data sets. Comparing columns (1) and (4), we find that a wealth transfer increases the probability of transitioning into homeownership by 2.09 and 7.62 percentage points in the SHP and the tax data, respectively. On average, after 2012, the probability of transitioning into homeownership decreases by 0.74 and 0.45 percentage points. Receiving a wealth transfer after 2012 increases the probability of transitioning into homeownership by 0.90 percentage points in the SHP data and by 0.50 in the tax data, although the SHP estimate is not very precise. Adding year $\times$ MS-region fixed effects in columns (2) and (5) changes the estimates only marginally and does not affect the significance level of the coefficients.

The estimates in Columns (3) and (6) show that the effect of a wealth transfer increases with their size in both data sets. After the introduction of the macroprudential policies, the effect is the strongest for transfers in between CHF 50,000 to 100,000 for the SHP data, which is similar to the observation of our main result in Table 7.1 with the tax data and predeath bequests. For wealth transfers in the tax data, the effects are strongest between CHF 25,000 to 50,000, closely followed by transfers between CHF 50,000 to 100,000.

Table J2 shows the results of the interaction of transfers with an age category dummy. Due to the lower number of observations in the SHP data, we separated the samples into two age categories using an age cut-off of 50 years. In both data sets, receiving a transfer is important for all households. It increases the probability to transition into homeownership significantly by 2.22 percentage points in the SHP data and 2.76 percentage points in the tax data, respectively. While the transfer is significantly more important for younger households in the tax data, the effect for the SHP data is insignificant and negative. Potentially, this is due to the lower number of reported large transfers in the SHP data. Nevertheless, in both data sets, the probability of young households to transition into homeownership decreases after 2012, by 1.53 and 0.47 percentage points for the SHP data and tax data, respectively. Similarly, receiving a transfer after 2012 increases the probability of younger households transitioning into homeownership in both data sets. For the SHP data, the probability increases by 1.11 percentage points, versus 1.52 percentage points for the tax data. In general, the results from the nationwide SHP data go in the same direction as in the tax data, even though the effects are less precisely measured due to the smaller sample size and presumably larger measurement error.

Table J1: Comparison of the Effect of a Wealth Transfer using SHP and Tax Data

	Wealth Transfer, SHP Data			Wealth Transfer, Tax Data		
	(1)	(2)	(3)	(4)	(5)	(6)
	ΔHO	ΔHO	ΔHO	ΔHO	ΔHO	ΔHO
Transfer = 1	2.09***	1.98***		7.62***	7.63***	
	(0.54)	(0.53)		(0.17)	(0.17)	
$\text{MP}_{t,12} = 1$	–0.74*			–0.45***		
	(0.38)			(0.04)		
$\text{MP}_{t,12} = 1 \times$ Transfer $= 1$	0.90	0.87		0.50**	0.47**	
	(0.74)	(0.74)		(0.23)	(0.23)	
kCHF 1 to 10			–0.25			0.28
			(0.62)			(0.21)
kCHF 10 to 25			–0.04			1.56***
			(0.78)			(0.23)
kCHF 25 to 50			3.37**			4.71***
			(1.49)			(0.33)
kCHF 50 to 100			3.85**			10.82***
			(1.90)			(0.44)
kCHF 100 or more			13.98***			20.81***
			(2.32)			(0.50)
$\text{MP}_{t,12} = 1\times$ kCHF 1 to 10			0.59			–0.23
			(0.83)			(0.29)
$\text{MP}_{t,12} = 1\times$ kCHF 10 to 25			1.36			0.44
			(1.12)			(0.32)
$\text{MP}_{t,12} = 1\times$ kCHF 25 to 50			–0.07			0.87*
			(1.95)			(0.49)
$\text{MP}_{t,12} = 1\times$ kCHF 50 to 100			4.51			0.67
			(2.96)			(0.62)
$\text{MP}_{t,12} = 1\times$ kCHF 100 or more			–0.74			–0.01
			(3.34)			(0.69)
Year FE × MS Region FE	No	Yes	Yes	No	Yes	Yes
Main Controls	Yes	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	13,786	13,697	13,697	780,955	780,955	780,955
$\bar{y}$	4.26	4.23	4.23	3.28	3.28	3.28

Note: The table shows the effect of a wealth transfer on the probability of transitioning into homeownership for the nationwide SHP data. In Column (1), (2), (4) and (5) is a dummy equal to one if a household received a transfer of at least CHF 1,000 in year t or $t-1$, and zero otherwise. Columns (3) and (6) use categorical variables for different transfer sizes (omitted category: households who receive no transfer). Main controls include lag income, having children, financial wealth (only for tax data) and income ventiles at first observation. $\text{MP}_{t,12}$ is a dummy indicating when macroprudential policies are active during our sample period (2012 to 2016). ΔHO indicates whether a household transitioned into homeownership. Regressions are calculated for households renting in the previous year $t-1$. $\bar{y}$ is the mean of the dependent variable. Compared to columns (2) and (5), Columns (1) and (4) have no year fixed effects but control for MS-Region Fixed Effects. * p<0.10, ** p<0.05, *** p<0.010. Standard errors are clustered at the household level.

Table J2: Comparison of the Effect of a Wealth Transfer Conditional on Age using SHP and Tax Data

	(1) SHP Data	(2) Tax Data
Transfer = 1	2.22***	2.76***
	(0.79)	(0.17)
$\text{MP}_{t,12} = 1\times$ Transfer = 1	0.20	–0.16
	(1.13)	(0.23)
Transfer = 1 $\times$ Age Category = 1	–0.38	9.30***
	(1.06)	(0.32)
$\text{MP}_{t,12} = 1\times$ Age Category = 1	–1.53*	–0.47***
	(0.79)	(0.07)
$\text{MP}_{t,12} = 1\times$ Transfer = 1 $\times$ Age Category = 1	1.11	1.52***
	(1.49)	(0.46)
Year FE $\times$ MS Region FE	Yes	Yes
Main Controls	Yes	Yes
Age FE	Yes	Yes
Observations	13,697	780,955
$\bar{y}$	4.23	3.28

Note: The table shows the effect of receiving a wealth transfer interacted with the age category on the probability of transitioning into homeownership. Age Category 1 refers to all households younger than 50 years. The omitted category are households aged 50 years and older. Column (1) shows the results for the nationwide SHP data and column (2) the results of a similar regression using the tax data from the canton of Bern. Main controls include lag income, having children, financial wealth (only for tax data) and income ventiles at first observation. $\text{MP}_{t,12}$ is a dummy indicating when the macroprudential policies are active during our sample period (2012 to 2016). ΔHO indicates whether a household transitioned into homeownership. Regressions calculated for all households renting in the previous year $t-1$. $\bar{y}$ is the mean of the dependent variable. * p<0.10, ** p<0.05, *** p<0.010. Standard errors are clustered at the household level.

Chapter 2

Do Local Forecasters Have Better Information

Kenza Benhima
University of Lausanne

Elio Bolliger
University of Lausanne

Abstract

Using individual inflation and GDP growth forecasts by professional forecasters for a panel of emerging and advanced economies, we provide direct evidence that foreign forecasters update their forecasts less frequently than local forecasters (about 10% less frequently) and make larger errors in absolute value (up to 9% larger). The local forecasters' more accurate expectations are not due to a more irrational expectation formation by foreigners, but to local forecasters' more precise information. The asymmetry is stronger at shorter horizons and when forecasting inflation. In general, the asymmetry is not weaker when forecasting is less uncertain. Taken together, our results provide a basis for disciplining international finance and trade models with heterogeneous information. On the methodological side, we provide tests that identify differences in information frictions across groups.

Keywords: Information asymmetries, Expectation formation.

JEL: E3, E7, D82.

1. Introduction

The informational advantage of locals over foreigners regarding macroeconomic fundamentals has far-reaching consequences. Information asymmetries are one of the main explanations for the home bias in asset holdings. The home bias in asset holdings, originally documented by French and Poterba (1991), refers to the fact that domestic assets constitute a disproportionate share of portfolios.[1] Information asymmetries are also a potential source of capital flow volatility, since disagreement between foreign and domestic investors can generate cross-border asset trade.[2] Beyond their impact on international asset markets, they also constitute a barrier to the international trade in goods, as highlighted by Anderson and van Wincoop (2004).[3] Finally, recent papers highlight their role in international business cycle comovement.[4] Yet, there is still a lack of direct proof of the existence of information asymmetry on macroeconomic fundamentals and of quantitative estimates of the extent of this asymmetry that could be used to discipline international finance and trade models with heterogeneous information.

We fill this gap using a unique dataset of yearly inflation and GDP growth forecasts by local and foreign forecasters. Unlike previous studies, the forecaster and country dimensions of the panel allows us to control for a rich set of fixed effects. We first show that foreign forecasters update their forecasts about 10% less frequently than local forecasters. They also make more mistakes than local forecasters, as foreign forecasters' excess absolute error can be as high as 9%, depending on the horizon and on the forecasted variable.[5] The local advantage is especially large when predicting inflation as opposed to GDP and it is stronger for shorter forecasting horizons.

We then investigate the role of information frictions and behavioral biases in explaining our results about errors. We do this in two steps. First, we rule out behavioral biases such as overreaction to new information as explanations of the foreigners' excess mistakes, by showing that the local and foreign behavioral biases do not differ systematically. Second, we test for the relative precision of local and foreign forecasters' private information, and find that

[1] See also Ahearne et al. (2004), Portes and Rey (2005) and Coeurdacier and Rey (2013). Work on asymmetric information and the home bias includes Pàstor (2000), Brennan and Cao (1997), Portes et al. (2001), Van Nieuwerburgh and Veldkamp (2009), Mondria (2010), De Marco et al. (2021).

[2] See Yuan (2005), Albuquerque et al. (2007), Albuquerque et al. (2009), Brennan and Cao (1997), Broner et al. (2013), Tille and van Wincoop (2010), Tille and van Wincoop (2014), Benhima and Cordonier (2022).

[3] See also Head and Mayer (2013), Allen (2014), Dasgupta and Mondria (2018), Eaton et al. (2021). Baley et al. (2020) show that cross-border uncertainty may sometimes increase trade.

[4] See Iliopulos et al. (2021) and Bui et al. (2021).

[5] Are these estimates economically significant? Take for instance the home bias in equity holdings. Van Nieuwerburgh and Veldkamp (2009) show that a difference in the variance of priors as small as 10% can generate empirically plausible levels of home bias when investors can choose what information to learn before they invest.

local forecasters have more precise private information. To do so, we build on and extend the fast-growing literature that uses model-based tests to identify frictions in the expectation formation of survey respondents (Coibion and Gorodnichenko, 2015; Bordalo et al., 2020; Kohlhas and Broer, 2022; Angeletos et al., 2021; Goldstein, 2021). In particular, we provide tests of asymmetric information that are robust to the presence of public signals (more on that below). These tests show that foreign forecasters have less precise information.

Finally, we explore some determinants of the information asymmetry between local and foreign forecasters, and examine whether the asymmetry is related to factors that drive forecasting uncertainty. Interestingly, the asymmetry is not reduced when forecasting is less uncertain. If anything, it is increased. Indeed, the local advantage is higher for short horizons and for inflation (as opposed to GDP growth), but also for large countries. In all these situations, the forecasting uncertainty (measured by the average forecast error) happens to be *smaller*. However, we find no evidence that the difference in forecast errors between local and foreign forecasters is linked to the development status of the country, to institutional quality, or to the volatility of business cycles or financial markets, despite the fact that these variables do affect the average forecasting uncertainty.[6] These results should help further discipline the link between uncertainty and information asymmetry in models of international finance and trade.

This evidence suggests that when information becomes available, it flows to local forecasters, and sometimes, but not always, to foreign forecasters. These results would be consistent with a better access to locally-produced information (by knowing when and where relevant information is released). Indeed, we show that the information asymmetry is stronger for nowcasting (when forecasting the current year's GDP growth or inflation), and that it increases in the course of a year (the asymmetry is higher in December than in January). This is consistent with the idea that local forecasters are exposed to the regular releases of partial GDP growth and inflation figures and integrate this information faster. Interestingly, inflation figures are typically available at a higher frequency and with a shorter lag than GDP, making the access to that information an even greater advantage. This is consistent with our finding that the difference in updating frequency is larger for inflation forecasts than for GDP growth forecasts.

As we don't measure the incentives to acquire information at the forecaster level, we cannot document the extent to which the local advantage is determined by those incentives. However, we do find evidence that the forecasts issued by the financial industry are on average more precise than the forecasts issued by the non-financial industry. This is consistent with the

[6]These findings echo the weak link between uncertainty and disagreement that has been documented in the literature (Lahiri and Sheng, 2010; Rich and Tracy, 2010).

idea that the finance industry has more incentives to produce accurate macroeconomic forecasts in order to better allocate portfolios between countries or between equity and bonds. However, there is no significant difference between the local advantage of the financial sector and that of the non-financial sector.

This paper contributes to the recent literature that uses professional forecasters' expectations to identify information frictions and behavioral biases. This literature has used reduced-form estimations as indicators of deviations from Full-Information Rational Expectations (FIRE). Coibion and Gorodnichenko (2015) (CG henceforth) use the estimated coefficient in the regression of the consensus error on the consensus revision as an indicator of deviations from Full Information (FI). Bordalo et al. (2020), (BGMS henceforth) Kohlhas and Broer (2022) (BK henceforth) and Angeletos et al. (2021) (AHS henceforth), use the estimated coefficient in the individual pooled regression as an indicator of deviations from Rational Expectations (RE).[7] We borrow this test directly from the literature to assess whether domestic and foreign behavioral biases differ.

However, CG's Full Information (FI) test, which has been commonly used in the literature, is not adapted to our purpose. Indeed, in the presence of public information, the CG coefficient, which is a common measure of information frictions, is biased. Importantly, the bias depends on the precision of the public signal and is not a monotonic function of the precision of private signals. Comparing the CG coefficient across local and foreign forecasters cannot indicate which group faces more frictions.[8] We thus provide two tests that are robust to the presence of public information. The first one relies on individual regressions à la BGMS but with country-time fixed effects to capture aggregate shocks and the public signals. This test is similar in spirit to Goldstein (2021), who proposes to use forecasters' deviations from the mean to measure information frictions robustly. The second test infers the relative precision of private information from the relative reaction of expectations to public signals.

This paper also belongs to the empirical literature documenting the local informational advantage. Many studies provide indirect evidence of asymmetric information between domestic and foreign investors by showing that location matters for portfolio composition and for portfolio returns.[9] However, based on investor choices and returns, some papers find that foreign

[7] An earlier literature has previously identified deviations from rationality by studying the joint behavior of actual on predicted values, the auto-correlation of forecasts revisions and the predictability of errors. See, for example, Mincer and Zarnowitz (1969), Zarnowitz (1983), Nordhaus (1987), Clements (1997), Lahiri and Sheng (2008).

[8] Both CG and Goldstein (2021) have emphasized that the CG coefficient is biased, but have not highlighted the implied non-monotonicity.

[9] See for instance Kang and Stulz (1997), Grinblatt and Keloharju (2001), Dvořák (2003), Portes and Rey (2005), Ahearne et al. (2004), Hamao and J. (2001), Hau (2001), Choe et al. (2005), Baik et al. (2010) and Sialm et al. (2020).

investors perform better than local investors (e.g. Grinblatt and M. (2000)).[10] In contrast to these studies, we investigate whether location affects the quality of information possessed by forecasters, thus providing direct evidence of information asymmetries. Closest to our study is the paper by Bae et al. (2008), which studies the performance of local and foreign analysts in forecasting earnings for firms. Our focus is different since we examine whether local forecasters outperform foreign ones in forecasting aggregate variables. Besides, we not only document the foreign forecasters' excess errors, but we also investigate whether these excess mistakes come from information frictions or behavioral biases. Finally, other studies document foreigners' lack of attention to domestic information.[11]

The paper is structured as follows. Section 2 describes our dataset. Section 3 focuses on the updating frequency of forecasts. Section 4 documents the foreign forecasters' excess mistakes. Section 5 lays down a model of expectation formation and tests for the sources of the foreigners' excess mistakes. Section 6 investigates drivers of forecast errors and asymmetric information. Section 7 provides several robustness checks. Finally, section 8 concludes.

2. The Data

Forecasts.

We use data from Consensus Economics. Consensus Economics is a survey firm polling individual economic forecasters on a monthly frequency. The survey covers 51 advanced and emerging countries and we focus on observations between 1998 and 2021.[12] Each month, forecasters provide estimates of several macroeconomic indicators for the current and the following year. An advantage of this dataset is that it allows for meaningful comparisons across both countries and forecasters.[13] In this paper, we focus on two indicators, namely inflation and GDP growth. The dataset discloses the name of the individual forecasters. There are 748 unique forecasters from which 149 conduct forecasts for at least 2 distinct countries. For each forecaster-country pair, the average (median) number of observations is 80 (60), which corresponds to approximately 7 (5) years. This leads to an unbalanced panel dataset.

[10] This could be explained by the specialization of some investors in some specific markets where they have an initial informational advantage. This informational advantage can be due to location, but not only. Therefore, information heterogeneity can also lead to specialization in non-domestic assets (see Van Nieuwerburgh and Veldkamp (2010) and De Marco et al. (2021)).

[11] See for instance Leuz et al. (2009), Mondria et al. (2010) Huang (2015) and Cziraki et al. (2021).

[12] For an overview of all advanced and emerging economies in our sample see table A1 in the appendix. Note that the survey provides forecasts as of 1989 for some countries. However, our sample period is limited by the GDP and inflation vintage series of actual outcomes provided by the IMF which starts in 1998.

[13] Consensus Economics clearly defines each macroeconomic indicator that it surveys.

Realized Outcomes.

Following the literature, we use first release data to compare forecast precision across forecasters. For each survey year, we use the realized outcome for yearly inflation and real GDP growth from the International Monetary Fund World Economic Outlook (IMF WEO) published in April of the subsequent year. This allows us to match the information set of the agents as closely as possible and avoids forecast errors that are due to data revisions. For example, to assess the accuracy of the 2013 real GDP growth forecast for Brazil from the January 2013 survey, we use the yearly GDP growth reported in the April 2014 IMF WEO as realized outcome. To assess the accuracy of the 2014 real GDP growth forecast for Brazil from the same January 2013 survey, we use the yearly GDP growth reported in April 2015. We conduct robustness checks with alternative vintages using IMF WEO published in September or in subsequent years.[14] Archived IMF WEO vintage data are available from 1998 onwards. Table A1 presents the list of variables and countries we study as well as the time range for which both forecast and realized data are available.

As is common in the literature, we trim observations, removing forecasts that are more than 5 interquartile ranges away from the median. The quantiles are calculated in two different ways. First, on the whole sample, but separately for emerging and advanced countries. Second, conditional on each country and date. This trimming ensures that our results are not driven by extreme outcomes, such as periods of hyperinflation, or by typos. It reduces the number of forecasts for current inflation and GDP by 4 and 1 percent, and those for future inflation and GDP by 10 and 7 percent, respectively. We conduct robustness checks with alternative trimming strategies.

Location.

Consensus Economics discloses the name of the forecasting institution. We use this name to match the Consensus Economics data to information about the location of the forecaster from Eikon (Refinitiv). Eikon provides the company tree structure of most forecasters in our dataset. The tree structure includes information about the countries of the headquarter as well as the subsidiaries and affiliates. If the forecaster was not listed in the Eikon database, we manually searched for this information on the internet. In the main analysis, we consider a forecaster to be foreign if neither its headquarter nor any of its subsidiaries are located in the country of the forecast. However, the information on the location is not time-varying and corresponds to the information accessed in 2021. This amounts to a measurement error that could bias downward the magnitude of the effect of location.

[14] Using alternative vintage series ensures that differences in forecasting precisions are not solely due to individual forecasters that anticipate revisions in actual GDP or inflation and therefore have a different forecasting target.

Table 2.1: Distribution of Observations across Forecasters conditional on Location and Scope

Location	**Scope**								
	National			**Multinational**			**Total**		
	N	Col %	Row %	N	Col %	Row %	N	Col %	Row %
Local	35,431	61.2	30.0	82,822	78.7	70.0	118,253	72.5	100.0
Foreign	22,452	38.8	50.0	22,446	21.3	50.0	44,898	27.5	100.0
Total	57,883	100.0	35.5	105,268	100.0	64.5	163,151	100.0	100.0

Note: The table shows the distribution of the forecasters conditional on their location and scope. “N” refers to the number of observations, “Col %” to the column percentage, and “Row %” to the row percentages, respectively. Forecasters are either local or foreign. Local forecasters have the headquarter or subsidiary in the country they forecast for, otherwise they are considered as a foreign forecaster. Multinational forecasters have subsidiaries in different countries than their headquarter is located in. National forecasters have only subsidiaries in the same country as the headquarter.

Forecast errors.

We use this information to construct forecast errors. The forecast errors with respect to the current year are defined as

$$Error^m_{ijt,t} = x_{jt} - E^m_{ijt}(x_{jt})$$

where t refers to the year, i is the forecaster, j is the country, $m = 1, .., 12$ is the month of the year when the forecast is produced, and x is either inflation of GDP growth. And the forecast errors with respect to the next year are defined as

$$Error^m_{ijt,t+1} = x_{jt+1} - E^m_{ijt}(x_{jt+1}).$$

Forecasters' Scope and Industry.

Furthermore, we characterize the scope of the forecasters. In more detail, we categorize forecasters with subsidiaries and headquarter all located in the same country as national forecasters. In contrast, we categorize forecasters with at least one subsidiary located in a country different from their headquarters as multinationals. Table 2.1 provides an overview of the distribution of observations across forecasters conditional on their location and scope.[15] Almost two third of the forecasts come from multinational forecasters, and almost three quarters are made by local forecasters. A higher proportion of forecasts by multinational forecasters is local, because multinationals are more likely to have a branch in the countries for which they produce forecasts.[16]

Besides data on location, Eikon provides information about the industry of the forecaster which we manually verified. We use industry information of the headquarter to distinguish non-financial from financial forecasters.

[15]As the scope variable is based on the location information, this variable is not time varying.

[16]We provide a similar distribution table for the number of country-forecaster pairs in Table A3 in the Appendix.

3. Foreign Forecasters Update their Forecasts Less Often

Before considering forecast errors, it is informative to examine forecast updating. Here, we explore the hypothesis that local forecasters update their forecasts more often than foreign forecasters. To do so, we compute the number of published forecasts for each year-forecaster-country unit, which we denote N_{ijt}. The distribution of these numbers of yearly forecasts is provided in Figure B1 in the Appendix. Most forecasters publish their forecasts 12 times a year, but some publish less often. A higher proportion of local forecasters publish a forecast at least 7 times a years.

We test formally whether foreign forecasters publish forecasts less often by taking the log of N_{ijt} and estimating

$$\ln(N_{ijt}) = \tilde{\delta}_{it} + \bar{\delta}_{jt} + \beta \text{Foreign}_{ij} + \varepsilon_{ijt}\,, \tag{3.1}$$

where $\tilde{\delta}_{it}$ and $\bar{\delta}_{jt}$ are respectively forecaster-year and country-year fixed effects. Foreign_{ij} is a dummy that takes the value of 1 if forecaster i is foreign to country j, and 0 otherwise.

The results are reported in Table 3.1. In the absence of fixed effects (column (2)), there is no significant difference in publication frequency between local and foreign forecasters. But as soon as we include country and forecaster fixed effects (column (3)), it appears that foreign forecasters publish significantly less: their number of releases are 12% to 14% smaller depending on the forecasted variable. When including the country-year and the forecaster-year fixed effects (column (4)), the foreign forecasters still appear to publish their forecasts 10% to 12% less often than local forecasters. The difference in publication frequency between local and foreign forecasters is smaller when considering GDP growth (as opposed to inflation), and when forecasting (as opposed to nowcasting).

Note that forecasters may publish a forecast without necessarily updating it, so the publication frequency is an imperfect measure of the updating frequency. We thus compute the number of yearly forecasts when considering only "distinct" forecasts, that is, forecasts that differ from the previous release. This is also an imperfect measure of updating frequency, since an identical forecast does not necessarily reflect the absence of new information. It may simply reflect the fact that the new information is not upsetting. But we assume that the arrival of new, upsetting information is time and country specific and will thus be captured by the fixed effects.

Figure B2 in the Appendix provides the distribution of the number of yearly forecasts using only "distinct" forecasts. Foreign forecasters are more likely to provide the same forecast for a whole year. But foreign forecasters are also slightly more likely to update their forecasts at the higher frequency of 11 to 12 times a year. We use this measure to estimate Equation (3.1) and

Table 3.1: Forecast Error conditional on Location of the Forecaster

	(1)	(2)	(3)	(4)	(5)
Variable	Coefficient				
CPI_t	Foreign	–0.03	–0.14***	–0.12***	–0.12***
		(0.06)	(0.04)	(0.04)	(0.04)
	N	16,427	16,346	10,857	10,822
	R^2	0.00	0.23	0.53	0.50
GDP_t	Foreign	–0.04	–0.13***	–0.10***	–0.10***
		(0.06)	(0.04)	(0.03)	(0.03)
	N	17,091	17,008	11,240	11,238
	R^2	0.00	0.23	0.54	0.52
CPI_{t+1}	Foreign	–0.02	–0.14***	–0.11***	–0.10**
		(0.06)	(0.05)	(0.04)	(0.04)
	N	15,371	15,286	10,082	9,950
	R^2	0.00	0.25	0.53	0.50
GDP_{t+1}	Foreign	–0.02	–0.12***	–0.10**	–0.08**
		(0.06)	(0.04)	(0.04)	(0.03)
	N	16,048	15,961	10,464	10,342
	R^2	0.00	0.25	0.53	0.51
	Country and Forecaster FE	No	Yes	No	No
	Country × Year FE	No	No	Yes	Yes
	Forecaster × Year FE	No	No	Yes	Yes

Note:The table shows the results of the regression of the number of forecast updates within a year on the location of the forecaster with different fixed-effects specifications. In columns (2) to (4), we consider that any available forecast is an update. In column (5), we only consider a forecast as an update if its value differs from the last available forecast. All standard errors are clustered at the country and forecaster level.

report the results in column (5) of Table 3.1. The results are, in fact, barely changed. All in all, foreign forecasters publish their forecasts about 10% less frequently than local forecasters.

4. Foreign Forecasters Make More Errors

In this section, we analyze the forecasters' errors and find that foreign forecasters make larger errors than local ones.

As preliminary descriptive evidence, Figure C1 shows the density of forecast errors for each group of forecasters. The forecast errors are distributed around 0 for both local and foreign forecasters. However, the distribution of forecast errors for foreign forecasters is wider than for local forecasters. A wider distribution of errors points towards less precise forecasts, as fewer errors are distributed close to zero.

In a first more formal test, we investigate whether the variance of forecast errors is indeed larger for foreign forecasters than the variance of local forecasters. For this, we perform a simple test of the equality of the variance of the annual average of forecast errors across locations, defined as $\frac{1}{12}\sum_{m=1}^{12} Error^m_{ijt,t+h}$, for $h = 0, 1$. We use the annual average here to take into account a potential high correlation of the errors within a year, which could bias the test. We implement Levene's test for equality of variances (Levene, 1960). The null hypothesis, H_0, is that variances are equal $\sigma^2_{\text{FE}_{\text{Local}}} = \sigma^2_{\text{FE}_{\text{Foreign}}}$, versus the alternative hypothesis of unequal variances, H_A, $\sigma^2_{\text{FE}_{\text{Local}}} \neq \sigma^2_{\text{FE}_{\text{Foreign}}}$.[17]

Table 4.1: Test for differences in Variance of Forecast Error

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variable	Sample	N Local	N Foreign	σ_{Local}	σ_{Foreign}	F-test	p-value
CPI_t	All sample	11,908	4,519	0.79	0.94	82.77	< 0.001
	Advanced Economies	5,655	1,278	0.42	0.49	29.39	< 0.001
	Emerging Economies	6,253	3,241	1.02	1.07	2.78	0.095
	Multinatonal firms	8,435	2,320	0.77	0.95	77.45	< 0.001
	National firms	3,473	2,199	0.86	0.93	12.74	< 0.001
	Financial Sector	8,005	1,274	0.78	1.04	69.99	< 0.001
	Non-Fincial Sector	1,828	2,158	0.74	0.83	19.10	< 0.001
GDP_t	All sample	12,390	4,701	1.15	1.44	131.49	< 0.001
	Advanced Economies	5,762	1,274	0.69	0.87	53.80	< 0.001
	Emerging Economies	6,628	3,427	1.44	1.60	15.36	< 0.001
	Multinatonal firms	8,690	2,424	1.11	1.51	148.38	< 0.001
	National firms	3,700	2,277	1.25	1.36	8.83	0.003
	Financial Sector	8,269	1,348	1.14	1.60	117.08	< 0.001
	Non-Fincial Sector	1,858	2,217	0.99	1.32	58.50	< 0.001
CPI_{t+1}	All sample	11,231	4,140	1.76	2.09	112.73	< 0.001
	Advanced Economies	5,382	1,171	0.91	1.04	22.85	< 0.001
	Emerging Economies	5,849	2,969	2.27	2.38	6.49	0.011
	Multinatonal firms	7,971	2,151	1.79	2.07	57.65	< 0.001
	National firms	3,260	1,989	1.68	2.10	60.22	< 0.001
	Financial Sector	7,582	1,192	1.81	2.17	44.28	< 0.001
	Non-Fincial Sector	1,711	1,964	1.66	2.00	45.50	< 0.001
GDP_{t+1}	All sample	11,707	4,341	2.45	3.10	109.10	< 0.001
	Advanced Economies	5,472	1,168	1.60	1.86	18.66	< 0.001
	Emerging Economies	6,235	3,173	3.00	3.45	15.99	< 0.001
	Multinatonal firms	8,206	2,275	2.36	3.24	123.84	< 0.001
	National firms	3,501	2,066	2.64	2.94	5.81	0.016
	Financial Sector	7,831	1,281	2.43	3.41	99.87	< 0.001
	Non-Fincial Sector	1,737	2,023	1.95	2.82	53.02	< 0.001

Note: The table shows Levens' test of equal variances of forecast errors between local and foreign forecasters. The Null hypothesis posits that the variance of the forecast errors of local forecasters is equal to the variance of foreign forecasters' forecast errors. The alternative hypothesis is that the variances are not equal. In the rows we report the test statistics for different subsamples.

Table 4.1 reports the results. In column (1), we define different sub-samples. We split the sample into advanced and emerging countries, multinational and national forecasters, financial and non-financial forecasters. Column (2) and (3) show the number of observations for local and foreign forecasters, respectively. Column (4) and (5) show the standard deviation of the forecast error conditional on the location. Column (6) reports the F-statistics and column (7) the corresponding p-value.

The null hypothesis of equal variances can be rejected at the 1% significance level for both GDP growth and inflation, and for both horizons. This result holds over the entire sample

[17]Note that there are different ways for calculating the test statistic for equal variances, namely using the mean, median or trimmed mean. We observe very little differences across these methods which is why we report the results of the test statistics calculating with the mean.

and for most subsamples. Only for a few subsamples, we observe a lower significance level, such as for current and future inflation forecasts over the subsample of emerging economies as well as current and future GDP forecasts over the subsample of national firms.

Note, however, that the test for equal variance does not allow to control for country- and forecaster-specific characteristics. For instance, Table 2.1 shows that a higher proportion of forecasts by multinational forecasters are local. Given that multinationals are also more likely to have well-endowed forecasting departments, local forecasts could artificially appear more accurate if we do not control for forecasters' characteristics.[18] For this reason, we estimate different fixed-effects model with alternative measures of the forecast error magnitude, and control for forecaster-, date- and country-specific characteristics by exploiting the panel structure of our data.

As a first measure of the forecast error distribution, we estimate the standard deviation $\sigma^m_{\mathrm{FE},i,j}$ of the forecast error for every forecaster-country-month triplet (m, i, j) for current and future forecasts separately. We discard forecaster-country-month triplets with less than 10 observations. We take the log of $\sigma^m_{\mathrm{FE},i,j}$ and estimate

$$\ln(\sigma^m_{\mathrm{FE},i,j}) = \delta^m + \tilde{\delta}_i + \bar{\delta}_j + \beta \text{Foreign}_{ij} + \varepsilon^m_{ij}\,, \tag{4.1}$$

where δ^m, $\tilde{\delta}_i$ and $\bar{\delta}_j$ are respectively month-of-year, forecaster and country fixed effects. Foreign_{ij} is a dummy that takes the value of 1 if forecaster i is foreign to country j, and 0 otherwise.

Table 4.2 reports the coefficient β for different specifications of the model. The standard deviation of forecast errors is higher when the forecasts are produced by a foreign forecaster than when they are produced by a local one. This finding is robust across different fixed-effect specifications. In the most conservative specification (with country, forecaster and month-of-year fixed effects), foreign forecasts are 6% to 14% higher than local forecasts. The difference between local and foreign forecasts' precision is larger for inflation than for GDP growth, and for the current than for the future year.

In this specification, we control for country, forecaster and month-of-year characteristics, but not for the time period. Ignoring time-specific characteristics could bias our results if, for instance, more foreign forecasts are produced in times of turmoil and uncertainty, where all forecasters will make more mistakes. Therefore, as a second measure of the forecast error distribution, we calculate the log absolute value of the forecast error, which is time-varying.[19]

[18]Similarly, even though Consensus Economics uses consistent definitions for macroeconomic indicators in their monthly survey for all forecasters, forecasters may employ divergent definitions that could bias the results.

[19]For absolute forecast errors smaller than 0.001 percentage point, we assign the value of $\ln(0.001)$ to keep all observations in the sample. The results are robust for different thresholds.

Table 4.2: Standard Deviation of the Forecast Error conditional on Location of the Forecaster

	(1)	(2)	(3)	(4)
Variable	Coefficient	$\ln(\sigma^m_{\text{FE},i,j})$	$\ln(\sigma^m_{\text{FE},i,j})$	$\ln(\sigma^m_{\text{FE},i,j})$
CPI_t	Foreign	0.12***	0.13**	0.14***
		(0.04)	(0.05)	(0.05)
	N	6,107	6,097	6,097
	R^2	0.47	0.50	0.81
GDP_t	Foreign	0.06***	0.12**	0.11***
		(0.02)	(0.05)	(0.03)
	N	6,544	6,535	6,535
	R^2	0.49	0.51	0.89
CPI_{t+1}	Foreign	0.07***	0.06	0.06*
		(0.02)	(0.04)	(0.04)
	N	6,107	6,097	6,097
	R^2	0.79	0.83	0.86
GDP_{t+1}	Foreign	0.07***	0.06**	0.06**
		(0.02)	(0.03)	(0.03)
	N	6,544	6,535	6,535
	R^2	0.77	0.81	0.86
	Country FE	Yes	Yes	Yes
	Forecaster FE	No	Yes	Yes
	Month FE	No	No	Yes

Note: The table shows the regression of the log standard deviation of current and future CPI and GDP on the location of the forecaster with different fixed-effects specifications. The standard deviation is calculated by forecaster-country pair for each month. We neglect forecasters that have less than 10 observations for a given month. All standard errors are clustered on the country and forecaster level.

We use the logarithm of the absolute forecast error to give more weight to small differences around zero. The model we estimate is as follows.

$$\ln(|Error^m_{ijt,t}|) = \delta^m_{it} + \tilde{\delta}^m_{jt} + \beta\text{Foreign}_{ij} + \varepsilon^m_{ij,t}\,, \tag{4.2}$$

δ^m_{it} are forecaster-date fixed effects and $\tilde{\delta}^m_{jt}$ are country-date fixed effects. These fixed effects enable us to control for country-specific trends in volatility and forecaster-specific trends in forecasting performance.

Table 4.3 displays the results for CPI and GDP. In all specifications, foreign forecasts are significantly larger in absolute value than local forecasts. In the most conservative specification with country-date and forecaster-date fixed effects, the absolute value of foreign forecast errors is 9% larger for current inflation. The difference is smaller for current GDP growth (6%) and for future inflation (7%). For future GDP growth, there is no significant difference between local and foreign forecasts.

Are these excess errors due to the relatively less frequent updating of foreign forecasters documented in Section 3? To answer this question, we repeat the last exercise using only the forecasts that differ from their previous release. The results are reported in Table C1 in the Appendix. The results are very similar, except that the coefficient for current GDP growth and future inflation are slightly lower (5% instead of 6% and 7% in the last column). This implies

Table 4.3: Forecast Error conditional on Location of the Forecaster

	(1)	(2)	(3)	(4)
Variable	Coefficient			
CPI_t	Foreign	0.26***	0.10***	0.09***
		(0.08)	(0.03)	(0.02)
	N	153,089	153,066	99,228
	R^2	0.01	0.14	0.62
GDP_t	Foreign	0.27***	0.11***	0.06**
		(0.08)	(0.03)	(0.02)
	N	160,971	160,947	103,866
	R^2	0.01	0.15	0.66
CPI_{t+1}	Foreign	0.27***	0.09***	0.07***
		(0.06)	(0.03)	(0.02)
	N	140,177	140,152	90,693
	R^2	0.01	0.14	0.67
GDP_{t+1}	Foreign	0.15*	0.08**	0.01
		(0.08)	(0.03)	(0.02)
	N	147,885	147,860	95,508
	R^2	0.00	0.16	0.72
	Country and Forecaster FE	No	Yes	Yes
	Country $\times$ Date	No	No	Yes
	Forecaster $\times$ Date FE	No	No	Yes

Note: The table shows the regression of the log absolute forecast error of current CPI and GDP on the location of the forecaster with different fixed-effects specifications. All standard errors are clustered on the country, forecaster and date level.

that foreign forecasters make more errors also conditional on updating.

5. What Explains the Foreigners' Errors?

To explore what explains the foreigners' errors, we lay down a simple noisy information model. We explore two potential sources of heterogeneity between local and foreign forecasters: behavioral biases and information asymmetry. We rule out differences in behavioral biases using rational expectation tests that are now common in the literature. We then establish the presence of asymmetric information by using two tests that are robust to common behavioral biases and to public signals.

5.1 A Simple Noisy Information Model

We consider a set of N professional forecasters indexed by $i = 1, .., N$ who form expectations on J countries indexed by $j = 1, .., J$. We denote by x_{jt} the variable that is forecasted. Denote by $S(j)$ the set of forecasters who form expectations on country j. Forecaster $i \in S(j)$ can belong either to the group of local forecasters $S^l(j)$ or to the group of foreign forecasters $S^f(j)$. We denote by $N(j)$, $N^l(j)$ and $N^f(j)$ the number of elements in $S(j)$, $S^l(j)$ and $S^f(j)$ respectively.

We assume that x_{jt}, the yearly realization of x_j, follows an AR(1):

$$x_{jt} = \rho_j x_{jt-1} + \epsilon_{jt} \tag{5.1}$$

with $\epsilon_{jt} \sim N(0, \gamma^{-1/2})$.

5.1.1 Information structure and behavioral biases

We consider an information structure and behavioral assumptions that are similar to Angeletos, Huo and Sastry (2020), except that we include public signals.

Information structure.

We assume that the information structure is country, month, and location-specific. Between month m of year $t-1$ and month m of year t, forecasters receive two types of signals: a public signal

$$\phi_{jt}^m = x_{jt} + (\kappa_j^m)^{-1/2} u_{jt}^m$$

observed by all forecasters, where $u_{jt}^m \sim N(0,1)$ is an i.i.d. aggregate noise shock and $\kappa_j^m > 0$ is the precision of the public signal, which is specific to country j and to month m, and a private signal

$$\varphi_{ijt}^m = x_{jt} + (\tau_{ij}^m)^{-1/2} e_{ijt}^m$$

that is observed only by forecaster i, where $e_{ijt}^m \sim N(0,1)$ is an i.i.d. idiosyncratic noise $\tau_{ij}^m > 0$ is the precision of the private signal, which is specific to country j, to month m, but also to forecaster i. Through the law of large numbers we have $\frac{1}{N(j)} \sum_{i \in S(j)} \epsilon_{ijt}^m = 0$, $\frac{1}{N^l(j)} \sum_{i \in S^l(j)} \epsilon_{ijt}^m = 0$ and $\frac{1}{N^f(j)} \sum_{i \in S^f(j)} \epsilon_{ijt}^m = 0$. Local and foreign forecasters differ through the precision of their private information τ_{ij}^m: $\tau_{ij}^m = \tau_{jl}^m$ if $i \in S^l(j)$ and $\tau_{ij}^m = \tau_{jf}^m$ if $i \in S^f(j)$. We assume that, for a given month m, ϵ_{ijt}^m and u_{jt}^m are mutually and serially independent. This means for instance that the noises in the signals of month m from year t are not correlated with the noises in the signals of month m from year $t-1$. But we do not impose that the noises are serially uncorrelated within a given year.[20]

Behavioral biases.

We consider two behavioral biases: over-extrapolation and over-confidence. Over-extrapolation (or under-extrapolation) consists in distorted beliefs about the persistence of shocks ρ_j. We denote forecaster i's belief about the persistence of x_{jt} by $\hat{\rho}_{ij}$. We assume that local and

[20] This type of information structure would arise if forecasters were receiving independent signals every month. In that case, the information received between month m of year $t-1$ and month m of year t would be represented by a 12-month moving average of the monthly signals, which is serially correlated on a month-on-month basis, but not on a year-on-year basis.

foreign forecasters may have different beliefs, so that $\hat{\rho}_{ij} = \hat{\rho}_{jl}$ if $i \in S^l(j)$ and $\hat{\rho}_{ij} = \hat{\rho}_{jf}$ if $i \in S^f(j)$. Over-confidence (or under-confidence) consists in distorted beliefs about the precision of private signals τ^m_{jk}. We denote forecaster i's belief about her precision by $\hat{\tau}^m_{ij}$. Again, we assume that local and foreign forecasters may have different beliefs, so that $\hat{\tau}^m_{ij} = \hat{\tau}^m_{jl}$ if $i \in S^l(j)$ and $\hat{\tau}^m_{ij} = \hat{\tau}^m_{jf}$ if $i \in S^f(j)$.

Expectations.

In month m of year t, forecasters build a "synthetic" signal out of the public and private signals:

$$\begin{aligned} s^m_{ijt} &= h^m_{ij}\phi^m_{jt} + (1 - h^m_{ij})\varphi^m_{ijt} \\ &= x_{jt} + v^m_{ijt} \end{aligned} \tag{5.2}$$

with

$$v^m_{ijt} = h^m_{ij}(\kappa^m_j)^{-1/2}u^m_{jt} + (1 - h^m_{ij})(\tau^m_{ij})^{-1/2}e^m_{ijt} \tag{5.3}$$

and $h^m_{ij} = \kappa^m_j/(\kappa^m_j + \hat{\tau}^m_{ij})$, so that $E^m_{ijt}(x_{jt}|\phi^m_{jt}, \varphi^m_{ijt}) = (\kappa^m_j + \hat{\tau}^m_{ij})/(\gamma_j + \kappa^m_j + \hat{\tau}^m_{ij})s^m_{ijt}$.
Between month m of year $t-1$ and month m of year t, the forecasters update their expectations in the following way:

$$E^m_{ijt}(x_{jt}) = (1 - G^m_{ij})\hat{\rho}_{ij}E^m_{ijt-1}(x_{jt-1}) + G^m_{ij}s^m_{ijt} \tag{5.4}$$

where G^m_{ij} is the Kalman gain that is consistent with forecaster i' beliefs about the persistence of x_{jt} and about the precision of their signal.
We define the forecast revisions between month m of year $t-1$ and month m of year t as

$$Revision^m_{ijt} = E^m_{ijt}(x_{jt}) - E^m_{ijt-1}(x_{jt}) \tag{5.5}$$

and the error as

$$Error^m_{ijt,t} = x_{jt} - E^m_{ijt}(x_{jt}) \tag{5.6}$$

5.1.2 The variance of errors

Consider the case with no behavioral biases. Forecasters with less precise information make more errors on average. This derives from the forecasters' optimal use of information. In fact, the variance of errors can be related to the Kalman gain, as stated in the following proposition (see the proof in Appendix G.1):

Proposition 1 *In the absence of behavioral biases ($\hat{\rho}_{ij} = \rho_j$ and $\hat{\tau}^m_{ij} = \tau^m_{ij}$), the variance of* *errors is given by:*

$$\begin{array}{lll} V(Error^m_{ijt,t-1}) & = V[x_{jt} - E^m_{ijt-1}(x_{jt})] & = \frac{\gamma^{-1}}{1-\rho_j^2(1-G^m_{ij})} \\ V(Error^m_{ijt,t}) & = V[x_{jt} - E^m_{ijt}(x_{jt})] & = \frac{\gamma^{-1}(1-G^m_{ij})}{1-\rho_j^2(1-G^m_{ij})} \end{array} \tag{5.7}$$

Both variances are decreasing in G^m_{jk}.

Since G^m_{ij} is increasing in τ^m_{ij}, then the variances are decreasing in τ^m_{ij}.
But asymmetric information is not the only potential source of differences in variances. Consider now the case with behavioral biases. The Kalman filter is a minimum mean-square error estimator. Therefore, mis-specified statistical and parametric inputs to the estimator will increase the variance of errors as compared to the well-specified estimator. Therefore, the difference in variances may be due to differences in behavioral biases. In the remainder of the section, we use model-based tests to detect differences in behavioral biases and differences in information.

5.2 Testing for Differences in Behavioral Biases

BGMS regressions.

Here we examine whether local and foreign forecasters differ systematically in the way they form expectations. Following Angeletos, Huo and Sastry (2020), we consider two behavioral biases that go a long way in explaining survey forecasts: over-extrapolation ($\hat{\rho}_{jk} \neq \rho_j$) and over-confidence ($\hat{\tau}^m_{ij} \neq \tau^m_{ij}$). We rely on regressions popularized by Bordalo et al. (2020) and Kohlhas and Broer (2020) to assess the presence of such biases among forecasters:

$$Error^m_{ijt} = \beta^{BGMSm}_{ij} Revision_{ijt} + \delta^m_{ij} + \lambda^m_{ijt} \tag{5.8}$$

where β^{BGMSm}_{ij} is a country, month and forecaster specific coefficient, δ^m_{ij} are country-month-forecaster fixed effects and λ^m_{ijt} is an error term.
Following Angeletos, Huo and Sastry (2020), we can show that these coefficients are related to the deviations of the beliefs $\hat{\rho}_{ij}$ and $\hat{\tau}_{ij}$ from their true counterparts (see the proof in Appendix G.2):

Proposition 2 *Estimating Equation (5.8) for each $i = 1,..N$, $j = 1,..J$ and $m = 1,..12$ by OLS gives the following coefficients:*

$$\beta^{BGMSm}_{ij} = -(\hat{\rho}_{ij} - \rho_j)\beta^m_{1ij} - [(\tau^m_{ij})^{-1} - (\hat{\tau}^m_{ij})^{-1}]\beta^m_{2ij}$$

β^m_{1ij} and β^m_{2ij} are described in the Appendix. They depend on the country-invariant parameters κ^m_j and ρ_j but also on the forecaster-specific beliefs $\hat{\tau}^m_{ij}$ and $\hat{\rho}_{ij}$.

A negative coefficient reflects an over-reaction of forecasters to their information. This over-reaction can arise from over-confidence ($\hat{\tau}_{ij}^m - \tau_{ij}^m > 0$) or from over-extrapolation ($\hat{\rho}_{ij} - \rho_j > 0$).[21]

While a non-zero coefficient can help detect the presence of behavioral biases, it suffers from one drawback in our context: the coefficient is a non-linear and potentially non-monotonic function of $\hat{\tau}_{ij} - \tau_{ij}$, $\hat{\rho}_{ij} - \rho_j$, the biases, but also of τ_{ij}, the precision of private signals. Interpreting differences in coefficients is therefore not easy.

To help our interpretation of the results, we consider a first-order expansion of the BGMS coefficient around close-to-zero and symmetric biases (see the proof in Appendix G.3):

Corollary 1 *The coefficient β_{ij}^{BGMSm} can be approximated at the first-order around $(\hat{\tau}_{ij}^m)^{-1} = (\tau_{ij}^m)^{-1} = (\tau_j^m)^{-1}$, where τ_j^m is the average level of precision and $\hat{\rho}_{ij} = \hat{\rho}_j = \rho_j$ as follows:*

$$\beta_{ij}^{BGMSm} \simeq -(\hat{\rho}_{ij} - \rho_j)\hat{\beta}_1^m - [(\tau_{ij}^m)^{-1} - (\hat{\tau}_{ij}^m)^{-1}]\hat{\beta}_2^m$$

where $\hat{\beta}_1^m$ and $\hat{\beta}_2^m$ are strictly positive and independent of $\hat{\rho}_{ij}$, τ_{ij}^m and $\hat{\tau}_{ij}^m$.

Therefore, a more negative BGMS coefficient will be interpreted as reflecting differences in either over-confidence or over-extrapolation.

We estimate Equation (5.8) using the mean-group methodology, under different assumptions about the homogeneity of the β^{BGMS} coefficient. We first assume that the coefficients only differ across countries and between local and foreign forecasters. We then allow the coefficients to differ across country-forecaster pairs. Finally, we allow the coefficients to differ across month within a country-forecaster pair. In each of these specifications, we collect the β^{BGMS} coefficients and test for significant differences between local and foreign forecasters by regressing the coefficient on the Foreign dummy, controlling for country, forecaster and month fixed effects, when possible. A significant coefficient for the Foreign dummy would indicate that there are systematic differences in behavioral biases. When allowing the coefficients to differ across country-forecaster pairs, we restrict the sample to the pairs providing forecasts for at least 10 years.

The results are displayed in Table 5.1. In all specifications, there is no systematic difference between local and foreign forecasters. Interestingly, the average coefficient is positive for both inflation and GDP growth in our most conservative specification (columns (5) and (6)), suggesting that forecasters under-react to news on average. This might seem in contradiction with previous evidence, which has found over-reaction, especially for inflation (Bordalo et al., 2020; Kohlhas and Broer, 2022; Angeletos et al., 2021). However, note that previous evidence has focused on the Survey of Professional Forecasters, which provides forecasts for the US.

[21] In Bordalo et al. (2020), this over-reaction can be due to diagnostic expectations.

Our estimated parameters are in fact highly heterogeneous (see Figure D1 in the Appendix), and in particular, they are heterogeneous across countries (see Figure D2 in the Appendix). Focusing on the US, we find that the inflation forecasts feature over-reaction on average, which is consistent with previous evidence. GDP growth forecasts do not feature systematic over- or under-reaction, which is also consistent with the existing evidence.

Table 5.1: Behavioral Biases - BGMS regressions

	(1)	(2)	(3)	(4)	(5)	(6)
Coefficient	CPI_t	GDP_t	CPI_t	GDP_t	CPI_t	GDP_t
Average Locals	–0.01**	0.06***	0.01	0.04***	0.03***	0.09***
	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)
Foreign	0.00	–0.02	0.00	0.03	0.00	0.04
	(0.01)	(0.01)	(0.02)	(0.02)	(0.03)	(0.03)
N	102	102	364	393	4,979	5,373
R^2	0.96	0.94	0.71	0.76	0.43	0.46
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Forecaster FE	No	No	Yes	Yes	Yes	Yes
Month FE	No	No	No	No	Yes	Yes
Mean-group by country and location	Yes	Yes	No	No	No	No
Mean-group by country and forecaster	No	No	Yes	Yes	No	No
Mean-group by cty, forc. and month	No	No	No	No	Yes	Yes

Note:The table shows the results of a regression of the β^{BGMS} coefficients on the Foreign dummy, where the β^{BGMS} are estimated using Equation (5.8) on different sub-groups of our sample. *Average Locals* corresponds to the constant term (or average fixed effect). *Foreign* corresponds to the coefficient of the Foreign dummy. The observations are clustered at the country level in specifications (1) and (2), and at the country and forecaster levels in specifications (3) to (6).

Perceived persistence.

A non-negative BGMS coefficient can arise both from distorted beliefs on the precision of private signals and from distorted beliefs on the persistence of the shocks. We have shown that these BGMS coefficients do not differ systematically between local and foreign forecasters. However, this does not imply that foreign forecasters have similar over-/under-confidence and over-/under-extrapolation. A similar result would arise if the relative over-/under-confidence of foreign forecasters compensates their relative over-/under-extrapolation. We examine more directly whether the beliefs on persistence are similar.

To do this, we use the relation between the forecasts on current and future variables implied by our model:

$$E^m_{ijt}(x_{jt+1}) = \hat{\rho}_{ij} E^m_{ijt}(x_{jt}) \tag{5.9}$$

We estimate Equation (5.9) using the same mean-group methodology. While in our model $\hat{\rho}_{ij}$ is specific to a country-forecaster pair and is independent of the month of the year, we allow it to differ across months as well. Indeed, while in our model, all the innovations to inflation have the same persistence, in reality, there could be some components of inflation that are purely transitory. We cannot exclude that forecasters learn about the transitory component over the year. That would affect the month-specific correlation between the nowcast and the forecast.

Table 5.2: Behavioral Biases - Over-extrapolation regressions

	(1)	(2)	(3)	(4)	(5)	(6)
Coefficient	CPI_t	GDP_t	CPI_t	GDP_t	CPI_t	GDP_t
Average Locals	0.41***	0.35***	0.41***	0.35***	0.39***	0.35***
	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)
Foreign	–0.00	0.01	0.02	0.04*	0.03	0.04
	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)
N	102	102	404	428	6,097	6,535
R^2	0.96	0.97	0.65	0.78	0.54	0.66
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Forecaster FE	No	No	Yes	Yes	Yes	Yes
Month-of-year FE	No	No	No	No	Yes	Yes
Mean-group by country and location	Yes	Yes	No	No	No	No
Mean-group by country and forecaster	No	No	Yes	Yes	No	No
Mean-group by country, forecaster and month	No	No	No	No	Yes	Yes

Note:The table shows the results of a regression of the perceived autocorrelation coefficients $\hat{\rho}$ on the Foreign dummy, where the $\hat{\rho}$ is estimated using Equation (5.9) on different sub-groups of our sample. *Average Locals* corresponds to the constant term (or average fixed effect). *Foreign* corresponds to the coefficient of the Foreign dummy. The obervations are clustered at the country level in specifications (1) and (2), and at the country and forecaster levels in specifications (3) to (6).

The results are reported in Table 5.2. In all specifications but one, the estimated perceived persistence is not significantly different for foreign forecasters. In column (6), where we allow the perceived persistence to vary across forecaster-country pairs, the foreign perceived persistence of GDP growth is significantly higher than the local one. However, when we allow the perceived persistence to vary across months as well, the difference is not significant anymore.

In the Appendix, we additionally examine whether forecasters differ in the way they use public news, since Kohlhas and Broer (2022) and Gemmi and Valchev (2022) show that forecasters typically under-react to public news. In Tables D1 and D2, we examine over-/under-reaction to public news, by examining regressions of forecast errors on public news, using two different measures of public news: the past consensus and the last vintage of realized outcome. A negative (positive) coefficient implies that forecasters over-react (under-react) to public news. Again, we do not find any systematic difference in behavioral biases.[22]

All in all, foreign and local forecasters do not have significantly different biases. From now on, we thus assume common behavioral parameters $\hat{\rho}_{jl} = \hat{\rho}_{jf} = \hat{\rho}_j$ and $\hat{\tau}^m_{jl} = \hat{\tau}^m_{jf} = \hat{\tau}^m_j$. In

[22]Interestingly, in our most conservative specification (columns (5) and (6)), we find systematic under-reaction to the past consensus (in Table D1, we can see that forecasters under-react to the past consensus on both GDP growth and inflation, as both average coefficients are positive), but not systematic under-reaction to the last vintage (in Table D2, we can see that forecasters only under-react to the last vintage of inflation and over-react to the last vintage of GDP growth, as only the average coefficient is positive for inflation and negative for GDP growth). This is consistent with the evidence provided by Gemmi and Valchev (2022), which suggests that forecasters tend to differentiate their forecasts from the other forecasters'.

the next sub-section, we examine differences in information frictions under this assumption.

5.3 Testing for Asymmetric Information

Consensus regressions that consist regressing the consensus error (i.e., the average error) on the consensus revision (i.e. the average revision) as in Coibion and Gorodnichenko (2015) are commonly used to detect information frictions. A positive coefficient indicates deviations from full information. Can we use these regressions to identify differences in information frictions between local and foreign forecasters? We show here that the relation between the precision of information and the coefficient of the consensus regression is non-monotonic in the presence of public signals. Therefore, even in the absence of behavioral biases, differences in the coefficient of the consensus regression are not a good indicator of the degree of information asymmetry. We propose two alternative tests that are robust to public signals.

5.3.1 Consensus regressions

Suppose that we performed the consensus regression as in Coibion and Gorodnichenko (2015) on both group of forecasters, that is, using the population of foreign forecasts on the one hand and the population of the local forecasts on the other, and then comparing the coefficients. In this case, what would we be identifying?

In our setup, this regression can be written, for each $j = 1, ..J$, $m = 1, ..12$ and $k = l, f$, where l refers to the local forecasts' population by l and f referes to the foreign forecasts's population:

$$Error^m_{jkt} = \beta^{CGm}_{jk} Revision^m_{jkt} + \delta^m_{jk} + \lambda^m_{jkt} \tag{5.10}$$

$Error^m_{jkt} = \frac{1}{N^k(j)} \sum_{i \in S^k(j)} Error^m_{ijt}$, $Revision^m_{jkt} = \frac{1}{N^k(j)} \sum_{i \in S^k(j)} Revision^m_{ijt}$, are the consensus error and the consensus revision in location $k = l, f$, δ^m_{jk} are country-month-location fixed effects and λ^m_{jkt} is an error term. The estimated parameter β^{CGm}_{jk} is a function of the deep parameters.

Table 5.3 displays the results of the estimation of β^{CGm}_{jk} using the mean-group estimator, under different assumptions on the heterogeneity of β^{CGm}_{jk}. In columns (1) and (2), we assume that β^{CGm}_{jk} differs across countries and locations. In columns (3) and (4), we assume that β^{CGm}_{jk} can also differ across months. While the β^{CGm}_{jk} coefficient is positive on average, as is expected, there does not appear to be any significant difference between foreign and local coefficients.

This does not necessarily mean that there are no information asymmetries between local and foreign forecasters. Indeed, the following proposition shows that, in the presence of public information, the relation between β^{CG} and the precision of private information is not monotonic (see the proof in Appendix G.4).

Table 5.3: Information Asymmetries - Consensus regressions

	(1)	(2)	(3)	(4)
Coefficient	CPI$_t$	GDP$_t$	CPI$_t$	GDP$_t$
Consensus	0.07***	0.12***	0.11***	0.16***
	(0.01)	(0.01)	(0.01)	(0.01)
Foreign	–0.01	–0.02	–0.00	–0.01
	(0.01)	(0.01)	(0.02)	(0.02)
N	102	102	1,223	1,224
R^2	0.93	0.94	0.50	0.53
Mean-group by country and location	Yes	Yes	Yes	Yes
Mean-group by country and month	No	No	No	No

Note:The table shows the results of a regression of the β^{CG} coefficients on the Foreign dummy, where the β^{CG} are estimated using equation (5.10) on different subgroups of our sample. *Consensus* corresponds to the constant term (or average fixed effect). *Foreign* corresponds to the coefficient of the Foreign dummy. The observations are clustered at the country level.

Proposition 3 *Suppose that there are no behavioral biases: $\hat{\rho}_{ij} = \rho_j$ and $\hat{\tau}_{ij}^m = \tau_j^m$, and that the precision parameters are identical within foreign forecasters and within local forecasters: $\tau_{ij}^m = \tau_{jl}^m$ if $i \in \mathcal{S}^l(j)$ and $\tau_{ij}^m = \tau_{jf}^m$ if $i \in \mathcal{S}^f(j)$, for all $j = 1, ..J$ and $m = 1, .., 12$. Estimating Equation (5.10) for each $j = 1, ..J$, $m = 1, .., 12$ and $k = l, f$ by OLS gives the following coefficients:*

$$\beta_{jk}^{CGm} = \frac{\frac{1-G_{jk}^m}{G_{jk}^m}\gamma^{-1} - [1 - \rho_j^2(1 - G_{jk}^m)]h_{jk}^2(\kappa_j^m)^{-1}}{\gamma^{-1} + [1 - \rho_j^2(1 - 2G_{jk}^m)](h_{jk}^m)^2(\kappa_j^m)^{-1}}$$

Note that $\beta_{jk}^{CGm} = (1 - G_{jk}^m)/G_{jk}^m$ when there is no public signal, which corresponds to the case studied by Coibion and Gorodnichenko (2015). The coefficient is directly related to the Kalman gain. A large coefficient implies a small Kalman gain and hence noisier information. Therefore, $\beta_{jl}^{CGm} < \beta_{jf}^{CGm}$ would imply that foreigners have noisier information ($\tau_{jf}^m > \tau_{jl}^m$). However, when $h_{jk}^m > 0$, β_{jk}^{CGm} depends on the variance of the fundamental shocks (γ^{-1}) and on the variance of the aggregate noise ($(\kappa_j^m)^{-1}$). β_{jk}^{CGm} is thus not a straightforward function of the information structure and it is not clear what to infer from $\beta_{jl}^{CGm} < \beta_{jf}^{CGm}$. This is due to the presence of aggregate noise. This aggregate noise, as discussed in Coibion and Gorodnichenko (2015), introduces a negative bias in the estimation of G_{jk}^m. Indeed, while the correlation between the error and the revision driven by the fundamental x_{jt} is positive, the public noise introduces a negative correlation. CG argue that because the bias is negative, a positive coefficient is still a sign of noisy information. However, in order to test for *differences* in the quality of private information by comparing β_{jl}^{CGm} and β_{jf}^{CGm}, we need that β_{jk}^{CGm} is a monotonic function of τ_{jk}^m.

Figure 5.1 shows that this is not the case. The figure describes how the precision of the private signal, τ_{jk}, affects the Kalman gain G_{jk}^m, the weight of public information h_{jk}^m and the coefficient β_{jk}^{CGm}. While the Kalman gain is increasing in the precision of private information, the weight of the public signal is decreasing. As a result, when the precision of the private signal goes to zero, forecasters put the highest possible weight on the public signal, and the coefficient is equal to zero. In this case, the public signal is the only valid source of information, so the individual forecasts correspond to the aggregate one. Rational expectations then imply a zero covariance between the aggregate revision and the aggregate error. When the precision of the private signal increases, the weight put on the public signal decreases, so the coefficient increases and becomes positive. Passed a certain threshold, the contribution of the public noise to the coefficient becomes negligible and the coefficient starts decreasing in τ_{jk}^m, driven by the increase in the Kalman gain, as in Coibion and Gorodnichenko (2015).

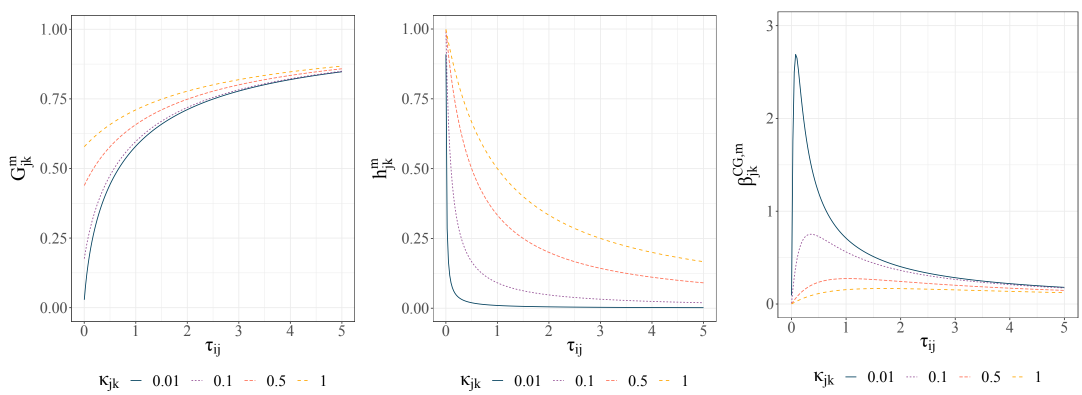

(A) G_{jk}^m (B) h_{jk}^m (C) $\beta_{jk}^{CG,m}$

Figure 5.1: The effect of τ_{jk}^m on β_{jk}^{CGm}

Note: The figure shows how the precision of the private signal, τ_{jk}, affects the Kalman gain G_{jk}^m, the weight of public information h_{jk}^m and the coefficient β_{jk}^{CGm}. The different colors in each plot correspond to different levels of the public signal precision κ_j^m.

We thus need tests that map to the degree of information frictions and that are robust to public information. We propose two such tests.

5.3.2 Fixed-effect regressions

For our first test of asymmetric information, we use an extension of the BGMS regression that controls for public noise. We use the following pooled regression, for each $j = 1, ..J$, $m = 1, .., 12$ and $k = l, f$:

$$Error_{ijkt}^m = \beta_{jk}^{FEm} Revision_{ijkt}^m + \delta_{jkt}^m + \lambda_{ijkt}^m \tag{5.11}$$

where δ_{jkt}^m are country-location-time fixed effects and λ_{ijkt}^m is an error term. The estimated parameter β_{jk}^{FEm} is a function of the deep parameters. We can show that, if $\hat{\rho}_{jk} = \hat{\rho}_j$ is homogeneous across groups, then differences in the estimated parameter $\beta_{jk}^{FE^m}$ across locations depend only on differences in G_{jk}^m (see the proof in Appendix G.5).

Proposition 4 *Suppose that the parameters are homogeneous within foreign forecasters and within local forecasters:* $\hat{\rho}_{ij} = \hat{\rho}_{jl}$, $\tau_{ij} = \tau_{jl}$ *and* $\hat{\tau}_{ij} = \hat{\tau}_j$, *if* $i \in \mathcal{S}^l(j)$, *and* $\hat{\rho}_{ij} = \hat{\rho}_{jf}$, $\tau_{ij} = \tau_{jf}$ *and* $\hat{\tau}_{ij} = \hat{\tau}_{jf}$, *if* $i \in \mathcal{S}^f(j)$. *Estimating Equation (5.11) for each* $j = 1, ..J$, $m = 1, .., 12$ *and* $k = l, f$ *by OLS gives the following coefficients:*

$$\beta_{jk}^{FEm} = -\frac{1 - \hat{\rho}_{jk}(1 - G_{jk}^m)}{1 - \hat{\rho}_{jk}(1 - 2G_{jk}^m)}$$

If forecasters have identical behavioral biases, i.e. $\hat{\rho}_{jl} = \hat{\rho}_{jf} = \hat{\rho}_j$ *and* $(\hat{\tau}_{jl}^m)^{-1} - (\tau_{jl}^m)^{-1} = (\hat{\tau}_{jf}^m)^{-1} - (\tau_{jf}^m)^{-1}$, *and if* $0 < \hat{\rho}_j < 1$, *then* $\beta_{jf}^{FEm} < \beta_{jl}^{FEm}$ *if and only if* $\tau_{jl}^m > \tau_{jf}^m$.

If the foreign and local forecasters have similar behavioral biases and if forecasters believe that there is some persistence in the process, then $\beta_{jf}^{FEm} < \beta_{jl}^{FEm}$ reflects an informational advantage for locals.

The estimated coefficient depends on the covariance between the error and the revision that is driven by idiosyncratic shocks. This covariance is necessarily negative: optimistic forecasters make a more negative error than pessimistic forecasters. As long as $\hat{\rho}_j$ is positive, this coefficient is more negative when information frictions are stronger (when the Kalman gain G_{jk}^m is lower). The lower G_{jk}^m, the more persistent is the forecast, as it incorporates new information in a slower fashion. This makes β_{jk}^{FEm} more negative because it increases the magnitude of the covariance between the revision and the forecast itself, which drives the error.[23,24]

We first estimate Equation (5.11) under the assumption that the β^{FE} coefficients differ across countries and locations, but not across months. We then regress these coefficients on the Foreign dummy and report the results in columns (1) and (2) of Table 5.4. We then estimate the equation under the assumption that the β^{FE} coefficients differ across countries, locations, and months. Similarly, we regress these coefficients on the Foreign dummy and report the results in columns (3) and (4). Note first that the estimated coefficients are negative on average, as predicted. Second, the coefficient for Foreign dummy is significantly negative for

[23] Note that the coefficient should be equal to β_{jk}^{BGMSm} in the absence of fixed effects. Why is it that adding fixed effects in the pooled regression results in a negative coefficient? It is because the fixed effects control for aggregate shocks (ϵ_{jt} and u_{jt}), which are not observed by forecasters at the time they revise their forecasts. A negative coefficient therefore is not a sign of a deviation from rational expectations.

[24] Note also that adding time fixed effects to the regression is equivalent to subtracting the cross-forecaster average from each side of the equation:

$$-\left(E_{ijkt}^m(x_{jt}) - E_{jkt}^m(x_{jt})\right) = \beta_{jk}^{FEm}(Revision_{ijkt}^m - Revision_{jkt}^m) + \lambda_{ijkt}^m$$

In that sense, this test is similar in spirit to Goldstein (2021), who proposes to measure information frictions by estimating the persistence of a forecaster's deviation from the mean:

$$\left(E_{ijkt}^m(x_{jt}) - E_{jkt}^m(x_{jt})\right) = \beta_{jk}^{Gm}\left(E_{ijkt-1}^m(x_{jt}) - E_{jkt-1}^m(x_{jt})\right) + \lambda_{ijkt}^m$$

$\beta_{jk}^{Gm} = 1 - G_{jk}^m$ is also directly and monotonically related to the degree of information frictions.

inflation. For GDP growth, it is negative as well, but smaller in magnitude and less significant (the p-value is higher than 10%). This is consistent with the preliminary evidence of Section 4 where we have shown that foreign forecasters made relatively more errors on inflation than on GDP growth.

Table 5.4: Information Asymmetries - Fixed-effect regressions

	(1)	(2)	(3)	(4)
Coefficient	CPI_t	GDP_t	CPI_t	GDP_t
Average Locals	–0.31***	–0.35***	–0.29***	–0.32***
	(0.00)	(0.00)	(0.00)	(0.00)
Foreign	–0.05***	–0.02	–0.05***	–0.02
	(0.01)	(0.01)	(0.01)	(0.01)
N	100	100	1,196	1,207
R^2	0.87	0.88	0.64	0.61
Country FE	Yes	Yes	Yes	Yes
Month FE	No	No	Yes	Yes
Mean-group by country and location	Yes	Yes	No	No
Mean-group by country, location and month	No	No	Yes	Yes

Note:The table shows the results of a regression of the β^{FE} coefficients on the Foreign dummy, where the β^{FE} are estimated using Equation (5.11) on different sub-groups of our sample. *Average Locals* corresponds to the constant term (or average fixed effect). *Foreign* corresponds to the coefficient of the Foreign dummy. The observations are clustered at the country level in specifications (1) to (4).

5.3.3 Foreign-local disagreement

Our second test of asymmetric information is based on disagreement between local and foreign forecasters. We define the disagreement between the local and foreign forecasters as follows:

$$Disagreement_{jt}^{m} = E_{jlt}^{m}(x_{jt}) - E_{jft}^{m}(x_{jt}) \tag{5.12}$$

where $E_{jkt}^{m}(x_{jt}) = \frac{1}{N(j)^k}\sum_{i\in\mathcal{S}^k(j)} E_{ijkt}(x_{jt})$ is the location-specific average expectation.
Consider now the following regression:

$$Disagreement_{jt}^{m} = \beta_j^{DISm} Revision_{jt}^{m} + \beta_j^{0m} x_{jt} + \beta_j^{1m} x_{jt-1} + \beta_j^{2m} E_{jlt-1}^{m}(x_{jt}) + \beta_j^{3m} E_{jft-1}^{m}(x_{jt}) + \delta_j^m + \lambda_{jt}^m \tag{5.13}$$

where $Revision_{jt}^{m} = \frac{1}{2}(Revision_{jlt}^{m} + Revision_{jft}^{m})$ is the average revision across locations for country j in year t and month m.
We can show that the sign of β_j^{DISm} depends on the relative precision of local forecasters versus foreign forecasters when the behavioral biases are homogeneous across locations (see the proof in Appendix G.6).

Proposition 5 *Suppose that the parameters are homogeneous within foreign forecasters and within local forecasters: $\hat{\rho}_{ij} = \hat{\rho}_{jl}$, $\tau_{ij} = \tau_{jl}$ and $\hat{\tau}_{ij} = \hat{\tau}_j$, if $i \in \mathcal{S}^l(j)$, and $\hat{\rho}_{ij} = \hat{\rho}_{jf}$, $\tau_{ij} = \tau_{jf}$ and $\hat{\tau}_{ij} = \hat{\tau}_{jf}$, if $i \in \mathcal{S}^f(j)$. Estimating Equation (5.13) for each $j = 1, ..J$ and $m = 1, .., 12$ by*

OLS gives the following coefficients:

$$\beta_j^{DISm} = \left(\frac{G_{jl}^m h_{jl}^m - G_{jf}^m h_{jf}^m}{G_j^m h_j^m} \right)$$

where $G_j^m h_j^m = \frac{1}{2}(G_{jl}h_{jl} + G_{jl}h_{jl})$.
If forecasters have identical behavioral biases, i.e. $\hat{\rho}_{jl} = \hat{\rho}_{jf} = \hat{\rho}_j$ *and* $(\hat{\tau}_{jl}^m)^{-1} - (\tau_{jl}^m)^{-1} = (\hat{\tau}_{jf}^m)^{-1} - (\tau_{jf}^m)^{-1}$, *then* $\beta_j^{DISm} < 0$ *if and only if* $\tau_{jl}^m > \tau_{jf}^m$.

Intuitively, β_j^{DISm} is negative if the foreign expectations are more sensitive to the public signal and hence to the public noise. This is the case if the foreign forecasters' private information is less informative than the local one.
We first estimate Equation (5.13) under the assumption that the β^{Dis} coefficients differ across countries, but not across months. We then test whether the coefficients are different from zero on average and report the results in columns (1) and (2) of Table 5.5. We then estimate the equation under the assumption that the coefficients differ across countries and months. Similarly, columns (3) and (4) report the significance tests. The disagreement coefficients are significantly negative on average for both inflation and GDP growth and in both specifications. Notably, the coefficient of GDP is smaller in magnitude, which is consistent with our previous results.

Table 5.5: Information Asymmetries - Disagreement regressions

	(1)	(2)	(3)	(4)
Coefficient	CPI_t	GDP_t	CPI_t	GDP_t
Disagreement	–0.09***	–0.07***	–0.09***	–0.07**
	(0.02)	(0.02)	(0.03)	(0.03)
N	51	51	611	612
R^2	0	0	–0.00	0
Mean-group by country	Yes	Yes	No	No
Mean-group by country and month	No	No	Yes	Yes

Note:The table shows the results of a regression of the β^{DIS} coefficients on the constant, where the β^{DIS} are estimated using Equation (5.13) on different sub-groups of our sample. corresponds to the constant term. In specifications (1) and (2), we show robust standard errors in specifications (3) and (4), standard errors are clustered at the country level.

6. What drives Asymmetric Information?

We have shown that foreign forecasters make more mistakes than local forecasters, and that their relative under-performance is explained by information asymmetries. In this subsection, we use our multi-country, multi-forecaster panel to explore the determinants of these asymmetries.

6.1 Errors

We first stack observations of inflation and GDP growth errors and errors at different horizon. We then regress the log of the absolute value of the error on the Foreign dummy and other variables, without fixed effects: a dummy that is equal to 1 if GDP growth is the forecasted variable and to 0 if inflation is, a dummy that is equal to 1 if the horizon is the next year and 0 if the horizon is the current year, and a variable that goes from 1 to 12 depending of the month of year, and a dummy equal to one if the country is an emerging market. We then examine the interaction between these variables and the Foreign dummy when including all the fixed effects.

Table 6.1: Forecast Error and Information Asymmetries - Drivers I

	(1)	(2)	(3)
Coefficient	$\ln(\lvert Error^m_{ijt,t}\rvert)$	$\ln(\lvert Error^m_{ijt,t}\rvert)$	$\ln(\lvert Error^m_{ijt,t}\rvert)$
Foreign	0.11**	0.06***	0.05**
	(0.04)	(0.02)	(0.03)
GDP	0.33***		
	(0.07)		
Future	0.96***		
	(0.05)		
Emerging	0.61***		
	(0.09)		
Month-of-year	–0.08***		
	(0.01)		
Foreign × GDP			–0.04**
			(0.02)
Foreign × Future			–0.03**
			(0.01)
Foreign × Emerging			0.01
			(0.02)
Foreign × Month-of-year			0.01**
			(0.00)
N	602,122	389,295	389,295
R^2	0.18	0.70	0.70
Country × Date × Variable × Horizon FE	No	Yes	Yes
Forecaster × Date × Variable × Horizon FE	No	Yes	Yes

Note: The table shows the regression of the log absolute forecast error of current and future CPI and GDP on regressors with different fixed-effects specifications. All standard errors are clustered on the country, forecaster and date level.

The results are reported in Table 6.1. Column (1), which does not include any fixed effect, shows that forecast errors are higher for GDP growth, for the future year and for Emerging economies. Noticeably, the forecast errors diminish over time within a given year, which suggests that information flows continuously during the year. Columns (2) and (3) include variable- and horizon-specific country-time fixed effects. Foreigners have a 6% penalty on average across all variables and horizons, as column (2) shows. Column (3) shows that this penalty is significantly lower for GDP growth and for the future year. Interestingly, the penalty increases over time within a given year. This evidence shows that, somehow paradoxically, the foreign penalty is higher when there is less forecasting uncertainty. Finally, the foreign penalty does not depend on the development status of a country. This last result is consistent with the evidence in Bae et al. (2008) on the local advantage of foreign analysts.

Table 6.2 further explores the role of country-specific, forecaster-specific and time-specific variables: log of distance, quality of institutions (from the World Development Indicators), country size (log of GDP evaluated at purchasing power parity), business cycle volatility (log of the yearly GDP growth rate or inflation rate standard deviation over the whole period), financial sector dummy (equal to one if the forecaster belongs to the financial sector), stock market volatility (log of the standard deviation of the return within the month) and the VIX.[25]
Columns (1) to (5) show how these variables affect the log of the absolute value of the forecast error with different fixed-effect specifications. Better institutions are negatively associated with the size of forecast errors, even when we control for country fixed effects. This means that countries with improving institutions have also declining forecast errors. Better institutions lead to more transparency, which affects the precision of forecasts. Larger countries have also lower forecast errors. This effect is mainly driven by the cross-country dimension since it becomes insignificant when we add country fixed effects. Indeed, large countries may attract the attention of forecasters more, or they may be producing more information. Volatility plays a role too: countries with more volatile business cycles or with higher stock market volatility have higher forecast errors. Global volatility (the VIX) is also positively associated with higher forecast errors worldwide. Hence, uncertain environments are generally associated with poorer forecasting performance. The effect of distance, which is positive in some specifications, is completely absorbed by the foreign dummy in Column (5), where we include all fixed effects. The effect of geography is negligible beyond the fact of being local or foreign. Finally, forecasters from the financial sector produce better forecasts, probably because they devote more resources to forecasting.

In Column (6), these variables are interacted with the Foreign dummy. While most of these variables have a significant effect on the precision of forecasts, they do not influence the foreign penalty. Better institutions and lower business cycle or market volatility benefit symmetrically local and foreign forecasters. Similarly, financial forecasters are better at forecasting both local and foreign countries, but still perform better when forecasting locally. Only the country size has an influence: the foreign penalty is larger for larger countries. In this case, as for the evidence in Table 6.1, lower uncertainty is associated with a larger foreign penalty.[26]

[25] The data sources are the following: distance and country size (Conte and Mayer, 2022), quality of institutions (World Bank, 2022), business cycle volatility (IMF WEO), financial sector (Eikon), stock market volatility and VIX (Datastream).

[26] In the Appendix tables F1 and F2, we show that the results are unchanged when we interact the Foreign dummy

6.2 The β coefficients

In the Appendix, we conduct a similar analysis, using the estimated coefficients from our asymmetric information tests, β^{FE} and β^{Dis}. The results, which are shown in Tables F3 and F4, are broadly consistent with the evidence on the errors.[27] First, according to Table F3, β^{FE} is more negative for GDP growth and emerging economies, and less negative in later months of the year, which implies that information frictions are more prevalent for the former, and less so for the latter. Consistently, we also find that the foreign penalty is lower for GDP growth (the interaction between the GDP growth dummy and the Foreign dummy has a positive coefficient) and stronger for later months (the interaction between the month variable and the Foreign dummy has a negative coefficient), but here, this penalty is only significant for the month variable. There is still no significant extra foreign penalty for Emerging economies. Consistently, the β^{Dis} coefficient, which directly measures the foreign penalty (a more negative coefficient implies a stronger foreign penalty), only depends significantly (and negatively) on the month variable.

In Table F4, β^{FE} is significantly less negative for countries with better institutions and for larger countries, but is not more negative in more volatile countries. The foreign penalty is still stronger in large countries, but not significantly so (the interaction between country size and the Foreign dummy has a negative coefficient). However, country size does make β^{Dis} significantly more negative, which means that it matters for the foreign penalty. All in all, our results are in line with the evidence on errors, except that they are less precisely estimated.

6.3 Discussion

The asymmetry of information between local and foreign forecasters regarding aggregate variables is a robust findings. It is not affected by the development status of the economy that is being forecasted, or by the quality of institutions. This is not surprising with regards to existing evidence. Indeed, Bae et al. (2008), who examine whether local analysts are better at forecasting local firms' earnings, find that the protection of investors' rights does not influence the locals' advantage, nor does the development status of the country where the firms are located.[28]

We do find that a few variables, like country size, the nature of the variable that is being forecasted, and the forecast horizon, do influence the locals' advantage. However, interestingly,

with one variable at a time.

[27] Note that, because these coefficients are estimated at the country level and do not vary across forecasters, we cannot estimate the effect of forecaster-specific variables like distance from the forecasted country or sector.

[28] In their paper, Bae et al. (2008) show that variables that improve the functioning of the local stock market do lower the local advantage (for instance, business disclosure). But these variables are not relevant when it comes to forecast aggregate outcomes.

that local advantage is typically higher in situations with less average forecasting uncertainty. It seems that when macroeconomic information is available, it flows to local forecasters.
These results would be consistent with a better access to locally-produced information (by knowing when and where relevant information is released). Indeed, the fact that the information asymmetry is stronger for nowcasting (when forecasting the current year's GDP growth or inflation) and that it increases in the course of a year (the asymmetry is higher in December than in January) is consistent with the assumption that local forecasters are exposed to the regular releases of partial GDP growth and inflation figures and integrate this information faster. Interestingly, inflation is typically available at a higher frequency and with a shorter lag than GDP, making the access to that information an even greater advantage. This is consistent with Table 3.1 in Section 3, where we can see that the difference in updating frequency is 2% larger for inflation forecasts than for GDP growth forecasts.

7. Robustness Checks

Different Vintage Series.

To calculate forecast errors, it is standard practice in the literature to use vintage series of actual outcomes for GDP and inflation. In the main text, we focus on the vintage series from the IMF that are published in April of the subsequent year. To show that our results do not depend on this specific vintage series, we provide a robustness check using two alternative series of the actual outcome of GDP and inflation.
As a first alternative, we use the vintage series published in September of the subsequent year. For example, if a forecast for the year 2011 was submitted in October 2011, we take the vintage Series posted in September 2012 to calculate the forecast error. Similarly, if a forecast for the year 2012 was submitted in October 2011, we use the vintage Series posted in September 2013. As a second alternative, we take the data published in April two years after the forecasted date. Therefore, for the same forecasts submitted in October 2011, we use the data published in April 2013 and in April 2014.
The results are displayed in columns (1) to (2) of Table E1 in the Appendix. We report the same regression results as in Tables 4.3, 5.1, 5.2, 5.4 and 5.5 using the vintage series published in September of the subsequent year. In columns (3) to (4), we replicate the same regressions using the vintage series published in April two years after the forecasted date. Overall, the results are robust across vintage series.

Forecasters forecasting for both Local and Foreign Countries.

The rich coverage of countries and individuals forecasters in our dataset allows us to focus exclusively on forecasters that are both local but also foreign with respect to the countries they forecast for. This allows for a more direct comparison of the forecast precision conditional on the location. With this restricted subsample, we re-estimate our main results from tables 4.3, 5.1, 5.2, 5.4 and 5.5. We report the results for the mean-group estimators with the most conservative fixed effects. Columns (5) and (6) of table E1 in the appendix report the results for inflation and GDP, respectively. Overall, the findings are very similar to the baseline results.

Alternative Trimming Strategy.

In the main text, we trim observations removing forecasts that are more than 5 interquartile ranges away from the median. We re-estimate our main results with a slightly less conservative trimming method. We trim observations that are more than 6 interquartile ranges away from the median, resulting in a loss of observations for current inflation and GDP of 3 and 0.6 percent, and for future inflation and GDP of 9 and 7 percent, respectively. The results are displayed in columns (7) and (8) of table E1 in the appendix and are similar.

Alternative Definition of Foreign Forecaster.

In the main text, a foreign forecaster is defined as a forecaster that has neither its headquarter nor any subsidiary located in the country it forecasts for. This definition suggests that there is an information flow even between subsidiaries and their headquarters, regardless of the size of these subsidiaries. In this robustness check, we use an alternative definition where we define a forecaster to be foreign if its headquarter is located in another country. Compared to the 28% of foreign forecasters in the baseline results, 64% of the forecasters are defined to be foreign according to the alternative definition. We re-estimate our main results, reported in columns (9) and (10) of table E1 in the appendix. Overall, our results remain robust to this alternative definition, even though they are slightly less pronounced and more imprecisely estimated. We conclude that the location of the headquarter seems to be relevant, but that there is some information flowing from local subsidiaries to foreign headquarters.

8. Conclusion

In this paper, we provide direct evidence of asymmetric information between domestic and foreign forecasters. Using professional forecasters' expectations data in which we determine the location of each forecaster-country pair, we show that foreign forecasters update their information less frequently compared to local forecasters and produce less precise forecasts, even conditional on updating their forecasts. These results hold across several different specifications of the measure of forecast precision as well when controlling for a rich set of fixed effects.

We analyze potential sources of the differences in forecasting precision using a model of expectation formation. We rule out over-confidence and over-extrapolation, and in general, behavioral biases, as drivers of the foreigners' excess mistakes: these biases are not significantly different between local and foreign forecasters. We then identify differences in information asymmetries between foreign and local forecasters using two newly developed tests.
The results are robust to alternative trimming strategies, the use of alternative vintage series to calculate the forecast errors, when restricting the sample to forecasters that produce forecasts for both local and foreign countries, and when we use alternative definitions of "local".
Finally, we explore some determinants of the information asymmetry between local and foreign forecasters. The asymmetry is stronger at shorter horizons and when forecasting inflation. In general, the asymmetry is not weaker when forecasting is less uncertain.
Our results have implications for the modeling and calibration of international trade and finance models with heterogeneous information. First, we provide estimates of the excess errors of foreign forecasters and their relative updating frequency. Second, we prove that the source of asymmetry between local and foreign forecasters is informational. Third, we provide evidence of an elusive link between forecasting uncertainty and information asymmetry.

Table 6.2: Forecast Error and Information Asymmetries - Drivers II

	(1)	(2)	(3)	(4)	(5)	(6)
Coefficient						
Foreign	–0.05	–0.02	–0.00	0.08***	0.06***	–0.17
	(0.04)	(0.03)	(0.03)	(0.02)	(0.01)	(0.29)
ln(Distance)	0.03***	0.03**	0.03	0.01*	0.01	0.00
	(0.01)	(0.01)	(0.02)	(0.00)	(0.01)	(0.01)
Institutions	–0.02	–0.04*	–0.04**	–0.25***		
	(0.02)	(0.02)	(0.02)	(0.07)		
ln(GDP)	–0.11***	–0.10***	–0.09***	–0.46		
	(0.02)	(0.02)	(0.03)	(0.37)		
ln(sd(Fundamental))	0.55***	0.47***	0.46***			
	(0.10)	(0.10)	(0.11)			
Finance	–0.07***	–0.07***				
	(0.02)	(0.02)				
ln(sd(Return))	0.29***	0.16***	0.12**	0.06*		
	(0.05)	(0.06)	(0.05)	(0.04)		
VIX	0.01***					
	(0.00)					
Foreign $\times$ ln(Distance)						–0.01
						(0.01)
Foreign $\times$ Institutions						–0.00
						(0.01)
Foreign $\times$ ln(GDP)						0.02*
						(0.01)
Foreign $\times$ ln(sd(Fundamental))						–0.03
						(0.03)
Foreign $\times$ Finance						–0.02
						(0.02)
Foreign $\times$ ln(sd(Return))						0.02
						(0.02)
Foreign $\times$ VIX						0.00
						(0.00)
N	529,067	529,067	529,004	529,004	388,415	347,278
R^2	0.09	0.30	0.33	0.37	0.70	0.70
Date $\times$ Variable $\times$ Horizon FE	No	Yes	Yes	Yes	No	No
Forecaster $\times$ Variable $\times$ Horizon FE	No	No	Yes	Yes	No	No
Country $\times$ Variable $\times$ Horizon FE	No	No	No	Yes	No	No
Country $\times$ Date $\times$ Variable $\times$ Horizon FE	No	No	No	No	Yes	Yes
Forecaster $\times$ Date $\times$ Variable $\times$ Horizon FE	No	No	No	No	Yes	Yes

Note: The table shows the regression of the log absolute forecast error of current and future CPI and GDP on regressors with different fixed-effects specifications. All standard errors are clustered on the country, forecaster and date level.

Bibliography

Ahearne, A. G., Griever, W. L., and Warnock, F. E. (2004). Information costs and home bias: an analysis of us holdings of foreign equities. *Journal of International Economics*, 62:313–336.

Albuquerque, R., Bauer, G. H., and Schneider, M. (2007). International equity flows and returns: a quantitative equilibrium approach. *The Review of Economic Studies*, 74(1):1–30.

Albuquerque, R., Bauer, G. H., and Schneider, M. (2009). Global private information in international equity markets. *Journal of Financial Economics*, 94(1):18–46.

Allen, T. (2014). Information Frictions in Trade. *Econometrica*, 82:2041–2083.

Anderson, J. E. and van Wincoop, E. (2004). Trade costs. *Journal of Economic Literature*, 42(3):691–751.

Angeletos, G.-M., Huo, Z., and Sastry, K. A. (2021). Imperfect macroeconomic expectations: Evidence and theory. *NBER Macroeconomics Annual*, 35(1):1–86.

Bae, K.-H., Stulz, R. M., and Tan, H. (2008). Do local analysts know more? a cross-country study of performance of local analysts and foreign analysts. *Journal of Financial Economics*, 88(3):581–606.

Baik, B., Kang, J.-K., and Kim, J.-M. (2010). Local institutional investors, information asymmetries, and equity returns. *Journal of Financial Economics*, 97(1):81–106.

Baley, I., Veldkamp, L., and Waugh, M. (2020). Can global uncertainty promote international trade? *Journal of International Economics*, 126(C).

Benhima, K. and Cordonier, R. (2022). News, sentiment and capital flows. *Journal of International Economics*, 137.

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2020). Overreaction in macroeconomic expectations. *American Economic Review*, 110(9):2748–82.

Brennan, M. J. and Cao, H. (1997). International portfolio investment flows. *Journal of Finance*, 52:1851–1880.

Broner, F., Didier, T., Erce, A., and Schmukler, S. L. (2013). Gross capital flows: Dynamics and crises. *Journal of Monetary Economics*, 60(1):113–133.

Bui, H., Huo, Z., Levchenko, A. A., and Pandalai-Nayar, N. (2021). News media and international fluctuations. Technical report.

Choe, H., Kho, B.-C., and Stulz, R. M. (2005). Do domestic investors have an edge? the trading experience of foreign investors in korea. *Review of Financial Studies*, 18(3):795–829.

Clements, M. P. (1997). Evaluating the rationality of fixed-event forecasts. *Journal of Forecasting*, 16(4):225–239.

Coeurdacier, N. and Rey, H. (2013). Home bias in open economy financial macroeconomics. *Journal of Economic Literature*, 51(1):63–115.

Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–78.

Conte, M., P. C. and Mayer, T. (2022). The cepii gravity database. CEPII Working Paper 2022-05.

Cziraki, P., Mondria, J., and Wu, T. (2021). Asymmetric attention and stock returns. *Management Science*, 67(1):48–71.

Dasgupta, K. and Mondria, J. (2018). Inattentive importers. *Journal of International Economics*, 112(C):150–165.

De Marco, F., Macchiavelli, M., and Valchev, R. (2021). Beyond Home Bias: International Portfolio Holdings and Information Heterogeneity. *The Review of Financial Studies*.

Dvořák, T. (2003). Gross capital flows and asymmetric information. *Journal of International Money and Finance*, 22(6):835–864.

Eaton, J., Eslava, M., Jinkins, D., Krizan, C. J., and Tybout, J. R. (2021). A Search and Learning Model of Export Dynamics. NBER Working Papers 29100, National Bureau of Economic Research, Inc.

French, K. R. and Poterba, J. M. (1991). Investor diversification and international equity markets. *The American Economic Review*, 81(2):222–226.

Gemmi, L. and Valchev, R. (2022). Biased Surveys. Technical report.

Goldstein, N. (2021). Tracking inattention. Ssrn working paper.

Grinblatt, M. and Keloharju, M. (2001). How distance, language, and culture influence stockholdings and trades. *The Journal of Finance*, 56(3):1053–1073.

Grinblatt, M. and M., K. (2000). The investment behavior and performance of various investor types: a study of finland's unique data set. *Journal of Financial Economics*, 55:43–67.

Hamao, Y. and J., M. (2001). Living with the "enemy": an analysis of foreign investment in the japanese equity market. *Journal of International Money and Finance*, 20:715–735.

Hau, H. (2001). Location matters: an examination of trading profits. *Journal of Finance*, 56:1959–1983.

Head, K. and Mayer, T. (2013). What separates us? Sources of resistance to globalization. *Canadian Journal of Economics*, 46(4):1196–1231.

Huang, X. (2015). Thinking Outside the Borders: Investors' Underreaction to Foreign Operations Information. *The Review of Financial Studies*, 28(11):3109–3152.

Iliopulos, E., Perego, E., and Sopraseuth, T. (2021). International business cycles: Information matters. *Journal of Monetary Economics*, 123:19–34.

Kang, J.-K. and Stulz, R. M. (1997). Why is there a home bias? an analysis of foreign portfolio equity ownership in japan. *Journal of Financial Economics*, 46(1):3–28.

Kohlhas, A. and Broer, T. (2022). Forecaster (Mis-)Behavior. *Review of Economics and Statistics*, pages 1–45.

Lahiri, K. and Sheng, X. (2008). Evolution of forecast disagreement in a bayesian learning model. *Journal of Econometrics*, 144(2):325–340.

Lahiri, K. and Sheng, X. (2010). Measuring forecast uncertainty by disagreement: The missing link. *Journal of Applied Econometrics*, 25(4):514–538.

Leuz, C., Lins, K. V., and Warnock, F. E. (2009). Do foreigners invest less in poorly governed firms? *The Review of Financial Studies*, 22(8):3245–3285.

Levene, H. (1960). Robust tests for equality of variances. *Contributions to probability and statistics*, pages 278–292.

Mincer, J. A. and Zarnowitz, V. (1969). *The Evaluation of Economic Forecasts*, pages 3–46. NBER.

Mondria, J. (2010). Portfolio choice, attention allocation, and price comovement. *Journal of Economic Theory*, 145(5):1837–1864.

Mondria, J., Wu, T., and Zhang, Y. (2010). The determinants of international investment and attention allocation: Using internet search query data. *Journal of International Economics*, 82(1):85–95.

Nordhaus, W. D. (1987). Forecasting efficiency: Concepts and applications. *The Review of Economics and Statistics*, 69(4):667–674.

Pàstor, L. (2000). Portfolio selection and asset pricing models. *The Journal of Finance*, 55(1):179–223.

Portes, R. and Rey, H. (2005). The determinants of cross-border equity flows. *Journal of Financial Economics*, 65:269–296.

Portes, R., Rey, H., and Oh, Y. (2001). Information and capital flows: The determinants of transactions in financial assets. *European Economic Review*, 45(4):783–796. 15th Annual Congress of the European Economic Association.

Rich, R. and Tracy, J. (2010). The Relationships among Expected Inflation, Disagreement, and Uncertainty: Evidence from Matched Point and Density Forecasts. *The Review of Economics and Statistics*, 92(1):200–207.

Sialm, C., Sun, Z., and Zheng, L. (2020). Home Bias and Local Contagion: Evidence from Funds of Hedge Funds. *Review of Financial Studies*, 33(10):4771–4810.

Tille, C. and van Wincoop, E. (2010). International capital flows. *Journal of International Economics*, 80(2):157–175.

Tille, C. and van Wincoop, E. (2014). International capital flows under dispersed private information. *Journal of International Economics*, 93(1):31–49.

Van Nieuwerburgh, S. and Veldkamp, L. (2009). Information immobility and the home bias puzzle. *The Journal of Finance*, 64(3):1187–1215.

Van Nieuwerburgh, S. and Veldkamp, L. (2010). Information Acquisition and Under-Diversification. *The Review of Economic Studies*, 77(2):779–805.

World Bank (2022). World development indicators.
https://databank.worldbank.org/source/world-development-indicators#.

Yuan, K. (2005). Asymmetric price movements and borrowing constraints: A rational expectations equilibrium model of crises, contagion, and confusion. *The Journal of Finance*, 60(1):379–411.

Zarnowitz, V. (1983). Rational expectations and macroeconomic forecasts. Working Paper 1070, National Bureau of Economic Research.

Appendix

A. Data Appendix

Table A1: Range of Observation Periods for each Country

	Country	GDP	CPI
1	Argentina	1998m2 – 2019m12	1998m2 – 2013m12
2	Austria	2005m1 – 2019m12	2005m1 – 2019m12
3	Belgium	2005m1 – 2019m12	2005m1 – 2019m12
4	Brazil	1998m2 – 2019m12	1998m2 – 2019m12
5	Bulgaria	2007m5 – 2019m12	2007m5 – 2019m12
6	Canada	1998m1 – 2019m12	1998m1 – 2019m12
7	Chile	1998m2 – 2019m12	1998m2 – 2019m12
8	China	1998m1 – 2019m12	1998m1 – 2019m12
9	Colombia	1998m2 – 2019m12	1998m2 – 2019m12
10	Croatia	2007m5 – 2019m12	2007m5 – 2019m12
11	Czech Republic	2002m1 – 2019m12	2002m1 – 2019m12
12	Denmark	2005m1 – 2019m12	2005m1 – 2019m12
13	Estonia	2007m5 – 2019m12	2007m5 – 2019m12
14	Finland	2005m1 – 2019m12	2005m1 – 2019m12
15	France	1998m1 – 2019m12	1998m1 – 2019m12
16	Germany	1998m1 – 2019m12	1998m1 – 2019m12
17	Greece	2005m1 – 2019m12	2005m1 – 2019m12
18	Hungary	2002m1 – 2019m12	2002m1 – 2019m12
19	India	1998m1 – 2019m12	1998m1 – 2019m12
20	Indonesia	1998m1 – 2019m12	1999m1 – 2019m12
21	Ireland	2005m1 – 2019m12	2005m1 – 2019m12
22	Israel	2005m1 – 2019m12	2005m1 – 2019m12
23	Italy	1998m1 – 2019m12	1998m1 – 2019m12
24	Japan	1998m1 – 2019m12	1998m1 – 2019m12
25	Latvia	2007m5 – 2019m12	2007m5 – 2019m12
26	Lithuania	2007m5 – 2019m12	2007m5 – 2019m12
27	Malaysia	1998m1 – 2019m12	1998m1 – 2019m12
28	Mexico	1998m2 – 2019m12	1998m2 – 2019m12
29	Netherlands	1998m1 – 2019m12	1998m1 – 2019m12
30	New Zealand	1998m1 – 2019m12	1998m1 – 2019m12
31	Nigeria	2005m1 – 2019m12	2005m1 – 2019m12
32	Norway	1998m6 – 2019m12	1998m6 – 2019m12
33	Peru	1998m2 – 2019m12	1998m2 – 2019m12
34	Philippines	1998m1 – 2019m12	1998m1 – 2019m12
35	Poland	2002m1 – 2019m12	2002m1 – 2019m12

to be continued on the next page

Table A1: (continued from previous page)

	Country	GDP	CPI
36	Portugal	2005m1 – 2019m12	2005m1 – 2019m12
37	Romania	2002m1 – 2019m12	2002m9 – 2019m12
38	Russia	2002m1 – 2019m12	2002m1 – 2019m12
39	Saudi Arabia	2005m1 – 2019m12	2005m1 – 2019m12
40	Slovakia	2002m1 – 2019m12	2002m1 – 2019m12
41	Slovenia	2007m5 – 2019m12	2007m5 – 2019m12
42	South Africa	2005m1 – 2019m12	2005m1 – 2019m12
43	South Korea	1998m1 – 2019m12	1998m1 – 2019m12
44	Spain	1998m1 – 2019m12	1998m1 – 2019m12
45	Sweden	1998m1 – 2019m12	1998m1 – 2019m12
46	Switzerland	1998m6 – 2019m12	1998m6 – 2019m12
47	Thailand	1998m1 – 2019m12	1998m1 – 2019m12
48	Turkey	2002m1 – 2019m12	2003m1 – 2019m12
49	United Kingdom	1998m1 – 2019m12	1998m1 – 2019m12
50	United States	1998m1 – 2019m12	1998m1 – 2019m12
51	Venezuela	1998m2 – 2017m12	1999m6 – 2012m12

Note: The table shows the first and last observation date for GDP and CPI for which forecasts and vintages are available. The data for forecasts comes from Consensus Economics, while actual outcomes are from the International Monetary Fund World Economic Outlook (IMF WEO).

Table A2: Development Status (DS) of all Countries

	Country	DS		Country	DS		Country	DS
1	Argentina	Emerging	18	Hungary	Emerging	35	Poland	Emerging
2	Austria	Advanced	19	India	Emerging	36	Portugal	Advanced
3	Belgium	Advanced	20	Indonesia	Emerging	37	Romania	Emerging
4	Brazil	Emerging	21	Ireland	Advanced	38	Russia	Emerging
5	Bulgaria	Emerging	22	Israel	Emerging	39	Saudi Arabia	Emerging
6	Canada	Advanced	23	Italy	Advanced	40	Slovakia	Emerging
7	Chile	Emerging	24	Japan	Advanced	41	Slovenia	Emerging
8	China	Emerging	25	Latvia	Emerging	42	South Africa	Emerging
9	Colombia	Emerging	26	Lithuania	Emerging	43	South Korea	Emerging
10	Croatia	Emerging	27	Malaysia	Emerging	44	Spain	Advanced
11	Czech Republic	Emerging	28	Mexico	Emerging	45	Sweden	Advanced
12	Denmark	Advanced	29	Netherlands	Advanced	46	Switzerland	Advanced
13	Estonia	Emerging	30	New Zealand	Advanced	47	Thailand	Emerging
14	Finland	Advanced	31	Nigeria	Emerging	48	Turkey	Emerging
15	France	Advanced	32	Norway	Advanced	49	United Kingdom	Advanced
16	Germany	Advanced	33	Peru	Emerging	50	United States	Advanced
17	Greece	Advanced	34	Philippines	Emerging	51	Venezuela	Emerging

Note: The table shows the development status (DS) of all countries in the sample.

Table A3: Distribution of Forecaster-Country Pairs by Location and Scope

Location	**Scope**								
	National			**Multinational**			**Total**		
	N	Col %	Row %	N	Col %	Row %	N	Col %	Row %
Local	408	53.4	31.2	899	71.2	68.8	1,307	64.5	100.0
Foreign	356	46.6	49.5	363	28.8	50.5	719	35.5	100.0
Total	764	100.0	37.7	1,262	100.0	62.3	2,026	100.0	100.0

Note:The table shows the distribution all forecaster-country pairs by location and scope. Out of the 748 unique forecasters, we identify 2,026 forecaster-country pairs. Each of this forecaster-country pair is either foreign or local. Local forecasters have the headquarter or subsidiary in the country they forecast for, otherwise they are considered as a foreign forecaster. For the scope of a forecaster, we define multinational forecasters to have subsidiaries in different countries than their headquarter is located in. National forecasters have only subsidiaries in the same country as the headquarter.

B. Updating Appendix

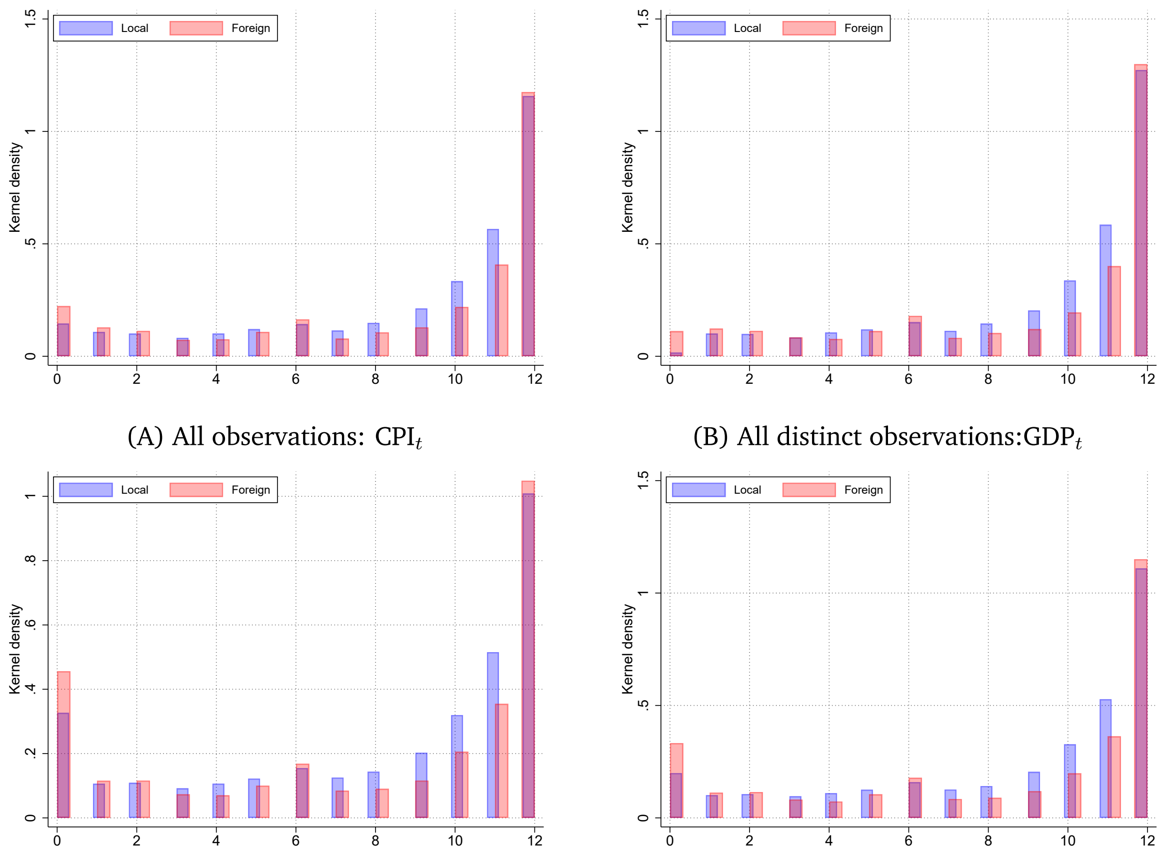

(A) All observations: CPI_t

(B) All distinct observations:GDP_t

(C) All observations: CPI_{t+1}

(D) All distinct observations:GDP_{t+1}

Figure B1: Distribution of the number of yearly updates I

Note: The figure displays the histograms of the number of updates by location. The population corresponds to all the country-forecaster-year units. We consider that any published forecast is an update.

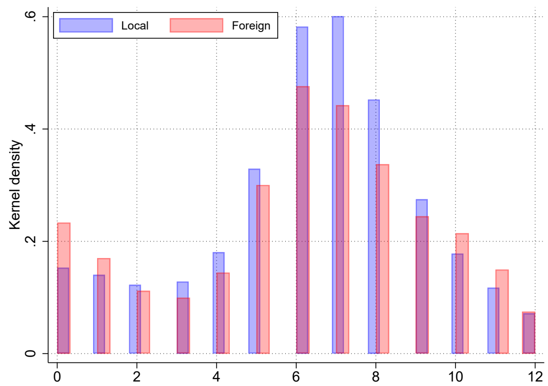

(A) All observations: CPI_t

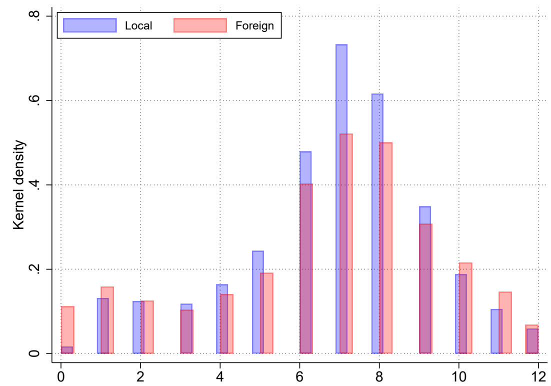

(B) All distinct observations: GDP_t

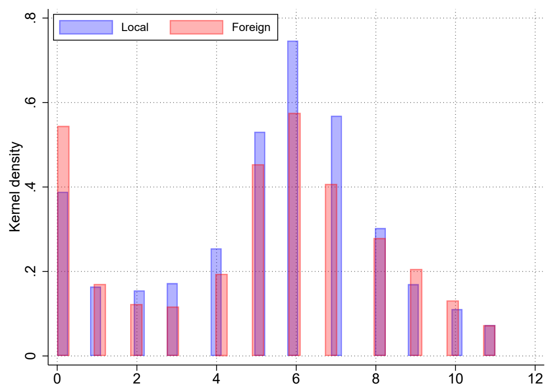

(C) All observations: CPI_{t+1}

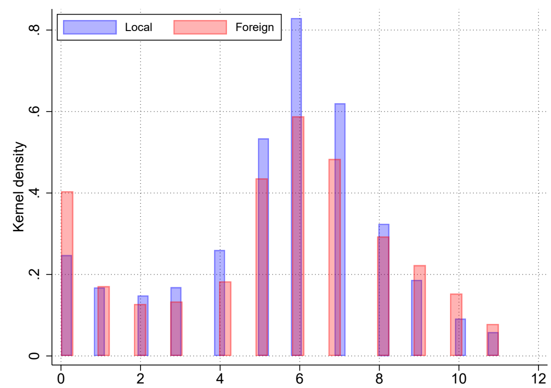

(D) All distinct observations: GDP_{t+1}

Figure B2: Distribution of the number of yearly updates II

Note: The figure displays the histograms of the number of updates by location. The population corresponds to all the country-forecaster-year units. We only consider a forecast as an update if its value differs from the last available forecast.

C. Errors Appendix

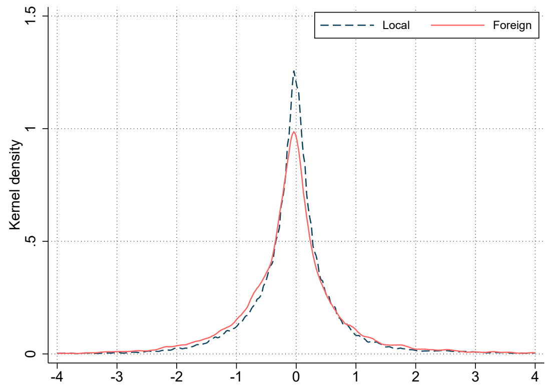

(A) CPI_t

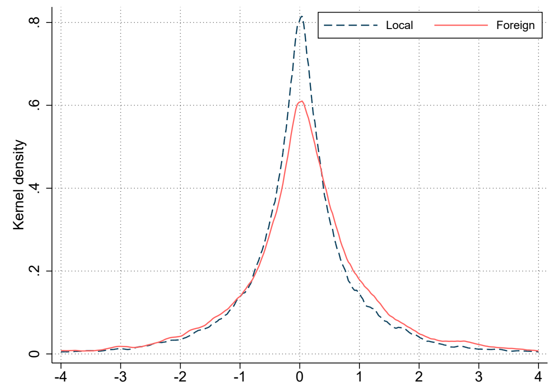

(B) GDP_t

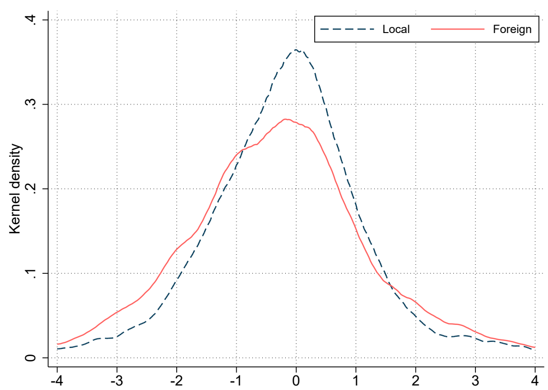

(C) CPI_{t+1}

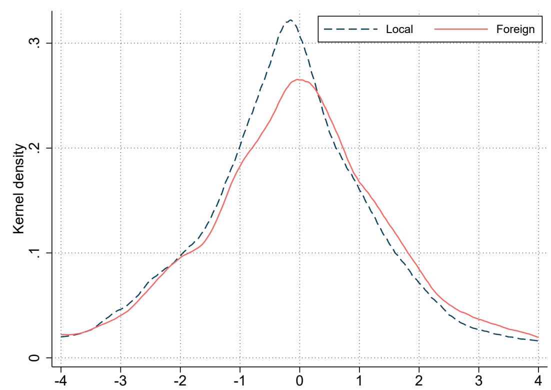

(D) GDP_{t+1}

Figure C1: Density plot of $Error^m_{ijt,t}$

Note: The figure displays the density of the forecast error $Error^m_{ijt,t}$ conditional on the location of the forecaster for current and future forecasts of inflation and GDP.

Table C1: Forecast Error conditional on Updating Forecasts

	(1)	(2)	(3)	(4)
Variable	Coefficient			
CPI_t	Foreign	0.25***	0.09***	0.09***
		(0.08)	(0.03)	(0.02)
	N	112,505	112,479	71,153
	R^2	0.01	0.14	0.65
GDP_t	Foreign	0.27***	0.11***	0.05*
		(0.09)	(0.02)	(0.02)
	N	116,080	116,054	73,067
	R^2	0.01	0.15	0.69
CPI_{t+1}	Foreign	0.26***	0.08**	0.05**
		(0.06)	(0.03)	(0.02)
	N	100,092	100,065	63,276
	R^2	0.01	0.14	0.69
GDP_{t+1}	Foreign	0.15*	0.09***	0.02
		(0.08)	(0.03)	(0.02)
	N	105,215	105,189	66,401
	R^2	0.00	0.16	0.74
	Country and Forecaster FE	No	Yes	Yes
	Country $\times$ Date	No	No	Yes
	Forecaster $\times$ Date FE	No	No	Yes

Note: The table shows the regression of the log absolute forecast error of current CPI and GDP on the location of the forecaster with different fixed-effects specifications, using only forecasts that are distinct from their last release. All standard errors are clustered on the country, forecaster and date level.

D. Biases Appendix

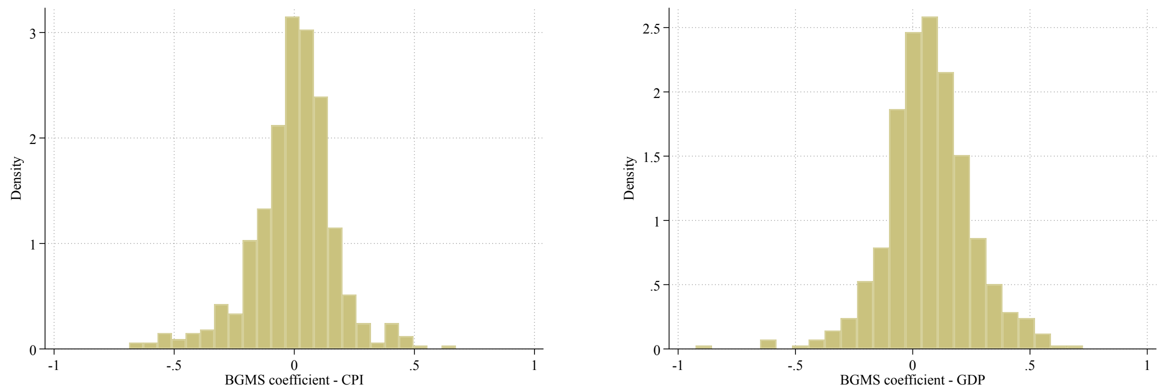

(A) CPI_t (B) GDP_t

Figure D1: Distribution of β^{BGMS} coefficients

Note: The figure displays the distribution of the β^{BGMS} coefficients estimated for each country-forecaster pair.

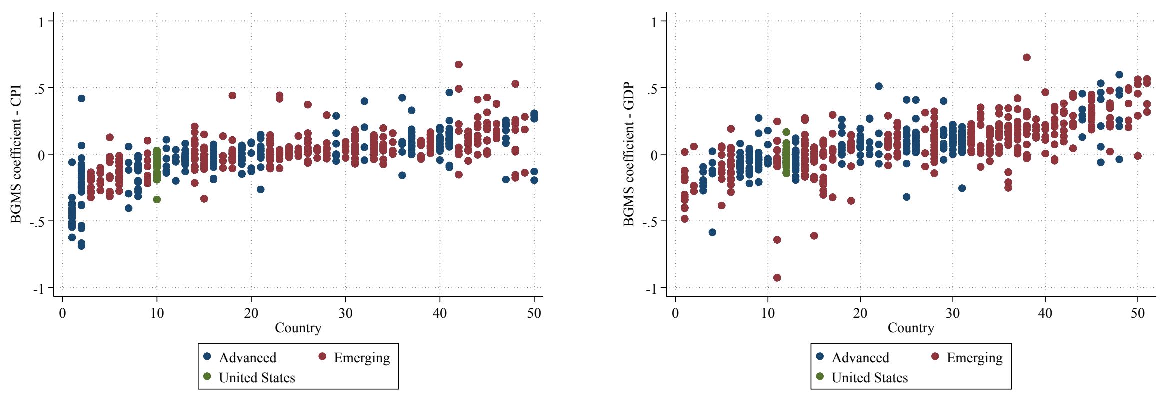

(A) CPI_t (B) GDP_t

Figure D2: β^{BGMS} coefficients by country

Note: The figure displays the β^{BGMS} coefficients estimated for each country-forecaster pair, by country, where countries are ranked by their median value.

Table D1: Behavioral Biases - Past consensus regressions

	(1)	(2)	(3)	(4)	(5)	(6)
Coefficient	CPI_t	GDP_t	CPI_t	GDP_t	CPI_t	GDP_t
Average Locals	0.01***	0.02***	0.02***	–0.01**	0.05***	0.02***
	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)	(0.00)
Foreign	0.02	–0.02	–0.02	–0.01	–0.01	–0.01
	(0.01)	(0.01)	(0.02)	(0.02)	(0.03)	(0.02)
N	102	102	390	411	6,213	6,655
R^2	0.95	0.91	0.71	0.73	0.36	0.34
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Forecaster FE	No	No	Yes	Yes	Yes	Yes
Month FE	No	No	No	No	Yes	Yes
Mean-group by cty and loc.	Yes	Yes	No	No	No	No
Mean-group by cty and for.	No	No	Yes	Yes	No	No
Mean-group by cty, for. and month	No	No	No	No	Yes	Yes

Note: The table shows the results of a regression of the $\beta^{PastConsensus}$ coefficients on the Foreign dummy, where the $\beta^{PastConsensus}$ are estimated on different sub-groups of our sample using $Error_{ijt}^m = \beta_{ij}^{PastConsensus,m} E_{jt}^{m-1}(x_{jt}) + \delta_{ij}^m + \lambda_{ijt}^m$, with $E_{jt}^m(x_{jt}) = \frac{1}{N(j)} \sum_{i \in S^j} E_{ijt}(x_{jt})$ is the average expectation across all forecasters and $E_{jt}^{m-1}(x_{jt}) = E_{jt-1}^{12}(x_{jt})$ if $m = 1$. *Average Locals* corresponds to the constant term (or average fixed effect). *Foreign* corresponds to the coefficient of the Foreign dummy. The observations are clustered at the country level in specifications (1) and (2), and at the country and forecaster levels in specifications (3) to (6).

Table D2: Behavioral Biases - Vintage regressions

	(1)	(2)	(3)	(4)	(5)	(6)
Coefficient	CPI_t	GDP_t	CPI_t	GDP_t	CPI_t	GDP_t
Average Locals	0.01***	–0.09***	0.02***	–0.09***	0.03***	–0.08***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Foreign	–0.00	–0.01*	–0.01	–0.01	–0.01	–0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
N	102	102	425	448	6,662	7,131
R^2	0.95	0.95	0.72	0.74	0.45	0.49
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Forecaster FE	No	No	Yes	Yes	Yes	Yes
Month FE	No	No	No	No	Yes	Yes
Mean-group by cty and loc.	Yes	Yes	No	No	No	No
Mean-group by country and for.	No	No	Yes	Yes	No	No
Mean-group by cty, for. and month	No	No	No	No	Yes	Yes

Note:The table shows the results of a regression of the $\beta^{LastVintage}$ coefficients on the Foreign dummy, where $\beta^{LastVintage}$ are estimated on different sub-groups of our sample using $Error_{ijt}^m = \beta_{ij}^{LastVintage,m} x_{jt-1} + \delta_{ij}^m + \lambda_{ijt}^m$. *Average Locals* corresponds to the constant term (or average fixed effect). *Foreign* corresponds to the coefficient of the Foreign dummy. The observations are clustered at the country level in specifications (1) and (2), and at the country and forecaster levels in specifications (3) to (6).

2. Do Local Forecasters Have Better Information

E. Robustness Appendix

Table E1: Robustness Checks - Summary Results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
		Vintages September		**Vintages April**		**Local and Foreign**		**Trimming**		**Headquarter**	
$\ln(\lvert Error^{m}_{ijt,t}\rvert)$	Foreign	0.09***	0.05**	0.08***	0.05**	0.08***	0.06**	0.10***	0.06**	0.06	0.08**
		(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.04)	(0.04)
	N	99,228	103,722	91,844	95,826	70,885	74,550	99,791	104,645	99,228	103,866
BGMS Regression	Average Locals	0.04***	0.10***	0.03***	0.11***	0.03***	0.08***	0.03***	0.08***	0.04**	0.11***
		(0.01)	(0.01)	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.01)
	Foreign	0.00	0.05	0.00	0.02	0.02	0.01	0.00	0.04	–0.02	–0.02
		(0.03)	(0.03)	(0.03)	(0.04)	(0.03)	(0.04)	(0.03)	(0.03)	(0.03)	(0.02)
	N	4,979	5,357	4,457	4,772	1,875	2,056	5,012	5,431	4,979	5,373
Over-Extrapolation	Average Locals	0.39***	0.37***	0.39***	0.36***	0.39***	0.35***	0.39***	0.35***	0.38***	0.37***
		(0.00)	(0.01)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)
	Foreign	0.03	0.03	0.03	0.05**	0.03	0.04	0.03	0.03	0.02	–0.02
		(0.02)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
	N	3,808	4,006	3,359	3,562	6,097	6,535	6,137	6,584	6,097	6,535
Information Asymmetries	Average Locals	–0.29***	–0.32***	–0.29***	–0.32***	–0.29***	–0.32***	–0.29***	–0.32***	–0.29***	–0.34***
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)	(0.02)
	Foreign	–0.05***	–0.02	–0.04***	–0.03	–0.05***	–0.02	–0.05***	–0.02	–0.02	0.03
		(0.01)	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.03)
	N	1,196	1,207	1,187	1,198	1,196	1,207	1,200	1,207	1,028	1,038
Disagreement	Disagreement	–0.09***	–0.07**	–0.10***	–0.07***	–0.06**	–0.06**	–0.10***	–0.07**	–0.04	–0.07*
		(0.03)	(0.03)	(0.03)	(0.02)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.04)
	N	611	612	611	612	592	597	611	612	548	550

Note:This table shows the results of several robustness checks. In columns (2) and (3), we use alternative vintage series that were published in September of the subsequent year of the forecast. In columns (4) and (5), we use vintage series that were published in April two years after the forecast. In columns (6) and (7), we only use forecasters that forecast for both countries where they are foreign and local. In columns (8) and (9), we use a less conservative trimming strategy to remove outliers for inflation and GDP forecasts. In columns (10) and (11), we only use the headquarter of the forecaster to identify whether the forecaster is local or foreign. For each of these robustness checks, we reproduce the results of tables 4.3, 5.1, 5.2, 5.4 and 5.5.

F. Determinants Appendix

Table F1: Forecast Error and Information Asymmetries - Drivers I, Separate Regressions

	(1)	(2)	(3)	(4)
Coefficient	$\ln(\|Error^m_{ijt,t}\|)$	$\ln(\|Error^m_{ijt,t}\|)$	$\ln(\|Error^m_{ijt,t}\|)$	$\ln(\|Error^m_{ijt,t}\|)$
Foreign	0.08***	0.07***	0.05**	0.02
	(0.02)	(0.02)	(0.02)	(0.02)
Foreign × GDP	–0.04**			
	(0.02)			
Foreign × Future		–0.03**		
		(0.01)		
Foreign × Emerging			0.01	
			(0.02)	
Foreign × Month-of-year				0.01**
				(0.00)
N	389,295	389,295	389,295	389,295
R^2	0.70	0.70	0.70	0.70
Country × Date × Variable × Horizon FE	Yes	Yes	Yes	Yes
Forecaster × Date × Variable × Horizon FE	Yes	Yes	Yes	Yes

Note: The table shows the regression of the log absolute forecast error of current and future CPI and GDP on regressors with different fixed-effects specifications. All standard errors are clustered on the country, forecaster and date level.

Table F2: Forecast Error and Information Asymmetries - Drivers II, Separate Regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Coefficient							
Foreign	0.09	0.06***	–0.26	0.07***	0.07***	0.05***	0.05*
	(0.10)	(0.02)	(0.17)	(0.02)	(0.02)	(0.02)	(0.02)
ln(Distance)	0.01						
	(0.01)						
Foreign × ln(Distance)	–0.00						
	(0.01)						
Foreign × Institutions		–0.00					
		(0.00)					
Foreign × ln(GDP)			0.02*				
			(0.01)				
Foreign × ln(sd(Fundamental))				–0.01			
				(0.02)			
Foreign × Finance					–0.01		
					(0.03)		
Foreign × ln(sd(Return))						0.02	
						(0.01)	
Foreign × VIX							0.00
							(0.00)
N	388,415	375,405	379,087	389,295	389,295	364,155	389,295
R^2	0.70	0.70	0.70	0.70	0.70	0.70	0.70
Cty × Date × Var. × Hor. FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Inst. × Date × Var. × Hor. FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note:The table shows the regression of the log absolute forecast error of current and future CPI and GDP on regressors with different fixed-effects specifications. All standard errors are clustered on the country, forecaster and date level.

Table F3: Determinants of information asymmetries - I

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Coefficient	β^{FE}	β^{FE}	β^{FE}	β^{Disag}	β^{Disag}	β^{Disag}	β^{Disag}
Foreign	–0.07***	–0.06**	0.00				
	(0.01)	(0.03)	(0.02)				
Month of year	0.03***	0.03***		–0.01*	–0.01**		
	(0.00)	(0.00)		(0.01)	(0.01)		
GDP	–0.03*	–0.02		0.02		0.02	
	(0.02)	(0.02)		(0.03)		(0.03)	
Emerging	–0.03	–0.04**		0.04			0.04
	(0.02)	(0.02)		(0.05)			(0.05)
Foreign × Month of year		–0.01***	–0.01***				
		(0.00)	(0.00)				
Foreign × GDP		–0.01	0.03				
		(0.02)	(0.02)				
Foreign × Emerging		0.06***	–0.00				
		(0.02)	(0.02)				
N	2,403	2,403	2,403	1,223	1,223	1,222	1,223
R^2	0.41	0.42	0.63	0.01	0.23	0.57	0.02
Country × Variable FE	No	No	Yes	No	Yes	No	No
Month-of-year × Variable FE	No	No	Yes	No	No	No	Yes
Country × Month-of-year FE	No	No	No	No	No	Yes	No

Note: The table shows the regression of β^{FE} and β^{Disag} on regressors with different fixed-effects specifications. All standard errors are clustered on the country level.

Table F4: Determinants of information asymmetries - II

	(1)	(2)	(3)	(4)
Coefficient	β^{FE}	β^{FE}	β^{FE}	β^{Disag}
Foreign	–0.04***	0.30	–0.17	
	(0.01)	(0.21)	(0.20)	
ln(sd(Fundamental))	0.01	0.00		0.04
	(0.02)	(0.02)		(0.05)
Institutions	0.01**	0.01**		–0.00
	(0.00)	(0.00)		(0.01)
ln(GDP)	0.03***	0.04***		–0.04**
	(0.01)	(0.01)		(0.02)
Foreign × ln(sd(Fundamental))		0.01	0.02	
		(0.03)	(0.03)	
Foreign × Institutions		–0.01	0.00	
		(0.01)	(0.01)	
Foreign × ln(GDP)		–0.02	0.01	
		(0.01)	(0.01)	
N	2,403	2,403	2,403	1,223
R^2	0.47	0.47	0.63	0.04
Month-of-year × Variable FE	Yes	Yes	Yes	Yes
Country × Variable FE	No	No	Yes	No

Note:The table shows the regression of β^{FE} and β^{Disag} on regressors with different fixed-effects specifications. Fundamental is either CPI or GDP. All standard errors are clustered on the country level.

G. Proofs

G.1 Proof of Proposition 1

The model can be written as follows:

$$\begin{aligned} x_{jt} &= \rho_j x_{jt-1} + \epsilon_{jt} \\ s^m_{ijt} &= x_{jt} + v^m_{ijt} \end{aligned} \tag{G.1}$$

with $v^m_{ijt} \sim N(0, (\kappa^m_j + \tau^m_{ij})^{-1/2})$. We denote $\lambda^m_{ij} = \kappa^m_j + \tau^m_{ij}$
Denote the one step-ahead forecast error for the forecast in the Kalman filter with $\Phi^m_{ij} = V(Error^m_{ijt,t-1}) = V[x_{jt} - E^m_{ijt-1}(x_{jt})]$. We can find Φ^m_{ij} from the Riccati equation

$$\Phi^m_{ij} = \rho_j^2[\Phi^m_{ij} - \Phi^m_{ij}(\Phi^m_{ij} + (\lambda^m_{ij})^{-1})^{-1}\Phi^m_{ij}] + \gamma_j^{-1}.$$

Denote the gain of the Kalman filter with

$$G^m_{ij} = \Phi^m_{ij}(\Phi^m_{ij} + (\lambda^m_{ij})^{-1})^{-1}.$$

Substituting in the Riccati equation, we obtain

$$\Phi^m_{ij} = \rho_j^2(1 - G^m_{ij})\Phi^m_{ij} + \gamma_j^{-1},$$

hence the first result.
Now denote the nowcast error in the Kalman filter with $\Omega^m_{ij} = V(Error^m_{ijt,t}) = V[x_{jt} - E^m_{ijt}(x_{jt})]$ We can use recursions of the Kalman filter to relate Ω^m_{ij} and Φ^m_{ij}:

$$\Omega^m_{ij} = \Phi^m_{ij} - G^m_{ij}(\Phi^m_{ij} + (\lambda^m_{ij})^{-1})G^{m\prime}_{ij}$$

Replacing $G^{m\prime}_{jk}$, we obtain

$$\begin{aligned} \Omega^m_{ij} &= \Phi^m_{ij} - G^m_{ij}(\Phi^m_{ij} + (\lambda^m_{ij})^{-1})[\Phi^m_{ij}(\Phi^m_{ij} + (\lambda^m_{ij})^{-1})^{-1}]' \\ &= \Phi^m_{ij} - G^m_{ij}\Phi^m_{ij} \\ &= (1 - G^m_{ij})\Phi^m_{ij} \end{aligned}$$

Hence the second result.
Note that solving the Riccati equation gives us an expression for Φ^m_{ij}:

$$\Phi^m_{ij} = \frac{1}{2}\left(\gamma_j^{-1} - (1-\rho_j^2)(\lambda^m_{ij})^{-1} + \sqrt{(\gamma_j^{-1} - (1-\rho_j^2)(\lambda^m_{ij})^{-1})^2 + 4\gamma_j^{-1}}\right) \tag{G.2}$$

and for G_{ij}:

$$G_{ij}^m = 1 - \frac{2}{\lambda_{ij}^m/\gamma_j + 1 + \rho_j^2 + \sqrt{(\lambda_{ij}^m/\gamma_j - (1-\rho_j^2))^2 + 4\lambda_{ij}^m/\gamma_j}}$$

which is an increasing function of λ_{ij}^m and hence of τ_{ij}^m.

G.2 Proof of Proposition 2

Notice that $E_{ijt}^m(x_{jt})$ can be rewritten in its moving-average form as follows:

$$E_{ijt}^m(x_{jt}) = \frac{G_{ij}^m}{1-(1-G_{ij}^m)\hat{\rho}_{ij}L} s_{ijt}^m \tag{G.3}$$

Forecast revision can then be written as

$$\begin{aligned} Revision_{ijt}^m &= E_{ijt}^m(x_{jt}) - E_{ijt-1}^m(x_{jt}) \\ &= E_{ijt}^m(x_{jt}) - \hat{\rho}_{ij} E_{ijt-1}^m(x_{jt-1}) \\ &= \tfrac{G_{ij}^m[1-\hat{\rho}_{ij}L]}{1-(1-G_{ij}^m)\hat{\rho}_{ij}L} s_{ijt}^m \\ &= \tfrac{G_{ij}^m[1-\hat{\rho}_{ij}L]}{1-(1-G_{ij}^m)\hat{\rho}_{ij}L} (x_{jt} + v_{ijt}^m) \end{aligned} \tag{G.4}$$

and the error as

$$\begin{aligned} Error_{ijt,t}^m &= x_{jt} - E_{ijt}^m(x_{jt}) \\ &= x_{jt} - \tfrac{G_{ij}^m}{1-(1-G_{ij}^m)\hat{\rho}_{ij}L} s_{ijt}^m \\ &= \left(1 - \tfrac{G_{ij}^m}{1-(1-G_{ij}^m)\rho_{ij}L}\right) x_{jt} - \tfrac{G_{ij}^m}{1-(1-G_{ij}^m)\hat{\rho}_{ij}L} v_{ijt}^m \end{aligned} \tag{G.5}$$

with $v_{ijt}^m = h_{ij}^m(\kappa_j^m)^{-1/2}u_{jt}^m + (1-h_{ij}^m)(\tau_{ij}^m)^{-1/2}e_{ijt}^m$ is the total noise.
The estimated OLS coefficient β_{ij}^{BGMSm} is given by

$$\beta_{ij}^{BGMSm} = \frac{Cov\left(Error_{ijt}^m, Revision_{ijt}^m\right)}{V\left(Revision_{ijt}^m\right)}$$

We define $\tilde{E}rror_{ijt,t}^m$ as the error if the persistence and private signal precisions were the ones corresponding to the forecaster's beliefs:

$$\tilde{E}rror_{ijt,t}^m = \left(1 - \frac{G_{ij}^m}{1-(1-G_{ij}^m)\hat{\rho}_{ij}L}\right)\tilde{x}_{ijt} - \frac{G_{ij}^m}{1-(1-G_{ij}^m)\hat{\rho}_{ij}L}\tilde{v}_{ijt}^m \tag{G.6}$$

with $\tilde{x}_{ijt} = \epsilon_{jt}/(1-\hat{\rho}_{ij}L)$ and $\tilde{v}_{ijt}^m = h_{ij}^m(\kappa_j^m)^{-1/2}u_{jt}^m + (1-h_{ij}^m)(\hat{\tau}_{ij}^m)^{-1/2}e_{ijt}^m$. We define

$\tilde{Revision}^m_{ijt,t}$ similarly:

$$\tilde{Revision}^m_{ijt} = \frac{G^m_{ij}[1-\hat{\rho}_{ij}L]}{1-(1-G^m_{ij})\hat{\rho}_{ij}L}(\tilde{x}_{ijt}+\tilde{v}^m_{ijt})$$

We then use the fact that the forecaster's expectations are rational conditional on their beliefs: $Cov(\tilde{Error}^m_{ijt,t}, \tilde{Revision}^m_{ijt}) = 0$ to determine the covariance of the actual errors and revisions:

$$\begin{aligned}
Cov\left(Error^m_{ijt}, Revision^m_{ijt}\right) = \;& Cov\left(Error^m_{ijt} - \tilde{Error}^m_{ijt}, \tilde{Revision}^m_{ijt}\right) \\
&+Cov\left(\tilde{Error}^m_{ijt}, Revision^m_{ijt} - \tilde{Revision}^m_{ijt}\right) \\
&+Cov\left(Error^m_{ijt} - \tilde{Error}^m_{ijt}, Revision^m_{ijt} - \tilde{Revision}^m_{ijt}\right) \\
= \;& Cov\left(\left(1-\tfrac{G^m_{ij}}{1-(1-G^m_{ij})\hat{\rho}_{ij}L}\right)(x_{jt}-\tilde{x}_{ijt}), \tfrac{G^m_{ij}(1-\hat{\rho}_{ij}L)}{1-(1-G^m_{ij})\hat{\rho}_{ij}L}\tilde{x}_{ijt}\right) \\
&+Cov\left(\left(1-\tfrac{G^m_{ij}}{1-(1-G^m_{ij})\hat{\rho}_{ij}L}\right)\tilde{x}_{ijt}, \tfrac{G^m_{ij}(1-\hat{\rho}_{ij}L)}{1-(1-G^m_{ij})\hat{\rho}_{ij}L}(x_{jt}-\tilde{x}_{ijt})\right) \\
&+Cov\left(\left(1-\tfrac{G^m_{ij}}{1-(1-G^m_{ij})\hat{\rho}_{ij}L}\right)(x_{jt}-\tilde{x}_{ijt}), \tfrac{G^m_{ij}(1-\hat{\rho}_{ij}L)}{1-(1-G^m_{ij})\hat{\rho}_{ij}L}(x_{jt}-\tilde{x}_{ijt})\right) \\
&-Cov\left(\tfrac{G^m_{ij}}{1-(1-G^m_{ij})\hat{\rho}_{ij}L}\tilde{v}^m_{ijt}, \tfrac{G^m_{ij}(1-\hat{\rho}_{ij}L)}{1-(1-G^m_{ij})\hat{\rho}_{ij}L}(v^m_{ijt}-\tilde{v}^m_{ijt})\right) \\
&-Cov\left(\tfrac{G^m_{ij}}{1-(1-G^m_{ij})\hat{\rho}_{ij}L}(v^m_{ijt}-\tilde{v}^m_{ijt}), \tfrac{G^m_{ij}(1-\hat{\rho}_{ij}L)}{1-(1-G^m_{ij})\hat{\rho}_{ij}L}\tilde{v}^m_{ijt}\right) \\
&-Cov\left(\tfrac{G^m_{ij}}{1-(1-G^m_{ij})\hat{\rho}_{ij}L}(v^m_{ijt}-\tilde{v}^m_{ijt}), \tfrac{G^m_{ij}(1-\hat{\rho}_{ij}L)}{1-(1-G^m_{ij})\hat{\rho}_{ij}L}(v^m_{ijt}-\tilde{v}^m_{ijt})\right) \\
= \;& -(\hat{\rho}_{ij}-\rho_j)G^m_{ij}(1-G^m_{ij})\tfrac{2\hat{\rho}_{ij}(1-G^m_{ij})(1-\rho_j^2)-(\hat{\rho}_{ij}-\rho_j)[1+\rho_j\hat{\rho}_{ij}(1-G^m_{ij})]}{[1-\rho_j\hat{\rho}_{ij}(1-G^m_{ij})][1-\rho_j^2][1-\hat{\rho}_{ij}^2(1-G^m_{ij})^2]} \\
&-[(\tau^m_{ij})^{-1}-(\hat{\tau}^m_{ij})^{-1}]((1-h^m_{ij})G^m_{ij})^2\tfrac{1-\hat{\rho}^2_{ij}(1-G^m_{ij})}{1-\hat{\rho}^2_{ij}(1-G^m_{ij})^2}
\end{aligned}$$

We used

$$\begin{aligned}
\tilde{Error}^m_{ijt} \;& = (1-G^m_{ij})\textstyle\sum_{s=0}^{+\infty}(1-G^m_{ij})^s\hat{\rho}^s_{ij}L^s\epsilon_{jt} \\
& -G^m_{ij}\textstyle\sum_{s=0}^{+\infty}(1-G^m_{ij})^s\hat{\rho}^s_{ij}L^s h^m_{ij}(\hat{\tau}^m_{ij})^{-1/2}e^m_{ijt} \\
\tilde{Revision}^m_{ijt} \;& = G^m_{ij}\textstyle\sum_{s=0}^{+\infty}(1-G^m_{ij})^s\hat{\rho}^s_{ij}L^s\epsilon_{jt} \\
& -G^m_{ij}\left(1-\tfrac{G^m_{ij}}{1-G^m_{ij}}\textstyle\sum_{s=1}^{+\infty}(1-G^m_{ij})^s\hat{\rho}^s_{ij}L^s\right)(1-h^m_{ij})(\hat{\tau}^m_{ij})^{-1/2}e^m_{ijt} \\
Error^m_{ijt}-\tilde{Error}^m_{ijt} \;& = \frac{-\left(\frac{\hat{\rho}_{ij}}{\rho_j}-1\right)(1-G^m_{ij})}{1-(1-G^m_{ij})\frac{\hat{\rho}_{ij}}{\rho_j}}\left(\textstyle\sum_{s=0}^{+\infty}\rho^s_{ij}L^s-\sum_{s=0}^{+\infty}(1-G^m_{ij})^s\hat{\rho}^s_{ij}L^s\right)\epsilon_{jt} \\
& -G^m_{ij}\textstyle\sum_{s=0}^{+\infty}(1-G^m_{ij})^s\hat{\rho}^s_{ij}L^s h^m_{ij}[(\tau^m_{ij})^{-1/2}-(\hat{\tau}^m_{ij})^{-1/2}]e^m_{ijt} \\
Revision^m_{ijt}-\tilde{Revision}^m_{ijt} \;& = \frac{-\left(\frac{\hat{\rho}_{ij}}{\rho_j}-1\right)G^m_{ij}}{1-(1-G^m_{ij})\frac{\hat{\rho}_{ij}}{\rho_j}}\left(\textstyle\sum_{s=0}^{+\infty}\rho^s_{ij}L^s-\sum_{s=0}^{+\infty}(1-G^m_{ij})^s\hat{\rho}^s_{ij}L^s\right)\epsilon_{jt} \\
& -G^m_{ij}\left(1-\tfrac{G^m_{ij}}{1-G^m_{ij}}\textstyle\sum_{s=1}^{+\infty}(1-G^m_{ij})^s\hat{\rho}^s_{ij}L^s\right)(1-h^m_{ij})[(\tau^m_{ij})^{-1/2}-(\hat{\tau}^m_{ij})^{-1/2}]e^m_{ijt}
\end{aligned}$$

We thus have

$$\beta_{1ij}^{m}=\frac{G_{ij}^{m}(1-G_{ij}^{m})\frac{2\hat{\rho}_{ij}(1-G_{ij}^{m})(1-\rho_{j}^{2})-(\hat{\rho}_{ij}-\rho_{j})[1+\rho_{j}\hat{\rho}_{ij}(1-G_{ij}^{m})]}{[1-\rho_{j}\hat{\rho}_{ij}(1-G_{ij}^{m})][1-\rho_{j}^{2}][1-\hat{\rho}_{ij}^{2}(1-G_{ij}^{m})^{2}]}}{V\left(Revision_{ijt}^{m}\right)}$$

and

$$\beta_{2ij}^{m}=\frac{(h_{ij}^{m}G_{ij}^{m})^{2}\frac{1-\hat{\rho}_{ij}^{2}(1-G_{ij}^{m})}{1-\hat{\rho}_{ij}^{2}(1-G_{ij}^{m})^{2}}}{V\left(Revision_{ijt}^{m}\right)}$$

with

$$\begin{aligned}V(Revision_{ijt}^{m})= & \frac{(G_{ij}^{m})^{2}}{1-\frac{\hat{\rho}_{ij}}{\rho_{j}}(1-G_{ij}^{m})}\left(\frac{G_{ij}^{m}\frac{\hat{\rho}_{ij}}{\rho_{j}}[1-\hat{\rho}_{ij}^{2}(1-G_{ij})]}{[1-\rho_{j}\hat{\rho}_{ij}(1-G_{ij})][1-\hat{\rho}_{ij}^{2}(1-G_{ij})^{2}]}-(\hat{\rho}_{ij}-\rho_{j})\frac{1-\rho_{j}\hat{\rho}_{ij}}{[1-\rho_{j}\hat{\rho}_{ij}(1-G_{ij})](1-\rho_{j}^{2})}\right)\\ & +(G_{ij}^{m})^{2}\left(1+\left(\frac{G_{ij}^{m}}{1-G_{ij}^{m}}\right)^{2}\frac{\hat{\rho}_{ij}^{2}(1-G_{ij}^{m})^{2}}{1-\hat{\rho}_{ij}^{2}(1-G_{ij}^{m})^{2}}\right)[(h_{ij}^{m})^{2}\kappa_{j}^{-1}+(1-h_{ij}^{m})^{2}\tau_{ij}^{-1}]\end{aligned}$$

Here we used

$$\begin{aligned}Revision_{ijt}^{m}= & \frac{G_{ij}^{m}}{1-\frac{\hat{\rho}_{ij}}{\rho_{j}}(1-G_{ij}^{m})}\left(\frac{\hat{\rho}_{ij}}{\rho_{j}}\sum_{s=0}^{+\infty}(1-G_{ij}^{m})^{s}\hat{\rho}_{ij}^{s}L^{s}-\left(\frac{\hat{\rho}_{ij}}{\rho_{j}}-1\right)\sum_{s=0}^{+\infty}\rho_{j}^{s}L^{s}\right)\epsilon_{jt}\\ & +G_{ij}^{m}\left(1-\frac{G_{ij}^{m}}{1-G_{ij}^{m}}\sum_{s=1}^{+\infty}(1-G_{ij}^{m})^{s}\hat{\rho}_{ij}^{s}L^{s}\right)v_{ijt}^{m}\end{aligned}$$

G.3 Proof of Corollary 1

We simply note here that β_{1ij}^{m} and β_{2ij}^{m}, evaluated at $(\hat{\tau}_{ij}^{m})^{-1}=(\tau_{ij}^{m})^{-1}=(\tau_{j}^{m})^{-1}$ and $\hat{\rho}_{ij}=\rho_{j}$, are both strictly positive, while $\hat{\rho}_{ij}-\rho_{j}$ and $(\tau_{ij}^{m})^{-1}-(\hat{\tau}_{ij}^{m})^{-1}$ are both equal to zero for $\hat{\tau}_{ij}^{m}=\tau_{ij}^{m}=\tau_{j}^{m}$ and $\hat{\rho}_{ij}=\rho_{j}$.
More specifically, note that β_{1ij}^{m} and β_{2ij}^{m} are functions of the parameters, so we denote $\beta_{1ij}^{m}=g_{1}\left((\hat{\tau}_{ij}^{m})^{-1},(\tau_{ij}^{m})^{-1},\hat{\rho}_{ij},\rho_{j}\right)$ and $\beta_{2ij}^{m}=g_{2}\left((\hat{\tau}_{ij}^{m})^{-1},(\tau_{ij}^{m})^{-1},\hat{\rho}_{ij},\rho_{j}\right)$. The first-order Taylor expansion for β_{ij}^{BGMSm} around $(\hat{\tau}_{ij}^{m})^{-1}=(\tau_{ij}^{m})^{-1}=(\tau_{j}^{m})^{-1}$ and $\hat{\rho}_{ij}=\rho_{j}$ is

$$\beta_{ij}^{BGMSm}\simeq-(\hat{\rho}_{ij}-\rho_{j})g_{1}\left((\tau_{j}^{m})^{-1},(\tau_{j}^{m})^{-1},\rho_{j},\rho_{j}\right)-[(\tau_{ij}^{m})^{-1}-(\hat{\tau}_{ij}^{m})^{-1}]g_{2}\left((\tau_{j}^{m})^{-1},(\tau_{j}^{m})^{-1},\rho_{j},\rho_{j}\right)$$

We can show that $\hat{\beta}_{1j}^{m}=g_{1}\left((\tau_{j}^{m})^{-1},(\tau_{j}^{m})^{-1},\rho_{j},\rho_{j}\right)$ and $\hat{\beta}_{2j}^{m}=g_{2}\left((\tau_{j}^{m})^{-1},(\tau_{j}^{m})^{-1},\rho_{j},\rho_{j}\right)$ are both strictly positive, hence the result.

G.4 Proof of Proposition 3

The estimated OLS coefficient β_{jk}^{CGm}, for $k=l,f$, $m=1,..,12$ and $j=1,..,J$, is given by

$$\beta_{jk}^{CGm}=\frac{Cov\left(Error_{jkt}^{m},Revision_{jkt}^{m}\right)}{V\left(Revision_{jkt}^{m}\right)}\tag{G.7}$$

And we can write:

$$\begin{aligned} Cov\left(Error^m_{jkt}, Revision^m_{jkt}\right) = \quad & Cov\left(\tilde{E}rror^m_{jkt}, \tilde{R}evision^m_{jkt}\right) \\ & + Cov\left(Error^m_{jkt} - \tilde{E}rror^m_{jkt}, \tilde{R}evision^m_{jkt}\right) \\ & + Cov\left(\tilde{E}rror^m_{jkt}, Revision^m_{jkt} - \tilde{R}evision^m_{jkt}\right) \\ & + Cov\left(Error^m_{jkt} - \tilde{E}rror^m_{jkt}, Revision^m_{jkt} - \tilde{R}evision^m_{jkt}\right) \end{aligned}$$

with $\tilde{E}rror^m_{jkt} = \frac{1}{N^k(j)} \sum_{i \in S^k(j)} \tilde{E}rror^m_{ijt}$ where $\tilde{E}rror^m_{ijt}$ is defined in (G.6).
We have

$$\begin{aligned} Cov\left(\tilde{E}rror^m_{jkt}, \tilde{R}evision^m_{jkt}\right) & = Cov\left(\left(1 - \frac{G^m_{jk}}{1-(1-G^m_{jk})\hat{\rho}_{jk}L}\right)\frac{1}{1-\hat{\rho}_{jk}L}\epsilon_{jt}, \frac{G^m_{jk}}{1-(1-G^m_{jk})\hat{\rho}_{jk}L}\epsilon_{jt}\right) \\ & + Cov\left(-\frac{G^m_{jk}}{1-(1-G^m_{jk})\hat{\rho}_{jk}L} h^m_{jk}(\kappa^m_j)^{-1/2}u^m_{jt}, \frac{G^m_{jk}[1-\hat{\rho}_{jk}L]}{1-(1-G^m_{jk})\hat{\rho}_{jk}L} h^m_{jk}(\kappa^m_j)^{-1/2}u^m_{jt}\right) \\ & = \frac{G^m_{jk}(1-G^m_{jk})}{1-\hat{\rho}^2_{jk}(1-G^m_{jk})^2}\gamma^{-1} - (G^m_{jk})^2\left(1 - \frac{G^m_{jk}}{1-G^m_{jk}}\frac{\hat{\rho}^2_{jk}(1-G^m_{jk})^2}{1-\hat{\rho}^2_{jk}(1-G^m_{jk})^2}\right)(h^m_{jk})^2(\kappa^m_j)^{-1} \end{aligned}$$

Here we used

$$\begin{aligned} \tilde{E}rror^m_{jkt} & = \left(1 - \frac{G^m_{jk}}{1-(1-G^m_{jk})\hat{\rho}_{jk}L}\right)\frac{1}{1-\hat{\rho}_{jk}L}\epsilon_{jt} \\ & - \frac{G^m_{jk}}{1-(1-G^m_{jk})\hat{\rho}_{jk}L} h^m_{jk}(\kappa^m_j)^{-1/2}u^m_{jt} \\ & = \left(\sum_{s=0}^{+\infty}\hat{\rho}^s_{jk}\left[1 - G^m_{jk}\left(\sum_{i=0}^{s}(1-G^m_{jk})^i\right)\right]L^s\right)\epsilon_{jt} \\ & - G^m_{jk}\sum_{s=0}^{+\infty}\hat{\rho}^s_{jk}(1-G^m_{jk})^s L^s h^m_{jk}(\kappa^m_j)^{-1/2}u^m_{jt} \\ \tilde{R}evision^m_{jkt} & = \frac{G^m_{jk}}{1-(1-G^m_{jk})\hat{\rho}_{jk}L}\epsilon_{jt} \\ & + \frac{G^m_{jk}[1-\hat{\rho}_{jk}L]}{1-(1-G^m_{jk})\hat{\rho}_{jk}L} h_{jk}(\kappa^m_j)^{-1/2}u^m_{jt} \\ & = G^m_{jk}\sum_{s=0}^{+\infty}\hat{\rho}^s_{jk}(1-G^m_{jk})^s L^s\epsilon_{jt} \\ & + G^m_{jk}\left(1 - \frac{G^m_{jk}}{1-G^m_{jk}}\sum_{s=1}^{+\infty}\hat{\rho}^s_{jk}(1-G^m_{jk})^s L^s\right)h^m_{jk}(\kappa^m_j)^{-1/2}u^m_{jt} \end{aligned}$$

Therefore,

$$\begin{aligned} Cov\left(Error^m_{jkt}, Revision^m_{jkt}\right) = \quad & \frac{G^m_{jk}(1-G^m_{jk})}{1-\hat{\rho}^2_{jk}(1-G^m_{jk})^2}\gamma^{-1} - (G^m_{jk})^2\left(1 - \frac{G^m_{jk}}{1-G^m_{jk}}\frac{\hat{\rho}^2_{jk}(1-G^m_{jk})^2}{1-\hat{\rho}^2_{jk}(1-G^m_{jk})^2}\right)(h^m_{jk})^2(\kappa^m_j)^{-1} \\ & - (\hat{\rho}_{jk} - \rho_j)G^m_{jk}(1-G^m_{jk})\frac{2\hat{\rho}_{jk}(1-G^m_{jk})(1-\rho^2_j) - (\hat{\rho}_{jk}-\rho_j)[1+\rho_j\hat{\rho}_{jk}(1-G^m_{jk})]}{[1-\rho_j\hat{\rho}_{jk}(1-G^m_{jk})][1-\rho^2_j][1-\hat{\rho}^2_{jk}(1-G^m_{jk})^2]}\gamma^{-1} \end{aligned}$$

and

$$
\begin{aligned}
V(Revision^m_{jkt}) = \quad & \frac{(G^m_{jk})^2}{1-\frac{\hat{\rho}_{jk}}{\rho_j}(1-G^m_{jk})}\left(\frac{G^m_{jk}\frac{\hat{\rho}_{jk}}{\rho_j}[1-\hat{\rho}^2_{jk}(1-G_{jk})]}{[1-\rho_j\hat{\rho}_{jk}(1-G_{jk})][1-\hat{\rho}^2_{jk}(1-G_{jk})^2]} - (\hat{\rho}_{jk}-\rho_j)\frac{1-\rho_j\hat{\rho}_{jk}}{[1-\rho_j\hat{\rho}_{jk}(1-G_{jk})](1-\rho_j^2)}\right) \\
& +(G^m_{jk})^2\left(1+\left(\frac{G^m_{jk}}{1-G^m_{jk}}\right)^2\frac{\hat{\rho}^2_{jk}(1-G^m_{jk})^2}{1-\hat{\rho}^2_{jk}(1-G^m_{jk})^2}\right)(h^m_{jk})^2\kappa_j^{-1} \\
& = (G^m_{jk})^2\frac{1}{1-\hat{\rho}^2_{jk}(1-G^m_{jk})^2}\gamma^{-1} + (G^m_{jk})^2\left(1+\left(\frac{G^m_{jk}}{1-G^m_{jk}}\right)^2\frac{\hat{\rho}^2_{jk}(1-G^m_{jk})^2}{1-\hat{\rho}^2_{jk}(1-G^m_{jk})^2}\right)(h^m_{jk})^2(\kappa^m_j)^{-1} \\
& -(\hat{\rho}_{jk}-\rho_j)(G^m_{jk})^2\frac{2\hat{\rho}_{jk}(1-G^m_{jk})(1-\rho_j^2)-(\hat{\rho}_{jk}-\rho_j)[1+\rho_j\hat{\rho}_{jk}(1-G^m_{jk})]}{[1-\rho_j\hat{\rho}_{jk}(1-G^m_{jk})][1-\rho_j^2][1-\hat{\rho}^2_{jk}(1-G^m_{jk})^2]}\gamma^{-1}
\end{aligned}
$$

Therefore, if $\hat{\rho}_{jk} = \rho_j$, then

$$
\begin{aligned}
\beta^{CGm}_{jk} = \beta^{CG}(\rho_j) \quad & = \frac{\frac{G^m_{jk}(1-G^m_{jk})}{1-\rho_j^2(1-G^m_{jk})^2}\gamma^{-1}-(G^m_{jk})^2\left(1-\frac{G^m_{jk}}{1-G^m_{jk}}\frac{\rho_j^2(1-G^m_{jk})^2}{1-\rho_j^2(1-G^m_{jk})^2}\right)(h^m_{jk})^2(\kappa^m_j)^{-1}}{(G^m_{jk})^2\frac{1}{1-\rho_j^2(1-G^m_{jk})^2}\gamma^{-1}+(G^m_{jk})^2\left(1+\left(\frac{G^m_{jk}}{1-G^m_{jk}}\right)^2\frac{\rho_j^2(1-G^m_{jk})^2}{1-\rho_j^2(1-G^m_{jk})^2}\right)(h^m_{jk})^2(\kappa^m_j)^{-1}} \\
& = \frac{\frac{1-G^m_{jk}}{G^m_{jk}}\gamma^{-1}-[1-\rho_j^2(1-G^m_{jk})](h^m_{jk})^2(\kappa^m_j)^{-1}}{\gamma^{-1}+[1-\rho_j^2(1-2G^m_{jk})](h^m_{jk})^2(\kappa^m_j)^{-1}}
\end{aligned}
$$

If $\hat{\rho}_{jk} \neq \rho_j$, then

$$
\beta^{CGm}_{jk} \quad = \beta^{CG}(\hat{\rho}_{jk}) - \frac{(\hat{\rho}_{jk}-\rho_j)\chi}{V(\tilde{R}evision^m_{jkt})-(\hat{\rho}_{jk}-\rho_j)\chi}[1-\beta^{CG}(\hat{\rho}_{jk})]
$$

with $\chi = (G^m_{jk})^2\frac{2\hat{\rho}_{jk}(1-G^m_{jk})(1-\rho_j^2)-(\hat{\rho}_{jk}-\rho_j)[1+\rho_j\hat{\rho}_{jk}(1-G^m_{jk})]}{[1-\rho_j\hat{\rho}_{jk}(1-G^m_{jk})][1-\rho_j^2][1-\hat{\rho}^2_{jk}(1-G^m_{jk})^2]}\gamma^{-1}$.

G.5 Proof of Proposition 4

Consider Equations (G.4) and (G.5). We can rewrite them as follows:

$$
\begin{aligned}
Revision^m_{ijkt} \quad & = E^m_{ijkt}(x_{jt}) - E^m_{ijkt-1}(x_{jt-1}) \\
& = \frac{G^m_{jk}[1-\hat{\rho}_{jk}L]}{1-(1-G^m_{jk})\hat{\rho}_{jk}L}(1-h^m_{jk})(\tau^m_{jk})^{-1/2}e^m_{ijkt} + \text{terms specific to } \{j,k,m,t\} \\
Error^m_{ijkt} \quad & = x_{jt} - E^m_{ijkt}(x_{jt}) \\
& = -\frac{G^m_{jk}}{1-(1-G^m_{jk})\hat{\rho}_{jk}L}(1-h^m_{jk})(\tau^m_{jk})^{-1/2}e^m_{ijkt} + \text{terms specific to } \{j,k,m,t\}
\end{aligned}
$$

The estimated coefficient is then equal to the covariance between the error and the revision conditional on all the terms that are country-location-time specific, divided by the variance of the revision conditional on all the terms that are country-location-time specific

$$
\begin{aligned}
\beta_{jk}^{FEm} &= \frac{Cov\left(-\frac{G_{jk}^m}{1-(1-G_{jk}^m)\hat{\rho}_{jk}L}(1-h_{jk}^m)(\tau_{jk}^m)^{-1/2}e_{ijkt}^m, \frac{G_{jk}^m[1-\hat{\rho}_{jk}L]}{1-(1-G_{jk}^m)\hat{\rho}_{jk}L}(1-h_{jk}^m)(\tau_{jk}^m)^{-1/2}e_{ijkt}^m\right)}{V\left(\frac{G_{jk}^m[1-\hat{\rho}_{jk}L]}{1-(1-G_{jk}^m)\hat{\rho}_{jk}L}(1-h_{jk}^m)(\tau_{jk}^m)^{-1/2}e_{ijkt}^m\right)} \\
&= \frac{-(G_{jk}^m)^2\left(1-\frac{G_{jk}^m}{1-G_{jk}^m}\frac{\hat{\rho}_{jk}^2(1-G_{jk}^m)^2}{1-\hat{\rho}_{jk}^2(1-G_{jk}^m)^2}\right)(1-h_{jk}^m)^2(\tau_{jk}^m)^{-1}}{(G_{jk}^m)^2\left(1+\left(\frac{G_{jk}^m}{1-G_{jk}^m}\right)^2\frac{\hat{\rho}_{jk}^2(1-G_{jk}^m)^2}{1-\hat{\rho}_{jk}^2(1-G_{jk}^m)^2}\right)(1-h_{jk}^m)^2(\tau_{jk}^m)^{-1}}
\end{aligned}
$$

Hence the result.

G.6 Proof of Proposition 5

We can write $Disagreement_{jt}$, $Revisionjt$ and x_{jt} as a function of the current shocks and past variables:

$$
\begin{aligned}
Disagreement_{jt}^m = \;& G_{jl}^m(x_{jt} + h_{jl}^m(\kappa_j^m)^{-1/2}u_{jt}^m) + (1 - G_{jl}^m)E_{jlt-1}^m(x_t) \\
& -G_{jf}^m(x_{jt} + h_{jf}^m(\kappa_j^m)^{-1/2}u_{jt}^m) - (1 - G_{jf}^m)E_{jft-1}^m(x_t) \\
= \;& G_{jl}^m(\epsilon_{jt} + \rho_j x_{jt-1} + h_{jl}^m(\kappa_j^m)^{-1/2}u_{jt}^m) + (1 - G_{jl}^m)E_{jlt-1}^m(x_t) \\
& -G_{jf}^m(\epsilon_{jt} + \rho_j x_{jt-1} + h_{jf}^m(\kappa_j^m)^{-1/2}u_{jt}) - (1 - G_{jf}^m)E_{jft-1}^m(x_t) \\
= \;& (G_{jl}^m - G_{jf}^m)\epsilon_{jt} + (G_{jl}^m h_{jl}^m - G_{jf}^m h_{jk}^m)(\kappa_j^m)^{-1/2}u_{jt}^m \\
& +\rho_j(G_{jl}^m - G_{jf}^m)x_{jt-1} + (1 - G_{jl}^m)E_{jlt-1}^m(x_t) - (1 - G_{jf}^m)E_{jft-1}^m(x_t) \\
Revision_{jt}^m = \;& G_j^m[(x_{jt} + h_j^m(\kappa_j^m)^{-1/2}u_{jt}^m) - E_{jt-1}^m(x_t)] \\
= \;& G_j^m[\epsilon_{jt} + \rho_j x_{jt-1} + h_j^m(\kappa_j^m)^{-1/2}u_{jt}^m) - E_{jt-1}^m(x_t)] \\
= \;& G_j^m\epsilon_{jt} + G_j^m h_j^m(\kappa_j^m)^{-1/2}u_{jt} + \rho_j G_j^m x_{jt-1} - G_j^m E_{jt-1}^m(x_t) \\
x_{jt} = \;& \epsilon_{jt} + \rho_j x_{t-1}
\end{aligned}
$$

The estimated coefficient is given by

$$
\begin{aligned}
\beta_j^{DISm} &= \frac{Cov\left(h_j^m G_j^m(\kappa_j^m)^{-1/2}u_{jt}^m, (h_{jl}^m G_{jl}^m - h_{jf}^m G_{jf}^m)(\kappa_j^m)^{-1/2}u_{jt}^m\right)}{V\left(h_j^m G_j^m(\kappa_j^m)^{-1/2}u_{jt}^m\right)} \\
&= \frac{h_{jl}^m G_{jl}^m - h_{jf}^m G_{jf}^m}{h_j^m G_j^m}
\end{aligned}
$$

Hence the result.
Consider the rational expectations case. The Kalman filter is given by: $G_{jk}^m = \Phi_{jk}(\Phi_{jk} + (\lambda_{jk}^m)^{-1})^{-1}$ and $h_{jk}^m = \kappa_j^m/\lambda_{jk}^m$. We can thus rewrite:

$$
h_{jk}^m G_{jk}^m = \frac{\kappa_j^m}{\lambda_{jk}^m + \Phi_{jk}^{-1}}
$$

For $h^m_{jk}G^m_{jk}$ to be decreasing in τ^m_{jk}, it is enough that $\lambda^m_{jk}+\Phi^{-1}_{jk}$ is increasing in λ^m_{jk}. We use the definition of Φ_{jk} in (G.2) to compute this derivative:

$$\begin{aligned}
\frac{\partial(\lambda^m_{jk}+\Phi^{-1}_{jk})}{\partial\lambda^m_{jk}} &= 1+\frac{1}{2}(1-\rho^2_j)\frac{1}{(\lambda^m_{jk})^2}\left(1-\frac{(1-\rho^2_j)(\lambda^m_{jk})^{-1}-\gamma^{-1}_j}{\sqrt{(\gamma^{-1}_j-(1-\rho^2_j)(\lambda^m_{jk})^{-1})^2+4\gamma^{-1}_j}}\right)\\
&= 1+\frac{1}{2}(1-\rho^2_j)\frac{1}{(\lambda^m_{jk})^2}\left(\underbrace{\frac{\sqrt{(\gamma^{-1}_j-(1-\rho^2_j)(\lambda^m_{jk})^{-1})^2+4\gamma^{-1}_j}+\gamma^{-1}_j-(1-\rho^2_j)(\lambda^m_{jk})^{-1}}{\sqrt{(\gamma^{-1}_j-(1-\rho^2_j)(\lambda^m_{jk})^{-1})^2+4\gamma^{-1}_j}}}_{>0}\right)
\end{aligned}$$

$h^m_{jk}G^m_{jk}$ is therefore decreasing in τ^m_{jk}.

Consider the case with behavioral biases. h_{jk} and G_{jk} are identical except that they reflect the forecasters' perceived parameters $\hat{\rho}_{jk}$ and $\hat{\tau}^m_{jk}$. As a consequence, $h^m_{jk}G^m_{jk}$ is decreasing in $\hat{\tau}^m_{jk}$. Therefore, for a given $(\hat{\tau}^m_{jk})^{-1}-(\tau^m_{jk})^{-1}$, $h^m_{jk}G^m_{jk}$ is decreasing in τ^m_{jk}. If the foreign and local forecasters have the same behavioral biases $\hat{\rho}_{jk}$ and $(\hat{\tau}^m_{jk})^{-1}-(\tau^m_{jk})^{-1}$, then differences in $h^m_{jk}G^m_{jk}$ reflect differences in τ^m_{jk}

CHAPTER 3

Inflation Expectations, Perceptions and News Media: Regional Differences in Switzerland

Elio Bolliger

University of Lausanne

Abstract

This paper studies newspaper inflation reporting and its effects on inflation expectations and perceptions in Switzerland. We create a standard quantitative inflation news measure and a novel qualitative measure of inflation sentiment for newspapers written in two national languages. To study the effects of news on inflation expectations and perceptions, we first check for the existence of a negativity bias in inflation news reporting. Second, we exploit the language barrier in Switzerland to analyse the effects of inflation media shocks on regional inflation expectations and perceptions. We highlight three findings. First, we find no evidence of a negativity bias in French- and German-written newspapers. Second, both the quantitative and qualitative news significantly affect expectations and perceptions. Third, we document socio-demographic differences in the effect of news across the language border.

Keywords: Inflation News Coverage, Inflation Expectations and Perceptions, Natural Language Processing.

JEL: C32, E31, D84.

1. Introduction

Since the Covid-19 crisis, inflation rates have reached historical heights in many countries around the entire globe. Theoretically and empirically, inflation expectations are an important driver of inflation. For households, expectations are strongly linked to inflation perceptions (Weber et al., 2022). Survey evidence shows that households draw their information about the current price levels from their personal shopping experience as well as social media and news from television and newspapers (Blinder and Krueger, 2004; Kumar et al., 2015; Cavallo et al., 2017; D'Acunto et al., 2021).

This paper focuses on the latter aspect. Using a comprehensive data set of newspaper articles in Switzerland, we study the inflation news coverage and its effects on households' inflation expectations and perceptions. First, we construct a quantitative inflation news measure and a novel inflation news sentiment for newspapers written in two of the national languages in Switzerland, French and German.[1] Second, we follow an empirical approach developed by Gambetti et al. (2023), henceforth GMZ, using a Threshold Structural Autoregression (TSVAR) to study differences in inflation news reporting. Third, we exploit the language barrier in Switzerland as a quasi-natural experiment to analyse how news reporting affects households' inflation expectations and perceptions.

Since the onset of the Covid-19 crisis, inflation rates in Switzerland have steadily increased, reaching a peak of 3.5% in August 2022. While this may seem relatively low compared to other countries such as the United States, which saw an inflation rate of 8.2% during the same period, this level of inflation is the highest that Switzerland has experienced in the past 25 years, and is comparable only to the levels observed during the financial crisis.

We start by discussing the two inflation news measures and descriptive results. The quantitative inflation news measure is defined as the difference of the number of articles reporting an inflation increase minus the number of articles reporting an inflation decrease. This measure is inspired by GMZ and Soroka (2006). The difference of articles reporting an increase versus decrease of inflation reflects the prevailing tone of inflation reported in newsarticles. We show that this measure correlates well with the actual inflation rate for news articles in both languages, French and German.

The quantitative inflation news measure, however, is not a sentiment-based indicator. A decrease (increase) in the inflation rate is not necessarily assessed as good (bad) by newspapers, as argued in Soroka (2006). We highlight this in an example. Switzerland experienced a rather strong inflation decrease after 2015 when the inflation rate plunged into negative ter-

[1] Switzerland has four national languages: German, French, Italian and Romansh which are reported by 62.3%, 22.8%, 8% and 0.5% of the population in Switzerland as their main language, respectively. Respondents could report up to 3 main languages. 23.1% reported a different main language (Federal Statistical Office, 2022b).

ritory. During this time, newspapers in both languages reported in a negative way about the level of the inflation rate, even warning about the risk of a deflationary spiral. To better assess the sentiment of news reporting, we construct a novel qualitative inflation news measure to explore differences in sentiment news reporting and its effects on households' expectations and perceptions.

This qualitative news measure is the positive or negative sentiment derived from the same set of articles using a state-of-the-art Natural Language Processing (NLP) Model. The model, called BERT and short for Bidirectional Encoder Representations, was originally developed at Google by Devlin et al. (2018b). It has been used extensively for various NLP tasks, including sentiment analysis (see, among others, Nemes and Kiss (2021) or Lee (2020)). We fine-tune this model on a set of 1000 self-annotated inflation articles for each language and show that this measure captures different information than the quantitative news measure of inflation.

We proceed by analysing newspaper inflation reporting using a TSVAR. In more detail, we study the dynamic responses of the quantitative and qualitative inflation news measure to unexpected positive and negative changes in the inflation rate. We do this for two reasons.

First, we want to analyse whether newspapers written in French and German report in a systematically different way about an inflation increase or decrease. Such differences are referred to as negativity bias in newspapers, where newspapers tend to over-report negative outcomes (inflation increase) over positive (inflation decrease) to attract more attention from news-readers (Soroka, 2006)Soroka 2006, GMZ). We find no evidence that such a negativity bias is present in Swiss newspapers written in French and German. Furthermore, we find no systematic differences in the quantitative and qualitative news measure across regions. These results will be helpful in interpreting the effects of news on inflation expectations and perceptions (more on that below).

Second, we use the TSVAR to identify media shocks that are unrelated to the current inflation dynamics but affect the inflation news reporting. For the identification, we exploit the high frequency of the newspaper data. We then aggregate these shocks such that they match the quarterly frequency of our survey data to study the effects of news on inflation expectations and perceptions. Arguably, Swiss newspapers written in French and in German report about the same national inflation rate. We assume that households living in the French-speaking part of Switzerland mostly consume news written in French and those in the German-speaking part news written in German. This allows us to exploit regional differences in news reporting to investigate how the media shocks affect households' expectations and perceptions while controlling for household characteristics as well as time fixed-effects.

We find that the quantitative news measure has a small but significant effect on inflation expectations and perceptions. When relatively more news about an inflation increase versus an inflation decrease are published, inflation expectations and perceptions increase. This effect is mostly significant during times when inflation goes up. For the qualitative sentiment measure, we find that the effect is countercyclical to the current inflation environment. When inflation is positive, a more positive assessment of inflation in news reduces expectations and perceptions, whereas when inflation is negative, a positive sentiment lifts up expectations and perceptions.

Furthermore, we document differences in the effect of news across the language border and age. We find that both the quantitative and qualitative media shocks have a significantly higher effect for households living in the German-speaking part compared to the households living in the French-speaking part. While our results from the TSVAR rule out that this effect is due to systematically different news reporting across regions, it is possible that other unobserved characteristics may explain this result. For example, Jost (2018) found that households in the German-speaking part have a significantly higher level of inflation aversion than households in the French-speaking part, which may make them more receptive to news about inflation.

Alternative explanations could be that differences in shopping experience (on a regional and cross-border level) might drive these results or that households directly observe different inflation rates in their cantons. However, Kluser (2023) found no evidence of regional price discrimination for one of the largest Swiss retailer and inflation rates in France and Germany are very similar. In addition, the official cantonal inflation rates of the three biggest Swiss cities, Basel, Zürich and Geneva are much the same. Therefore, it is unlikely that differences in shopping experience or inflation rates drive the effect of news across the language border.

With respect to age, we find some evidence that the effect of the quantitative and qualitative media shock is stronger for relatively elderly households compared to younger households. This is in line with studies that show elderly people spend more time reading news in newspapers than younger people (see, for example, Lee and Delli Carpini (2010) or Federal Office of Communications (2016)). An alternative explanation for the higher effect across age is that they pay more attention to inflation news due to their higher inflation exposure through their savings (Doepke and Schneider, 2006) or consumption expenditures (Basso et al., 2023).

These results are policy relevant. As news affects households' expectations and perceptions, it may have several real effects. First, well-anchored inflation expectations increase the effectiveness of monetary policy (Lamla and Lein, 2014; Nautz and Strohsal, 2015). Second, inflation expectations may be self-fulfilling (Leduc et al., 2007). Third, our results provide guidance for policy communication. While we find a significant number of articles mentioning the Swiss National Bank in the analysed news, the higher effect among elderly households suggests that policy makers should exploit various communication channels to reach all types of households. Fourth, we provide a new inflation sentiment measure that tracks how newspapers assess inflation for newspapers written in both French and German. This measure may be valuable for policy institutions to summarise how inflation is assessed by the news media in Switzerland.

Our paper broadly relates to three different strands of literature. First, it contributes to the literature in political science and economics about news reporting of economic events. Most findings in this literature document a negativity bias, i.e., negative news receive a relatively higher coverage than positive news (see, for example, Goidel and Langley (1995), Fogarty (2005), Soroka (2006), Soroka (2012)). However, for unemployment news coverage in the United States, GMZ find that the negativity bias disappears when taking into account the non-linearity in the unemployment rate response to economic shocks itself. Using the approach of GMZ, we are the first to explore the news coverage of inflation and find no evidence for the negativity bias in Switzerland for newspapers written in French and German.

Second, our paper belongs to a literature that studies the role of information for expectations and perceptions. Motivated by increasing empirical evidence about the rejection of the full-information rational expectation (FIRE) model, there is a growing body of literature studying different sources of information economic agents use and how they affect their economic decisions, perceptions and expectations.[2] For instance, recent survey evidence finds that shopping experience is an important driver of inflation perception and expectations (Cavallo et al., 2017; D'Acunto et al., 2021). For newspapers, Larsen et al. (2021) find that news media coverage plays an important role in the expectation formation process. Moreover, households update their expectations more often during periods of high news coverage, for example during recessions (Carroll, 2003; Doms and Morin, 2004). Switzerland provides an interesting framework to study the effects of media to inflation expectations and perceptions. While Swiss newspapers report about the same national inflation rate, households in the different language regions are likely to consume news in the language predominant in their region. This assumption allows us to be more restrictive in our econometric set-up than the above mentioned studies by including time fixed-effects. To the best of our knowledge, we are the first to exploit this language barrier to study how inflation news reporting affect inflation expectations and perceptions.

Third, our paper contributes to the literature that studies cultural aspects of economies and its link to economic outcomes. For example, a high inflation aversion is often attributed to Germany due to its experience of hyperinflation in 1923 (Cukierman, 1992; Hayo, 1998; Issing, 2005; Beyer et al., 2008; Ehrmann and Tzamourani, 2012). In Switzerland, Eugster et al. (2017) and Jost (2018) both exploit the language border to study how culture affects unemployment and monetary policy preferences, respectively. Jost (2018) document that

[2]For the rejection of the FIRE, see, for example, Mankiw et al. (2003); Coibion and Gorodnichenko (2015); Bordalo et al. (2020); Kohlhas and Broer (2019) or Angeletos et al. (2020).

households in the French-speaking part have a significantly lower inflation aversion than households in the German-speaking part. We show that news has a significantly higher effect for households in the German-speaking part of Switzerland compared to those in the French-speaking part. Potentially, this difference in the effect of news is linked to the different level of inflation aversion across the language border.

The rest of the paper is organized as follows. In section 2, we describe the newspaper article data used and construct the quantitative and qualitative inflation news measure. In Section 3, we check for the existence of a negativity bias in inflation news reporting and systematic differences across regions. Section 4 discusses the household survey used for inflation expectations and perceptions. In Section 5, we analyse how media shocks affect inflation expectations and perceptions. Section 6 provides several robustness checks. Finally, Section 7 concludes.

2. Newspaper Articles Data

In this section, we outline the newspaper article data and how we derive the quantitative and qualitative news measures of inflation. Moreover, we provide some descriptive results.

2.1 Selection of Articles

To retrieve newspaper articles in Switzerland written in German and French we use a novel database, called "Swissdox". Swissdox is an online media database that aims to provide the broadest possible coverage of the Swiss media landscape Swissdox (2022). The database contains more than 20 million news articles. In this study, we focus on the largest regional media outlets in the French-speaking and German-speaking part of Switzerland that have the longest time coverage in the database. For newspapers written in German, we include articles from *Berner Zeitung, Tages Anzeiger, Blick* and *20 Minuten*.[3] French-written newspapers encompass articles from *Le Temps, 24 heures, Tribune de Genève, 20 minutes* and *Le Matin*. For all newspapers, we focus on printed articles.

We first clean the text of the articles, remove the most frequent stopwords and apply lemmatization to the remaining words. After that, we follow closely GMZ to select articles that indicate an inflation increase or decrease. For this, we search for economic articles that contain the words inflation or prices. If these words are preceded or followed by indicators of an increase or decrease in the distance of five-words, the article is classified as inflation in-

[3]Note that we exclude the German-written Neue Zürcher Zeitung (NZZ) from the main analysis in this paper. Swissdox classifies NZZ as a national newspaper. As one of our assumption is that households in the French (German) speaking part consume only French (German) written newspapers, we prefer to focus on more regional newspapers and believe leaving out the NZZ is a more prudent approach for this.

crease or inflation decrease, respectively.[4] Appendix A provides a detailed description of all words considered for both languages and provides an overview of the articles collected using wordclouds. In this paper, and similar to GMZ, we do not distinguish whether the articles write about past, current and/or future inflation. For newspapers written in German, we find a total of 10,520 articles and 13,407 articles for newspapers written in French.

2.2 Quantitative News Measure

The quantitative news measure is defined as the difference between the number of articles writing about an increase of inflation and those that report a decrease of inflation. An article is counted as an article writing about a decrease of inflation if the article mentions a decrease more often than an increase of inflation.[5]

The quantitative news measure is an indicator of the prevailing information on inflation in the newspaper at a specific point in time. If the number of articles indicating an increase (decrease) in inflation dominates the number of articles writing about a decrease (increase) in inflation, the indicator takes on a more positive (negative) value.

Over the entire sample period, we observe 8,180 German-written articles that write about an inflation increase and 2,340 that write about an inflation decrease. For French-written articles, we have 7,646 and 5,761, respectively. Overall, we identify more articles to write about an inflation increase relative to an inflation decrease for German-written articles.

2.3 Descriptive Results

Figure 2.1 plots the time series of news articles reporting an inflation increase or a decrease together with the observed inflation rate for each language separately. Panel 2.1A plots the inflation rate and the number of articles that report an inflation increase. In general, the number of articles reporting an inflation increase in German-written newspapers increases when inflation rises. The reporting of articles about an inflation increase peaked during the financial crisis with 101 articles published in June 2008.

[4] Note that the results are very similar using a different word distance.

[5] Note that, similar to GMZ, we could also only consider articles that write exclusively about increases or decreases of inflation. The results are robust to this alternative definition. However, to increase the total number of articles considered in this analysis we preferred sticking to this measure.

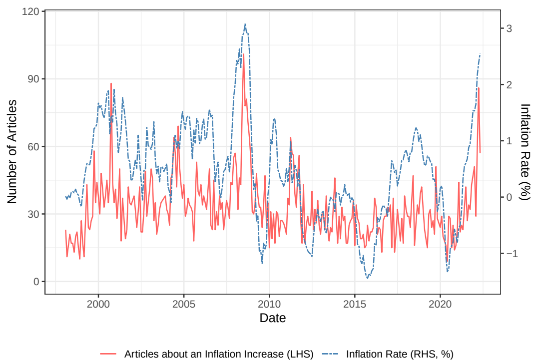

(A) German - Increase

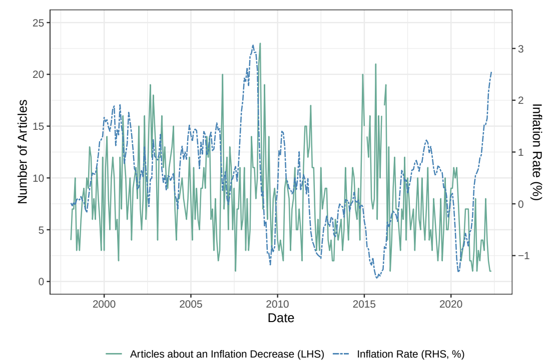

(B) German - Decrease

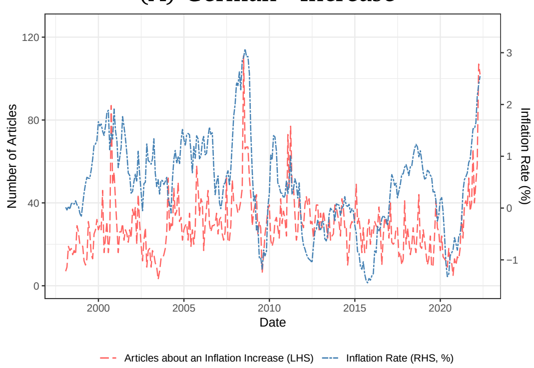

(C) French - Increase

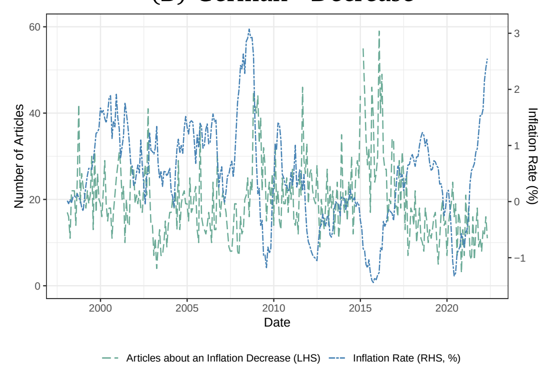

(D) French - Decrease

Figure 2.1: Articles writing about an Inflation Increase and Decrease

Note: On the left-hand side, the panels show the number of articles that write about an inflation increase (left-hand scale) together with the actual observed inflation rate (right-hand scale) for both newspapers written in German and French, respectively. The observed inflation rate is the change in percentage compared to the same month in the previous year. On the right-hand side, the panels show the number of articles writing about an inflation decrease (left-hand scale) together with the observed inflation rate (right-hand scale) for both newspapers written in German and French, respectively.

Panel 2.1B plots the inflation rate together with the number of articles writing about an inflation decrease. During the financial crisis, the maximum articles mentioning an inflation decrease was 21 in December 2009. The highest number of articles published in a given month was reached in January 2015 with 37 articles, which was during a period of relatively strong negative inflation in Switzerland.

Panel 2.1C and 2.1D plot the same indices for all articles from French-written newspapers. In general, the reporting of inflation increase and decrease follows the evolution of the inflation rate. One exception is the period between January 2003 to December 2003, with unusually low levels of articles writing about inflation increase or decrease. This is due to a reporting issue in the Swissdox database. For this reason, we will discard this time period from our sample for both, the German articles and the French articles.

Figure 2.2 plots the quantitative inflation news measure for both languages. This measure is the difference between the number of articles writing about an increase minus the number of articles writing about a decrease of inflation. When this measure increases, there are relatively more articles published that report an inflation increase versus an inflation decrease. Therefore, it summarises the prevailing information about inflation for newspapers written in German and French.

As Soroka (2006) argues, for households, news about an inflation decrease are good news, whereas news about an inflation increase are bad news. Hence, an increase in the measure of quantitative inflation news could be interpreted as relatively more bad news versus good news. However, for inflation in Switzerland, this is not necessarily always the case. For example, the period around 2015 is described by a negative inflation and, at the same time, a decrease in economic activity. Newspapers reported negatively about the inflation decrease, warning about a deflationary spiral.[6]

If the quantitative news measure represents indeed the prevailing information on inflation, we would expect a positive correlation of this measure with the actual inflation rate. For German articles in panel 2.2A, the quantitative news measure tracks the evolution of the inflation rate considerably well. Over the entire sample period, the correlation between the quantitative news measure and the observed inflation rate is 0.59. A similar pattern can be observed for the French-written articles in panel 2.2B. The correlation between the quantitative news measure and the observed inflation rate is only slightly lower at 0.52.

Both the quantitative news measure for German and French-written newspaper show similar spikes in September 2000, the global financial crisis and during the Covid-19 crisis. The quantitative new measure for German-written articles is more often positive compared to the quantitative news measure from the French-written articles. This follows from the relatively higher number of articles published about an inflation increase compared to a decrease for German-written newspapers.

2.4 Qualitative Inflation News Measure

As mentioned by GMZ, the quantitative inflation news measure is not a sentiment measure. It summarizes the prevailing information about inflation by counting the articles that write about an inflation increase compared to those that write about an inflation decrease, but this information does not necessarily reflect the sentiment attached to inflation in the news articles (i.e., whether newspapers assess inflation to be positive or negative).

There is a growing literature that exploits text information from newspapers that extract the general economic sentiment and how this sentiment effects macroeconomic outcomes, such as consumer confidence or consumption (see, for example, Shapiro et al. (2022) or Starr

[6]For example, the title and lead of an article in the German-written newspaper 20 Minuten in December 2014 translates as follows: *"How falling prices hurt the economy - The Swiss National Bank expects inflation to fall into negative territory in 2015. Even if falling prices seem positive - they do not always have good consequences."*, Frommberg (2014).

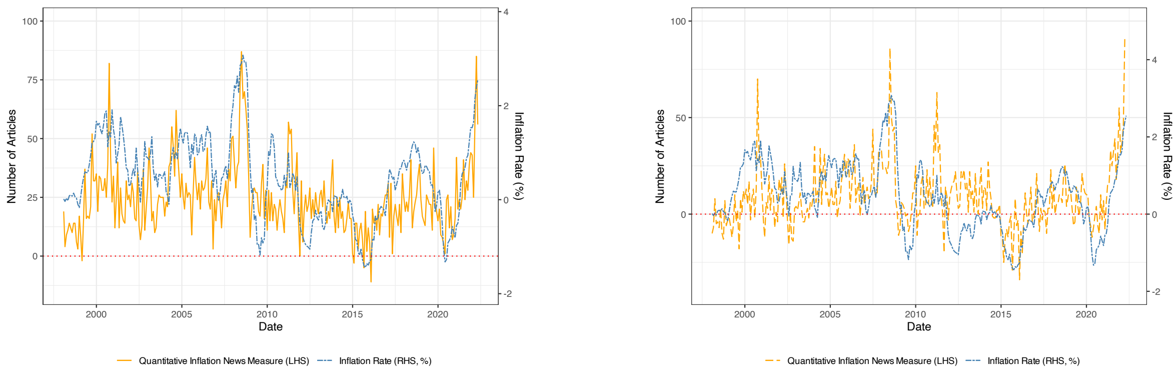

(A) Newspapers written in German (B) Newspapers written in French

Figure 2.2: Quantitative Inflation News Measure

Note: In the left panel, the figure shows the quantitative inflation news measure (right-hand scale) of news calculated as the difference between the number of articles writing about an inflation increase minus the number of articles writing about an inflation decrease using only newspapers written in German, together with the observed inflation rate (right-hand scale). The left panel shows the same time series of quantitative inflation news measure using newspapers written in French (left-hand scale) together with the actual inflation rate (right-hand scale).

(2012)). However, to the best of our knowledge, no such sentiment indicator has been created with a model specifically trained on inflation articles.

For this reason, we develop the qualitative inflation news measure which represents the sentiment of the inflation news articles. To derive this sentiment, we use a state-of-the-art NLP model called BERT developed at Google by Devlin et al. (2018b). Their seminal contribution is a model that generates context-aware embeddings for words and documents. In contrast to alternative NLP models that process text only in one direction (either from left to right or right to left), BERT uses a bidirectional approach.

Context aware embeddings are achieved by pre-training the model on two specific tasks called Masked Language Modelling (MLM) and Next Sentence Prediction (NSP), respectively. In MLM, some percentage of the input text is masked at random. The goal of the model is then to predict these masked words. In NSP, the model learns about the relationship between two sentences. As input, the model receives two sentences A and B. The goal of the model is to decide whether sentence B follows A, where this is in 50% of the cases true and in 50% B is a random sentence from the Corpus.

BERT has been used extensively for various NLP tasks in economics and finance, including sentiment analysis (see, for example, Araci (2019) or Sousa et al. (2019)).[7] For our sentiment analysis, we use the pre-trained multilingual base model (Devlin et al., 2018a) for the French language and the BERT model provided by Deepset (Deepset, 2019) for the German language.

[7]Other popular approaches, such as sentiment analysis via dictionary method would require comparable dictionaries for the two languages that classify sentiment similarly. The most common dictionary from Loughran and McDonald (2011) has only been translated to German (see Bannier et al. (2019)), but not to French. Moreover, these dictionaries usually capture the general sentiment, non-specific to inflation.

To fine-tune the model, we create a random sample of 1,000 articles for each language. In these articles, we focus on the paragraphs that write about inflation. We then self-classify these articles into two categories, positive and negative. An article is classified as positive (negative) if inflation is assessed to be positive (negative). For example, in February 2008, a newspaper article of “Le Matin” translates to "*[...] The economy does not need it. Inflationary pressures have just returned to a level not seen for more than a decade. [...]* ". In this article, inflation is negatively assessed. In contrast, the following article of the “NZZ” in March 2002 that translates to "*[...] The Swiss National Bank (SNB) considers the current interest rate level [...] is appropriate for a sustainable and inflation-free economic development. [...] An increase would be inappropriate in view of the favourable inflation outlook [...]*". Here, the level of inflation is positively connoted.

Our training data set consists of 85% of the articles. Validating the predictions on our test data set, we achieve an accuracy of 72% and 68% to predict the sentiment for the French and German language, respectively. Due to the novelty of this inflation sentiment measure, the accuracy can be compared most closely with the financial sentiment prediction of Araci (2019). In their paper, Araci (2019) achieved an accuracy of 86%. However, the sentiment prediction of Araci (2019) is based on single sentences and the model was pre-trained on financial vocabulary. In contrast, due to the different languages considered, we use two models that were trained on non-specialized text corpora and use entire paragraphs to predict the sentiment of inflation.

Using the fine-tuned model, we predict the sentiment of the same articles considered for the quantitative inflation news measure. Then, we average the prediction of negative and positive sentiment for a given month. We standardize the time series by subtracting the mean over the entire sample period and dividing by its standard deviation.

2.5 Descriptive Results

In Figure 2.3, we plot the qualitative inflation news measure for the newspapers written in German and French, respectively as a four-month moving average (left-hand-scale) together with the actual inflation rate (right-hand-scale). The correlation between the qualitative news measure and the actual inflation rate is not consistently positive or negative.

For example, during the financial crisis, the sentiment for both newspapers written in German and French decreased, while this decrease is more marked for German-written newspapers. At the same time, the inflation rate increased. Similarly, after 2021, the recent spike in inflation correlates negatively with the sentiment measures from German and French-written newspapers. The correlation is -0.54 for the German and -0.37 for the French-written newspapers.

In contrast, there are periods of positive correlation between sentiment and the inflation rate.

At the beginning of the sample from 1998 to 2001, the correlations are 0.37 and 0.46 for the German and French-written newspapers, respectively. Another example is the period from 2010 to 2012 where rising inflation rates are associated with a more positive sentiment for both newspapers.

Inflation is not always assessed similarly across newspapers written in German and French. For example, during a 12-month period after January 2012, French-written newspapers assessed the current level of inflation to be positive, while the newspapers written in German assess it to be more negative. However, there are also periods where both newspapers written in German or French assess the level of inflation in a similar fashion. For example, both sentiment measures are in the negative territory during the Covid 19 crisis.

Our qualitative news measure indicates that a rise (fall) in the quantitative news measure does not always correspond to negative (positive) news for households, which challenges the assessment by Soroka (2006).[8]

For example, from 2015 to 2016, the level of inflation decreased strongly in Switzerland. The qualitative sentiment measure decreases and falls into negative territory during this time. Accordingly, newspapers assess the level of inflation negatively. This is reflected in the text of the newsarticles. For example, the "Tribune de Genève" wrote in August 2018: "[...] *Our central bank had the opportunity to continue to print money, especially since the risk of inflation is zero and we suffer rather from an even more insidious evil, deflation.[...]*".

In the line of argument of Soroka (2006), this period of an inflation decrease should correspond to good news from the perspective of households. As the number of articles writing about an increase of inflation is lower than the number of articles writing about a decrease in inflation, the quantitative inflation news measure decreases during this time. During such periods, our qualitative news measure can provide a more nuanced picture of the actual sentiment of inflation in newspapers.

We stress that our qualitative inflation sentiment measure is not tracking a general economic sentiment. This could be a valid concern if the paragraphs writing about inflation also feature prominent information about GDP and unemployment. However, in that case, we would expect a strong positive correlation of the sentiment with GDP (Shapiro et al., 2022; Starr, 2012). The Figure A2 in the appendix plots the qualitative inflation news measure together with the real GDP growth rate of Switzerland. For both newspapers written in German and in French, the correlation is low at -0.18 and 0.08, respectively.

[8]Note that GMZ follow a similar line of argument to study good and bad news reporting and the households' reaction to these news in the case of unemployment. Arguably, the relation of bad and good news with the actual unemployment rate is less ambiguous compared to inflation. For example, Doepke and Schneider (2006) show that elderly rich households are most affected by inflation whereas younger, middle-class households with mortgage debt might benefit from it.

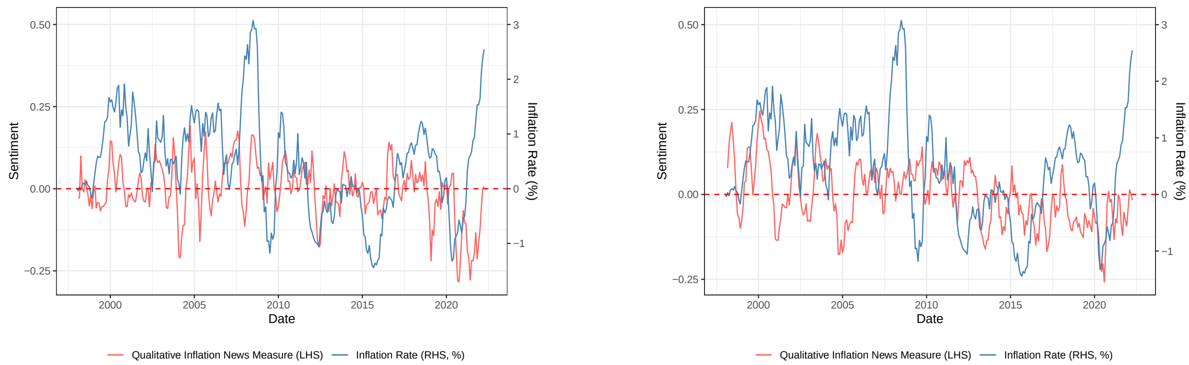

(A) Newspapers written in German

(B) Newspapers written in French

Figure 2.3: Qualitative Inflation News Measure

Note: In the left panel, the figure plots the qualitative inflation news measure for newspapers written in German (left-hand scale) together with the observed inflation rate (right-hand scale). The qualitative news measure is the average sentiment resulting from the prediction of the BERT model. In the right panel, we plot the qualitative news measure for newspapers written in French (left-hand scale) together with the actual inflation rate (right-hand scale). The qualitative inflation news measure is standardized and displayed here with a four-month moving average. Values above (below) the horizontal zero line indicate periods of a predominantly positive (negative) sentiment with respect to inflation.

3. Inflation Reporting of Newspapers

3.1 Model

In this section, we describe the framework from GMZ to analyse how the quantitative and qualitative inflation news measures in newspapers react to changes in inflation rate using a TSVAR model.

We use the novel empirical approach from GMZ for two main reasons. First, it allows to check for a negativity bias in news reporting. Such a negativity bias has been reported in the literature for both unemployment and inflation in US newspapers, where news coverage is relatively higher during periods where inflation and unemployment increases compared to times when inflation and unemployment decreases. News about inflation and unemployment increase are generally referred to as bad news for households (see, for example, Soroka (2006)).

A negativity bias could shape people's exceptions as well as perceptions and lead them to have overly pessimistic news during times when unemployment or inflation increases. Ultimately, this may affect how households respond to these economic conditions (GMZ). Therefore, this analysis will be helpful later on when we interpret the effects of news on inflation expectations and perceptions.

Second, with their approach, we can exploit the frequency of our news measures and inflation rates to identify (non-structural) media shocks. In contrast to our survey data about inflation expectations and perceptions, news data and inflation rates are available at monthly frequency. We then aggregate these shocks to match the quarterly frequency of the survey

data, to study the effects of news on expectations and perceptions.
y_t is a $m \times 1$ time series vector with m endogenous variables of interest and $\tilde{y}_{t-1}$ being a $m \times (mp+1)$ matrix, where p is the lag-order such that $\tilde{y}_{t-1} = (1, y_{t-1}, \dots, y_{t-p})$. Then, a threshold VAR can be written as

$$y_t = (1 - \Gamma(z_t))[\tilde{y}_{t-1}\beta_1] + \Gamma(z_t)[\tilde{y}_{t-1}\beta_2] + \varepsilon_t \tag{3.1}$$

where ε_t is a $m \times 1$ disturbance term with $\varepsilon_t \sim WN(0, \Sigma)$. z_t is a scalar and $\Gamma(\cdot)$ a function taking the value 0 or 1. Similar to GMZ, we set $z_t = \Delta\pi_{t-1}$, the lagged change in the inflation rate. The lag ensures unconfoundedness with the disturbance term ε_t. As we are interested in potential differences in media reporting in cases of an inflation decrease or increase, we set $\Gamma(z_t) = 0$ if $\Delta\pi_{t-1} \leq 0$ and $\Gamma(z_t) = 1$ if $\Delta\pi_{t-1} > 0$. The model can then be estimated as an OLS on two distinct samples, with β_1 being the coefficients that describe the dynamics in case of an inflation decrease, and β_2 the coefficients for the dynamics in case of an inflation increase.
To investigate whether the inflation news measures derived from the newspaper articles react differently during times of an inflation increase or decrease, we analyse the dynamic response of the quantitative and qualitative news measure to an innovation in the inflation rate that is not related to the other shocks in the system. To identify the shock, we follow GMZ by using a Cholesky decomposition. We define C to be the Cholesky factor of Σ. C is a lower triangular matrix such that $\Sigma = CC'$. Therefore, the identified shocks can be calculated by $v_t = C^{-1}\varepsilon_t$. Our first endogenous variable in the system is the change in the inflation rate and the second variable is the inflation news measure. Consequently, the first identified shock v_{1t} is orthogonal to v_{2t}. In that case, v_{1t} represents unexpected changes in the inflation rate, unrelated to past changes in the inflation rate or the inflation news measure.
It's important to note that this shock is a combination of different structural shocks. These shocks drive the forecast error in the inflation rate change for the upcoming month. The impulse response to this shock shows how the variables evolve when the inflation rate unexpectedly changes, meaning when they are higher or lower than expected. The identified shock, though not structural, helps analyze how the inflation news measure reacts to positive or negative changes in the inflation rate, regardless of the underlying shock's nature.
As the observed inflation rate is the same for the French-speaking and German-speaking region, the regime indicator z_t as well as the impact effects in both regimes ($\Gamma(z_t) = 0, \Gamma(z_t) = 1$) are the same. This makes it easier to compare the dynamics in the systems across the regimes, but also across regions. As in GMZ, we condition on the sign of the shock to retrieve the regime-specific impulse response functions. For confidence intervals, we use bias-corrected bootstrap confidence bands as proposed by Kilian (1998).

3.2 Results

In the baseline specification, we set $y_t = [\Delta\pi_t, \text{News Measure}_{t,r}]$, where News Measure is the quantitative or qualitative inflation news measure. $r \in$ [DE,FR] stands for the region and refers to whether the news measure is derived from the German-written media ("DE") or the French-written media ("FR"), respectively. Using the Schwarz Information Criterion (SIC) for lag selection, we obtain a value of $p = 2$ for the models with the quantitative and qualitative inflation news measure for both newspapers written in German and French.

In a first step, we focus on the results of the quantitative news measure and compare the results across the newspapers written in German versus French. In a second step, we repeat the analysis for the qualitative news measure.

Figure 3.1 summarizes the results for the quantitative inflation news measure. We discuss the panels from left to right.

Panel 3.1A and 3.1C plot the change in the inflation rate to a positive (red) and negative (blue) shock for the models estimated with the German and French-written newspaper measures, respectively. The impulse response functions to a positive or negative change in inflation are very similar in terms of size and persistence. This is in contrast to what GMZ find for unemployment where positive shocks to unemployment (that is, an increase in unemployment) are significantly more persistent than negative shocks. Naturally, the impulse response functions of inflation for the model estimated on German-written newspapers and those with French-written newspapers are similar as we use the same national inflation rate in the model.

Panel 3.1B and 3.1D show the impulse response functions of the quantitative news measure to a positive and negative one standard deviation inflation innovation (0.28 percentage points). We discuss two main observations.

First, following a positive shock in the inflation rate, the quantitative inflation news measure increases significantly. Similarly, following an unexpected decrease in the inflation rate, the quantitative news measure decreases significantly. This means that after an unexpected positive (negative) increase in the inflation rate, relatively more news about an inflation increase (decrease) are published.

Second, the quantitative inflation news measure reacts very similar to an increase or decrease in the inflation rate. If inflation media reporting were to be biased, that is, over-reporting of articles writing about an inflation increase versus decrease, the reaction of the quantitative inflation news measure would be asymmetrical. We get a similar picture if we focus on the cumulative response functions and calculate the differences in the respective cumulative impulse resopnse functions. The results are plotted in Figure C1 in the appendix for German-written newspapers and in Figure C2 for French-written newspapers.

However, as brought forward by GMZ, to correctly analyse whether media reporting is indeed (un-)biased, we need to look at the normalized impulse response functions and their differences. In more detail, reporting in newspapers may differ just because the shocks have different persistences. For the United States, GMZ show that after taking into account the potential non-linearity of the response of unemployment to an unemployment shock, the preliminary media bias of unemployment news reporting disappears. We calculate these multipliers and their difference as a robustness check and report them in the appendix in Figure C3. Overall, the dynamic multipliers confirm our previous results that we find no evidence of under- or over-reporting following an unexpected increase or decrease in the inflation rate for both newspapers written in German and French.

Besides checking for a negativtiy bias "within" German and French-written newspapers, we can also investigate whether we find diferences in news reporting across German and French newspapers. This allows us to analyse whether newspapers written in German or French report systematically different following an inflation shock. To do this, we take the normalized cumulative impulse response functions of the quantitative news measure of German-written articles that follows a positive unexpected shock in the inflation rate and calculate the difference to the same impulse response function of the quantitative news measure of French-written articles. Then, we do the same for a negative unexpected shock in the inflation rate. The results are displayed in Figure C5 in the appendix. We find no significant difference, neither for the positive nor the negative inflation shock.

These results are important. For both the quantitative news measure derived from German-written newspapers and from French-written newspapers, we find no evidence that they i) systematically over- or under-report after an unexpected change in the inflation rate (i.e., no evidence for a negativity bias of inflation) and ii), we find no evidence for systematically different news-reporting across newspapers written in French and German. Next, we repeat this analysis for the qualitative news measure.

We are interested in analysing how newspapers' sentiment about inflation reacts to an unexpected change in inflation. Furthermore, we investigate whether newspapers written in German report differently about these changes compared to the French-written newspapers. For brevity, we only report the results of cumulative impulse response functions for both newspapers written in German and French in the main text and let the reader refer to the Figure C4. in the appendix C for the normalised impulse response functions.

The results are displayed in Figure 3.2. For German-written newspapers, Panel 3.2A shows that the qualitative inflation news measure is more positive in case of a unexpected decrease in the inflation rate. For an unexpected increase in the inflation rate, the sentiment is more negative.

Panel 3.2B shows similar results for French-written newspapers. An unexpected increase (decrease) in inflation leads to more negative (positive) reports, even though the effect is not

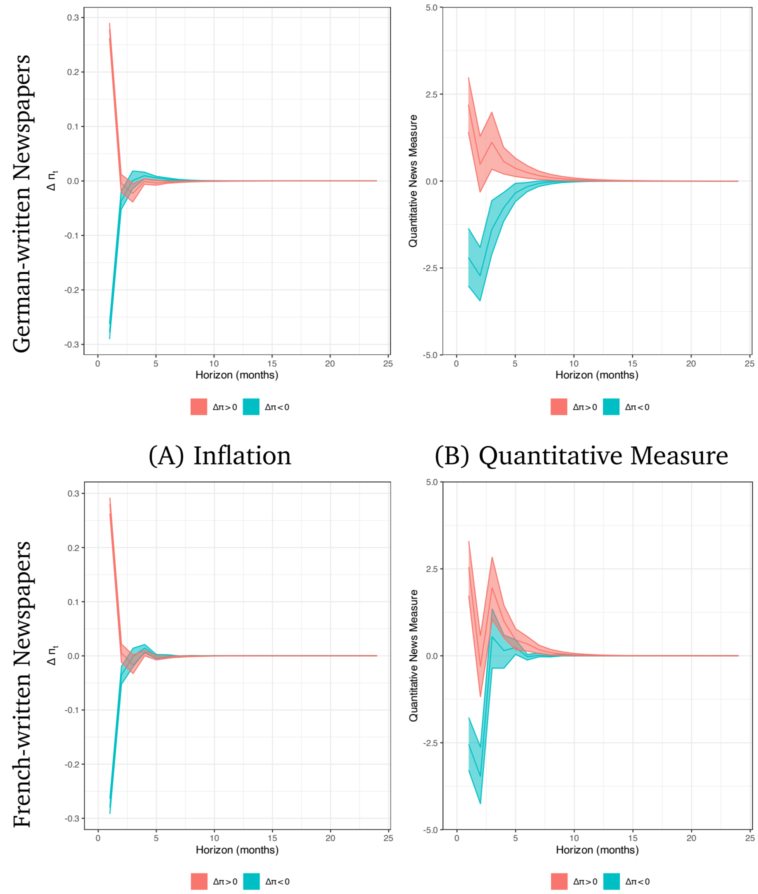

(A) Inflation (B) Quantitative Measure

(C) Inflation (D) Quantitative Measure

Figure 3.1: Results TSVAR for Quantitative News Measure

Note: The Figure shows the impulse response functions of the inflation rate and the quantitative news measure to a one standard deviation innovation in the inflation rate for both the German and French-written newspapers. Red coloured impulse response functions correspond to the response to a positive innovation in the inflation rate, and blue ones to a negative innovation in the inflation rate. The quantitative news measure is the difference of articles writing about an inflation increase minus the number of articles writing about an inflation decrease. The shadowed areas correspond to 68% confidence bands.

significant and less pronounced than for the German-written newspapers.

Finally, we test for differences in sentiment newspaper reporting across regions. In more detail, we analyse whether German-written newspapers report significantly more negative (positive) about an unexpected increase (decrease) in inflation. The results are displayed in the appendix Figure C6. In general, we find no strong differences in the reaction of inflation sentiment to an unexpected inflation shock across region. However, for an increase in inflation, sentiment tends to be slightly more negative for German-written newspapers than French-written newspapers for months 2 to 4, but remains insignificantly ever after.

In summary, the results in this section provide several insights. First, we find no evidence for a negativity bias in newspapers written in German and French in applying the novel approach of GMZ for inflation news reporting. Moreover, we find no significant difference in the news reporting across regional newspapers. Second, for the qualitative sentiment measure, we find that newspapers tend to report rather negatively about an inflation increase compared to a decrease, where this pattern is more pronounced for German-written newspapers. Broadly speaking, there are no strong differences in inflation sentiment reporting across regional newspapers. These results will be useful to help us interpreting the effects of news on inflation expectations and perceptions in the next section.

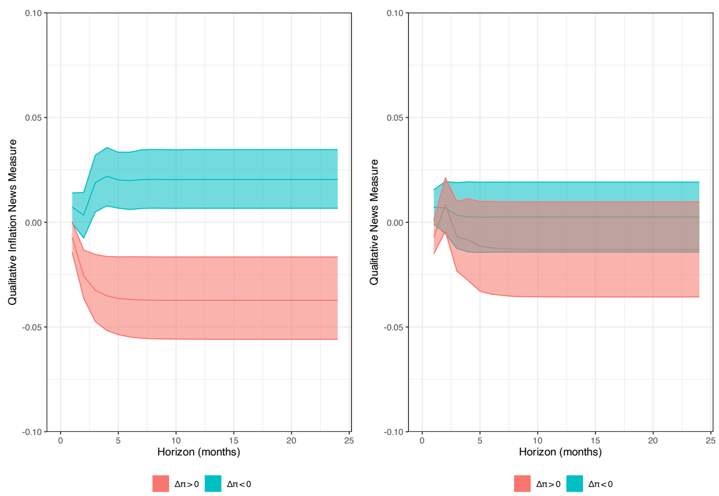

(A) German-written Newspapers

(B) French-written Newspapers

Figure 3.2: Cumulative Impulse Response Functions

Note: The Figure shows the cumulative impulse response functions of the TSVAR model described in section 3.1 from the qualitative sentiment news measure. Panel (a) shows the results for the qualitative sentiment measure derived from the German-written newspapers. Panel (b) shows the results for the French-written newspapers. The shadowed areas correspond to 68% confidence bands.

4. Inflation Expectations Data

In the next step, we analyse how news affects inflation expectations and perceptions. For this, we use data from the national Swiss Consumer survey (State Secretariat for Economic Affairs (SECO), 2022) conducted by the State Secretariat for Economic Affairs (SECO). This quarterly survey provides information of a repeated cross-section of representative households about their perceptions and expectations with respect to the economy, inflation, but also job security. The survey started with a sample size of approximately 500 households. Between the first quarter 1981 up until 2012 in the second quarter, the sample size increased to around 1100 households. From the third quarter 2012 up until now, the sample size encompasses roughly 1200 households.
In the survey, the households provide qualitative answers to the following questions about inflation perception and expectations:

- *Question:* How, in your view, have prices changed over the last 12 months? Have they:

- *Answers:* risen sharply; risen slightly; remained virtually unchanged; fallen slightly; fallen sharply; Don't know; No answer given.
- *Question:* How, in your view, will prices change over the next 12 months? Will they:
- *Answers:* rise sharply; rise slightly; remain virtually unchanged; fall slightly; fall sharply; Don't know; No answer given.

Importantly for us, the survey also includes socio-demographic information. The survey reports so-called WEMF regions. WEMF regions group communities into self-contained economic areas. This allows us to separately analyse the replies of households in the French-speaking part of Switzerland and the German-speaking part. However, in between 2012 and 2013 no information about the location is recorded. For this reason, we will discard this period from our sample.

While our main empirical analysis about the effects of news on expectations and perceptions focuses on individual household level data from the qualitative responses of the survey, we also quantify inflation expectations. The quantification of inflation expectations allows for a more direct interpretation of the size of the effect of news on expectations and serves as a robustness check.

To quantify inflation expectations from qualitative survey data, we follow Rosenblatt-Wisch and Scheufele (2015) who use a modified, more robust version of the widely used probability approach of Carlson and Parkin (1975). Rosenblatt-Wisch and Scheufele (2015) use the same survey data for Switzerland. All details about the approach are described in the appendix B.

Figures B2 and B3 in the appendix plot the shares of the replies to the question about inflation expectations and perceptions for households living in the German and French-speaking regions, respectively. Overall, the share of households with both high inflation expectations and perceptions is bigger for households in the French-speaking part than the German-speaking part.

Table 4.1 summarizes the mean shares of the qualitative survey replies for each of the two language regions. On average, 8.43% of the households expected prices to strongly increase compared to 5.32% of the households in the German-speaking part. Also for a small expected increase of prices, the share is higher for households in the French-speaking part with 50.47% compared to 45.60% in the German-speaking part. Inversely, more households in the German-speaking part expect prices to decrease (10.12%) compared to the French-speaking households (5.95%).

We find the same pattern for inflation perceptions. A higher share of French-speaking households perceive prices to have strongly increased (12.99%) or slightly increased (45.87%) in comparison with households in the German-speaking part (7.45% and 41.47%, respectively). Again, more households located in the German-speaking part perceived prices to have decreased (11.43%) compared to households in the French-speaking part (6.10%).

Table 4.1: Summary Table for Regional Inflation Expectations and Perceived Inflation

Variable	Reply	French Region	German Region
Inflation Expectations	Strong Increase	8.43	5.32
	Slight Increase	50.47	45.60
	Constant	35.15	38.96
	Decrease	5.95	10.12
Inflation Perceptions	Strong Increase	12.99	7.45
	Slight Increase	45.87	41.47
	Constant	35.04	39.64
	Decrease	6.10	11.43

Note: The table shows the shares of replies for each category with respect to expected inflation and perceived inflation from the qualitative survey data for German and French-speaking households. The shares are in percentages and calculated over the entire sample period.

For brevity, the results for the quantified average inflation expectations are discussed and summarized in the appendix B.3.

5. News Media and Inflation Expectations and Perceptions

5.1 Model

In this section, we discuss the approach to derive the quantitative and qualitative inflation media shocks, following GMZ, and how we use them to investigate the effects on inflation expectations and perceptions. To analyse the effect of news on expectations and perceptions, we need changes in the news measures that are unrelated to the dynamics of the inflation rate. Otherwise, variations of expectations and perceptions due to the media shocks could simply reflect a stronger increase or decrease in the inflation rate.

For this reason, we make use of the series of shocks from our model in section 3.1. In more detail, the shocks v_{2t} represent variations in the news measure that are unrelated to the (current and past) inflation rate.

Arguably, households in the French-speaking part of Switzerland consume more news written in French whereas households in the German-speaking part consume more news in German. We exploit this language barrier in our empirical set-up. In more detail, it allows us to control for time-fixed effects that capture confounds that affect the media shock as well as the perceptions and expectations at the same time.[9] For example, we control for a national shock that increases news coverage but also expectations and perceptions at the same time. The

[9] In effect, we exploit the regional variation in the shock series v_{2t}. These series of shocks are plotted in Figure C7 in the appendix.

main linear probability model used can be written as follows.

$$\text{reply}_{i,r,t,q} = \alpha + \sum_{j=1}^{2} \beta_j \text{QMS}_{r,t+1-j} + \sum_{j=1}^{2} \tilde{\beta}_j \text{SMS}_{r,t+1-j} + \phi \text{HH}_{i,r,t} + \mu \text{Region}_r + \gamma_t + \varepsilon_{i,r,t} \tag{5.1}$$

The subscript i refers to household-level observations, subscript r refers to the region (German versus French-speaking region). We analyse quantitative media shocks (QMS) and qualitative sentiment media shocks (SMS) at time t and their lag $t-1$. QMS and SMS correspond to the series of shocks v_{2t} identified from the model estimated in section 3.1 using the quantitative and qualitative media news measures in the specification of y_t, respectively. $\text{HH}_{i,r,t}$ are household level controls that include gender and age, Region_r is an indicator of whether the household is located in the French or German-speaking part of Switzerland, and γ_t are time fixed-effects.[10] $\text{reply}_{i,r,t,q}$ is a dummy equal to 0 if the household expects or perceives prices to decrease or have decreased, respectively, and 100 if the household expects or perceives prices to increase or have increased, with subscript q indicating whether the answer refers to the question about expectations (e) or perceptions (p). All media shocks are scaled by their respective standard deviation, such that a one unit increase corresponds to a one standard-deviation increase. For all regressions, we cluster the errors at the date$\times$region level.

5.2 Results

In Table 5.1 we show the effects of a quantitative and qualitative media shock on the share of households expecting or perceiving an increase in prices versus a decrease in prices. In Columns (1) to (4), we estimate the regression from equation 5.1 with quantitative and qualitative media shocks separately, where in columns (5) and (6), all media shocks and their lags are included. Note that an increase in the quantitative shock corresponds to relatively more news articles writing about an inflation increase compared to decrease. An increase in the qualitative media shock corresponds to a more positive inflation sentiment.

For the quantitative media shock, in column (1), a one standard deviation increase in the quantitative media shock leads to a 1.33 percentage point increase in the share of households indicating an increase in price expectations. Relative to the mean share, this corresponds to an increase of 1.6 percent. This effect is small but significant at the 5% level. We find no significant effect of the lag of the quantitative media shock. Column (2) shows the effect of the quantitative media shock on inflation perceptions. The coefficients have positive signs but

[10] Note that we also conduct robustness checks where we include additional controls such as the job and education level. However, these types of information are less populated which results in a significant lower number of total observations. As the results remain similar after controlling for these additional variables, we prefer to stick with a higher number of observations to retain a representative sample of households.

are insignificant.
For the qualitative media shock at time t and $t-1$, in column (3), we find no significant effect on the share of households indicating an increase in price expectations versus those that indicate a decrease. However, in column (4) we find a significant effect on price perceptions. A one standard deviation increase in the qualitative media shock increases price perceptions by 1.09 percentage points. Again, this effect is small but similar in magnitude to the effects of a quantitative shock on expectations.
In Columns (5) and (6), we introduce both media shocks and their lags in the regression, as the quantitative and qualitative media shocks might be correlated. However, the results are largely unchanged and remain similar in their magnitude as well as their significance.

Table 5.1: Effect of Media Shocks on Inflation Expectations and Perceptions

	(1) reply_e	(2) reply_p	(3) reply_e	(4) reply_p	(5) reply_e	(6) reply_p
Quantitative Media Shock_t	1.33**	0.79			1.34**	0.85
	(0.64)	(0.96)			(0.62)	(0.95)
Quantitative Media Shock_{t-1}	0.23	0.43			0.24	0.47
	(0.68)	(0.92)			(0.68)	(0.90)
Qualitative Media Shock_t			0.22	1.09**	0.19	1.08**
			(0.49)	(0.50)	(0.48)	(0.48)
Qualitative Media Shock_{t-1}			-0.07	0.29	0.08	0.36
			(0.50)	(0.46)	(0.49)	(0.49)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Region_r	Yes	Yes	Yes	Yes	Yes	Yes
$\text{HH}_{i,r,t}$	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,447	34,747	32,447	34,747	32,447	34,747
$\bar{y}$	85	86	85	86	85	86

Note: The table shows the effects of quantitative and qualitative media shocks derived in section 3.1 on the share of households expectations (reply_e) and perceptions (reply_p) that indicate an increase versus a decrease. Household controls $\text{HH}_{i,r,t}$ include gender and age fixed effects. The media shocks are standardized such that a one unit increase corresponds to a one standard-deviation increase in the media shock. $\bar{y}$ corresponds to the mean of the dependant variable. * p<0.10, ** p<0.05, *** p<0.010. Standard Errors are clustered at the date × region level.

Next, we investigate whether expectations and perceptions react differently conditional on whether the change in inflation is positive or negative. This analysis is motivated by the findings in section 3. In periods of a positive (negative) change in inflation, the quantitative news measure tends to increase (decrease), which corresponds to relatively more (less) news about an inflation increase compared to news about a decrease. Furthermore, the sentiment reported during a period of inflation increase is rather negative compared to a positive sentiment for an inflation decrease. For this, we estimate the model conditional on times when inflation increases $\Delta\pi_t > 0$ and times when inflation decreases $\Delta\pi_t \leq 0$ (or stays constant).
The results are displayed in table 5.2. We first focus on the results in columns (1) and (2) where the model is estimated conditional on times where $\Delta\pi_t > 0$. With respect to expecta-

tions, column (1) shows that relatively more news about an inflation increase than decrease have a positive effect on inflation expectations. In contrast, a more positive sentiment with respect to inflation during this period has a decreasing effect on inflation expectations. For perceptions, we find that the lag of the quantitative media shock can have a significant positive effect on perceptions but find no effect of the qualitative news sentiment.

Columns (3) and (4) focus on periods where $\Delta\pi_t \leq 0$. While the quantitative media shock has no significant effect during these periods, the qualitative sentiment shock has a positive effect on both expectations and perceptions. A one standard deviation increase in the qualitative media shock at time t increases expectations and perceptions by 1.7 and 3.3 percentage points, respectively.

Table 5.2 underlines two observations. First, the quantitative news measure has the highest effects when the change in inflation is positive. As seen in section 3, this effect is not driven by systematically different news reporting during these periods. We find that the qualitative sentiment measure of inflation has a significant effect during all periods. Moreover, the alternating sign of the effect of sentiment point towards a potential counter-cyclical effect of sentiment on expectations and perceptions.

The quantitatively rather small effects of news on expectations and perceptions are in line with the findings of Dräger (2015) about news reporting and their effect on expectations and perceptions in Sweden. Also, sticky expectations of households could be another explanation for the rather small effects of news. As pointed out by Rosenblatt-Wisch and Scheufele (2015), who use the same survey data for Switzerland, households only update their expectations once a year.

Finally, we investigate whether the effects of news on expectations and perceptions vary conditional on socio-demographic characteristics.

In Figure 5.1, we plot the interaction effects of the media shocks with different age categories. We split the sample into households aged between 15 and 40 (age category 1), 40 to 60 (age category 2), and households aged 60 and older (age category 3). We then interact the age category dummy with the contemporaneous and lagged media shocks. The reference group in the regressions are the youngest households aged between 15 and 40. The results are displayed in Figure 5.1.

For expectations, we find some evidence that the interaction effect of news on the share of households indicating an inflation increase versus decrease is generally more positive for households in older age categories compared to the households in the youngest age category. For the quantitative media shock, it is especially the lagged media shock that shows a higher effect for older age categories. For sentiment media shocks, we find that it is more positive for the contemporaneous media shock compared to the lagged shock.

For perceptions, however, we find only weak evidence of a higher effect of media shocks

Table 5.2: Effect of Quantitative and Qualitative Media Shocks conditional on Change in Inflation

	$\Delta\pi > 0$		$\Delta\pi \leq 0$	
	(1) reply$_e$	(2) reply$_p$	(3) reply$_e$	(4) reply$_p$
Quantitative Media Shock$_t$	1.41**	1.64	1.17	0.38
	(0.69)	(1.41)	(1.05)	(1.10)
Quantitative Media Shock$_{t-1}$	1.36*	2.36**	-2.11	-0.41
	(0.70)	(1.01)	(1.40)	(0.58)
Qualitative Media Shock$_t$	-0.87**	-0.14	1.71	3.30***
	(0.36)	(0.44)	(1.10)	(0.77)
Qualitative Media Shock$_{t-1}$	0.10	0.42	-0.79	0.03
	(0.50)	(0.60)	(1.33)	(1.01)
Date FE	Yes	Yes	Yes	Yes
Region$_r$	Yes	Yes	Yes	Yes
HH$_{i,r,t}$	Yes	Yes	Yes	Yes
Observations	17,379	18,545	15,064	16,200
$\bar{y}$	89	89	81	83

Note: The table shows the effects of quantitative and qualitative media shocks derived in section 3.1 on the share of households expectations (reply$_e$) and perceptions (reply$_p$) that indicate an increase versus a decrease. The models are estimated conditional on a positive or negative change in inflation. Household controls HH$_{i,r,t}$ include gender and age fixed effects. The media shocks are standardized such that a one unit increase corresponds to a one standard-deviation increase in the media shock. $\bar{y}$ corresponds to the mean of the dependant variable. * p<0.10, ** p<0.05, *** p<0.010. Standard Errors are clustered at the date $\times$ region level.

across age groups. The effect is only significantly positive for the lagged quantitative media shock. However, most coefficients tend to be slightly positive.

Such higher effects of news on expectations and perceptions with increasing age may be related to differences in news consumption among younger and elderly respondents. In various studies, age has been found to be an important driver of newspaper consumption (see, for example, Lauf (2001), or Elvestad and Blekesaune (2008)), even after taking into account online articles of newspapers (Thurman and Fletcher, 2019).

An alternative explanation for the higher effect of elderly households is that they might pay more attention to inflation news as they have a higher exposure through their savings (Doepke and Schneider, 2006) or due to differences in their relative shares of food and housing on the consumption basked (Basso et al., 2023).

In Table 5.3, we analyse whether the effects of news differ across the language border. More precisely, we interact the media shocks with the Region$_r$ variable. For households living in the German-speaking region, most of the coefficients for the both the quantitative and qualitative media shocks are positive and significant. In contrast to that, we find that for households in the French-speaking region, news tends to have a lower effect on expectations and perceptions. Note that the higher effect of news on expectations and perceptions for households in the German-speaking part remains robust when estimating the models conditional on periods when inflation increases or decreases.

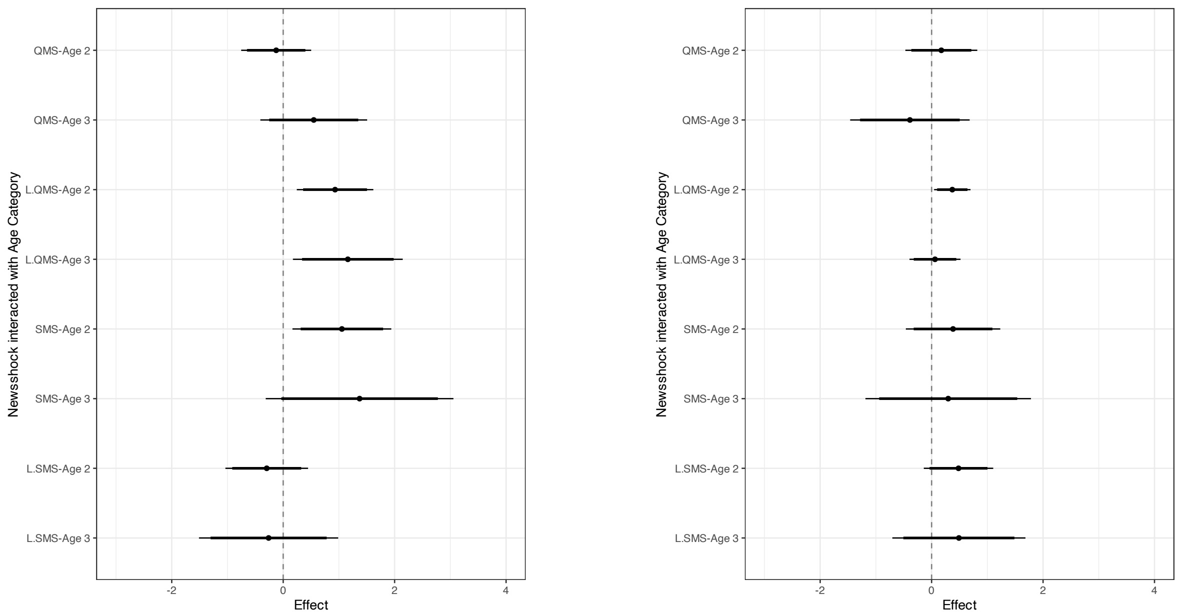

(A) Expectations

(B) Perceptions

Figure 5.1: Media shocks interacted with Age Categories

Note: The Figure shows the marginal effects of the media shock interacted with the Age Categories. The marginal effects are effects relative to the reference group, which are households aged between 15 and 40. The second age category are households aged between 40 and 60. The third age category are households aged 60 and above. Panel (a) shows the results for the interactions of the quantitative media shock (QMS), its lag (L.QMS) and the qualitative sentiment shock (SMS) and its lag (L.SMS) on the share of households expecting an inflation increase versus decrease. Panel (b) the same interaction effects but for inflation perceptions. The Figure shows 5% and 10% confidence intervals.

A potential explanation for a higher effect of inflation media shocks on households in the German-speaking part might be that those households are more receptive to inflation news than households in the French-speaking part. From the results in section B.3, we can rule out that the effect of news is purely driven by systematically different news reporting. Moreover, Jost (2018) has documented that households in the German-speaking part show a significantly higher inflation aversion than the households in the French-speaking part. This, in turn, could mean that households in the German-speaking part are more receptive to inflation news, prompting a stronger reaction.

Alternatively, it is possible that the differences observed across the language border might be driven by different shopping experiences both across Swiss regions but also across the border, or that households draw information on different inflation rates across region from official government statistics. However, recent evidence from Kluser (2023) shows no support for regional price discrimination of one of the largest Swiss retailers. In addition, Figure B5 shows that inflation rates in Germany are very similar to those in France.[11] Figure B4 shows a similar pattern for regional inflation rates of the three biggest cities in Switzerland, Zürich,

[11]Exceptions are the period in between 2003 to 2004 and after 2021. However, note that our sample excludes observations in between 2003 and 2004 due to a reporting issue of the Swissdox database during the period.

Geneva and Basel. It is therefore unlikely that these explanations are the main drivers of the different effect we observe across the language border.

In summary, exploiting regional variation in inflation news reporting, we find that both the quantitative media shocks and qualitative sentiment shocks have a significant effect on inflation expectations and perceptions. While relatively more news about an inflation increase compared to news about an inflation decrease has the strongest positive effect during periods of positive inflation, sentiment has a countercyclical effect on expectations and perceptions during times when inflation is positive or negative. Furthermore, we find some evidence that the effect of news is stronger for elderly households. Finally, we document differences in the reaction to news across the language border, where we find households located in the German-speaking part to be more reactive to news compared to households in the French-speaking part.

Table 5.3: Effect of Quantitative and Qualitative Media Shocks conditional on Region

	(1) reply_e	(2) reply_p
German-speaking Region		
Quantitative Media Shock$_t$	1.35*	1.06
	(0.69)	(1.04)
Quantitative Media Shock$_{t-1}$	0.69	1.36
	(0.65)	(0.82)
Qualitative Media Shock$_t$	1.25*	1.14**
	(0.74)	(0.53)
Qualitative Media Shock$_{t-1}$	-0.41	-0.56
	(0.71)	(0.67)
French-speaking Region		
French-speaking Region=1 $\times$ Quantitative Media Shock$_t$	-0.30	-1.20**
	(0.51)	(0.47)
French-speaking Region=1 $\times$ Quantitative Media Shock$_{t-1}$	-1.50***	-1.86***
	(0.55)	(0.58)
French-speaking Region=1 $\times$ Qualitative Media Shock$_t$	-2.40**	-0.52
	(1.10)	(0.90)
French-speaking Region=1 $\times$ Qualitative Media Shock$_{t-1}$	0.55	-0.90
	(1.17)	(0.83)
Date FE	Yes	Yes
Region$_r$	Yes	Yes
HH$_{i,r,t}$	Yes	Yes
Observations	32,447	34,747
$\bar{y}$	85	86

Note: The table shows the effects of quantitative and qualitative media shocks derived in section 3.1 interacted with a region dummy on the share of households expectations (reply_e) and perceptions (reply_p) that indicate an increase versus a decrease. The reference group are households located in the German-speaking part. Household controls HH$_{i,r,t}$ include gender and age fixed effects. The media shocks are standardized such that a one unit increase corresponds to a one standard-deviation increase in the media shock. $\bar{y}$ corresponds to the mean of the dependant variable. * p<0.10, ** p<0.05, *** p<0.010. Standard Errors are clustered at the date $\times$ region level.

6. Robustness Checks

In this section, we provide several robustness checks. First, we estimate the main results including a business cycle indicator, ordered after the inflation rate but before the quantitative or qualitative news measure. Arguably, inflation is a lagged variable that follows business cycle activity. For this reason, our shock identified in section 5.1 might be confounded by other shocks, following from the current economic conditions.[12] The results are displayed in appendix D.1. Overall, our results remain unchanged with respect to both newspaper reporting as well as for the effects of media shocks on inflation expectations and perceptions. Second, we add the stock prices growth rate from the Swiss Market Index to the baseline model specification in section 5.1. We order this variable at the last position, such that our shock series takes into account this forward-looking variable. We report the results in appendix D.1. In general, our results remain robust to this alternative specification.

Third, we use the quantitative inflation expectations derived from the qualitative survey replies to study the effects of news on expectations. The results indicate that the qualitative survey responses about inflation expectations are similar to the findings obtained from quantitative analysis, although the latter are less significant due to the limited number of observations. In general, it appears that a positive quantitative media shock can have a positive effect on expectations, and the effects of qualitative media shocks are consistent with those discussed in section 5.2. One notable exception is that we did not observe a similar negative effect of expectations from sentiment shocks during periods of positive inflation growth. However, it is worth noting that all media shocks had a relatively smaller impact on households in the French-speaking region. A detailed breakdown of these results is available in Table D.3 in the appendix.

7. Conclusion

In this paper, we use a novel data set of newspaper articles in Switzerland to shed light on the links between regional news reporting and households' inflation expectations and perceptions. We create a standard quantitative inflation news measure that reports the differences in the number of articles published about an inflation increase and those that report an inflation decrease. Then, we develop a novel qualitative inflation news measure. Using a state-of-the-art NLP model, we conduct a sentiment analysis on the news articles to predict whether inflation is positively or negatively assessed.

We proceed by investigating how newspapers in Switzerland report about inflation using a threshold SVAR. For the quantitative news measure, we check for the existence of a negativity bias, i.e. newspapers over-reporting during times when inflation increases versus times when inflation decreases. We find no evidence for a negativity bias in newspapers written in French nor in German. Moreover, we find no across-regional differences in quantitative news reporting. For the qualitative news sentiment measure, we find that newspapers assess inflation in

[12] GMZ conduct a similar robustness check where they use industrial production. However, due to data limitations, we use a business cycle indicator as an alternative to industrial production.

general positively during times of negative inflation and assess it rather negative during times of positive inflation. Again, we find that these differences do not differ significantly across newspapers written in French and German.

Deriving media shock for both the quantitative and qualitative news measure, we exploit a quasi-experimental setting in Switzerland to study the effects of news on expectations and perceptions. In this set-up, we find that both the quantitative media shocks and qualitative sentiment shocks have a significant effect on inflation expectations and perceptions. While relatively more news about an inflation increase compared to news about an inflation decrease has the strongest positive effect during periods of positive inflation, sentiment has a countercyclical effect on expectations and perceptions during times when inflation is positive or negative. Furthermore, we find some evidence that the effect of news is stronger for elderly households. Finally, we document differences in the reaction to news across the language border, where we find households located in the German-speaking part to be more reactive to news compared to households in the French-speaking part. While these observed differences across the language border cannot be explained by systematic differences in news reporting across the regions, a potential explanation is that households in the German-speaking part are more receptive to news as they have been found to be more inflation averse.

Overall, our findings are policy relevant for several reasons. As news coverage affects inflation expectations, real effects can be the consequence. Well-anchored inflation expectations improve the effectiveness of monetary policy. In our sample, 21% of the articles in the German-written newspapers and 15% of those in French-written newspapers write about the Swiss National Bank. This leaves a potential role for central bank communication which has been shown to be picked up by news media and affecting households' inflation expectations (Hirsch et al., 2023). If central bank communication can affect the sentiment, i.e. how newspapers assess inflation, it may be used as a channel to affect expectations and perceptions in line with the price stability target. In addition, our results suggest that central banks should exploit various communication channels to reach households of all age categories (and regions). Finally, our novel measure of inflation sentiment may be used as a policy indicator that summarizes how newspapers assess inflation.

We think that several directions of further research seem to be worth following. First of all, in this paper, we focused on printed newspaper articles. An in-depth analysis of online articles (and social platforms), changing media consumption habits and its effects on households' expectations and perceptions would contribute to our understanding of news coverage and its effects on households. Another question that we have left aside in this paper is whether different narratives might lead to asymmetric effects in households' expectations. Promising methods in topic modelling, as in Müller et al. (2022), provide a fruitful starting point for this question. Finally, analysing the linkages of media reporting to other expectations and different agents in the economy is a valuable topic for future research.

Bibliography

Angeletos, G.-M., Huo, Z., and Sastry, K. A. (2020). *Imperfect Macroeconomic Expectations: Evidence and Theory*, pages 1–86. University of Chicago Press.

Araci, D. (2019). Finbert: Financial sentiment analysis with pre-trained language models. *arXiv preprint arXiv:1908.10063*.

Bannier, C., Pauls, T., and Walter, A. (2019). Content analysis of business communication: introducing a german dictionary. *Journal of Business Economics*, 89(1):79–123.

Basel-Stadt (2022). Preise.

Basso, H. S., Dimakou, O., and Pidkuyko, M. (2023). How inflation varies across spanish households. *Banco de Espana Ocassional Paper*, (2307).

Berk, J. M. (1999). Measuring inflation expectations: a survey data approach. *Applied Economics*, 31(11):1467–1480.

Beyer, A., Gaspar, V., Gerberding, C., and Issing, O. (2008). Opting out of the great inflation: German monetary policy after the break down of bretton woods. Technical report, National Bureau of Economic Research.

Blinder, A. S. and Krueger, A. B. (2004). What does the public know about economic policy, and how does it know it?

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2020). Overreaction in macroeconomic expectations. *American Economic Review*, 110(9):2748–82.

Carlson, J. A. and Parkin, M. (1975). Inflation expectations. *Economica*, 42(166):123–138.

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.

Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–78.

Cooley, T. F. and Prescott, E. C. (1976). Estimation in the presence of stochastic parameter variation. *Econometrica: journal of the Econometric Society*, pages 167–184.

Cukierman, A. (1992). Central bank strategy, credibility, and independance: Theory and evidence. *Journal des Économistes et des Études Humaines*, 3(4):581–590.

Deepset (2019). German BERT.
https://huggingface.co/bert-base-german-cased.

Devlin, J., Chang, M., Lee, K., and Toutanova, K. (2018a). BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805.

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

Doepke, M. and Schneider, M. (2006). Inflation and the redistribution of nominal wealth. *Journal of Political Economy*, 114(6):1069–1097.

Doms, M. E. and Morin, N. J. (2004). Consumer sentiment, the economy, and the news media. *FRB of San Francisco Working Paper*, (2004-09).

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.

D'Acunto, F., Malmendier, U., Ospina, J., and Weber, M. (2021). Exposure to grocery prices and inflation expectations. *Journal of Political Economy*, 129(5):1615–1639.

Ehrmann, M. and Tzamourani, P. (2012). Memories of high inflation. *European Journal of Political Economy*, 28(2):174–191.

Elvestad, E. and Blekesaune, A. (2008). Newspaper readers in europe: A multilevel study of individual and national differences. *European Journal of Communication*, 23(4):425–447.

Eugster, B., Lalive, R., Steinhauer, A., and Zweimüller, J. (2017). Culture, work attitudes, and job search: Evidence from the swiss language border. *Journal of the European Economic Association*, 15(5):1056–1100.

Federal Office of Communications (2016). Bericht zur Überprüfung der Definition und der Leistungen des Service public der SRG unter Berücksichtigung der privaten elektronischen Medien. Accessed in July 2022, retrieved from https://www.uvek.admin.ch/dam/uvek/de/dokumente/kommunikation/faktenblatt-aktuelle-medienthemen.pdf.download.pdf/faktenblatt-aktuelle-medienthemen.pdf.

Federal Statistical Office (2022a). Printmedien: Auflage- und Leserzahlen ausgewählter Tages- und Sonntagszeitungen, nach Sprachregionen. Accessed in July 2022, retrieved from https://www.bfs.admin.ch/bfs/en/home/statistics/catalogues-databases.assetdetail.je-d-16.03.01.01.06.html.

Federal Statistical Office (2022b). Sprachen. Accessed in July 2022, retrieved from https://www.bfs.admin.ch/bfs/de/home/statistiken/bevoelkerung/sprachen-religionen/sprachen.html.

Fogarty, B. J. (2005). Determining economic news coverage. *International Journal of Public Opinion Research*, 17(2):149–172.

Frommberg, L. (2014). How falling prices hurt the economy. https://www.20min.ch/story/so-schaden-sinkende-preise-der-wirtschaft-956510657822.

Gambetti, L., Maffei-Faccioli, N., and Zoi, I. S. (2023). Bad news, good news: Coverage and response asymmetries. *Finance and Economics Discussion Series 2023-001*.

Goidel, R. K. and Langley, R. E. (1995). Media coverage of the economy and aggregate economic evaluations: Uncovering evidence of indirect media effects. *Political Research Quarterly*, 48(2):313–328.

Hayo, B. (1998). Inflation culture, central bank independence and price stability. *European Journal of Political Economy*, 14(2):241–263.

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.

Issing, O. (2005). Why did the great inflation not happen in germany? *Federal Reserve Bank of St. Louis Review*, 87(2 Part 2):329–35.

Jost, A. (2018). Cultural differences in monetary policy preferences. Technical Report 2, Swiss National Bank.

Kilian, L. (1998). Small-sample confidence intervals for impulse response functions. *Review of economics and statistics*, 80(2):218–230.

Kluser, F. (2023). Cross-border shopping: Evidence from household transaction records.

Kohlhas, A. and Broer, T. (2019). Forecaster (Mis-)Behavior. 2019 Meeting Papers 1171, Society for Economic Dynamics.

Kumar, S., Afrouzi, H., Coibion, O., and Gorodnichenko, Y. (2015). Inflation targeting does not anchor inflation expectations: Evidence from firms in new zealand. Technical report, National Bureau of Economic Research.

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.

Larsen, V. H., Thorsrud, L. A., and Zhulanova, J. (2021). News-driven inflation expectations and information rigidities. *Journal of Monetary Economics*, 117:507–520.

Lauf, E. (2001). Research note: The vanishing young reader: Sociodemographic determinants of newspaper use as a source of political information in europe, 1980-98. *European journal of communication*, 16(2):233–243.

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

Lee, A. M. and Delli Carpini, M. (2010). News consumption revisited: Examining the power of habits in the 21st century. Paper presented at the11th International Symposium on Online Journalism, 23–24 April, University of Texas: Austin. Retrieved from https://isoj.org/wp-content/uploads/2018/01/LeeCarpini10-1.pdf, accessed 12 January 2023.

Lee, H. S. (2020). Exploring the initial impact of covid-19 sentiment on us stock market using big data. *Sustainability*, 12(16):6648.

Lolić, I. and Sorić, P. (2018). A critical re-examination of the carlson–parkin method. *Applied Economics Letters*, 25(19):1360–1363.

Loughran, T. and McDonald, B. (2011). When is a liability not a liability? textual analysis, dictionaries, and 10-ks. *The Journal of finance*, 66(1):35–65.

Mankiw, N. G., Reis, R., and Wolfers, J. (2003). Disagreement about inflation expectations. *NBER macroeconomics annual*, 18:209–248.

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. Technical report, DoCMA Working Paper.

Nardo, M. (2003). The quantification of qualitative survey data: a critical assessment. *Journal of economic surveys*, 17(5):645–668.

Nautz, D. and Strohsal, T. (2015). Are us inflation expectations re-anchored? *Economics Letters*, 127:6–9.

Nemes, L. and Kiss, A. (2021). Prediction of stock values changes using sentiment analysis of stock news headlines. *Journal of Information and Telecommunication*, 5(3):375–394.

Organization for Economic Co-operation and Development (2023a). Consumer Price Index: All Items: Total: Total for France [cpaltt01frm659n]. retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/CPALTT01FRM659N.

Organization for Economic Co-operation and Development (2023b). Consumer Price Index: All Items: Total: Total for Germany [cpaltt01dem659n]. retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/CPALTT01DEM659N.

Organization for Economic Co-operation and Development (2023c). Consumer Price Index: Harmonized Prices: All Items: Total for France [cp0000frm086nest]. retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/CP0000FRM086NEST.

Organization for Economic Co-operation and Development (2023d). Consumer Price Index: Harmonized Prices: All Items: Total for Germany [cphptt01dem659n]. retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/CPHPTT01DEM659N.

République et Canton de Genève (2022). Prix à la consommation.

Rosenblatt-Wisch, R. and Scheufele, R. (2015). Quantification and characteristics of household inflation expectations in switzerland. *Applied Economics*, 47(26):2699–2716.

Seitz, H. (1988). The estimation of inflation forecasts from business survey data. *Applied Economics*, 20(4):427–438.

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

Soroka, S. N. (2006). Good news and bad news: Asymmetric responses to economic information. *The Journal of Politics*, 68(2):372–385.

Soroka, S. N. (2012). The gatekeeping function: Distributions of information in media and the real world. *The Journal of Politics*, 74(2):514–528.

Sousa, M. G., Sakiyama, K., de Souza Rodrigues, L., Moraes, P. H., Fernandes, E. R., and Matsubara, E. T. (2019). Bert for stock market sentiment analysis. In *2019 IEEE 31st International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 1597–1601. IEEE.

Stadt Zürich (2022). Konsumentenpreisindex.

Starr, M. A. (2012). Consumption, sentiment, and economic news. *Economic Inquiry*, 50(4):1097–1111.

State Secretariat for Economic Affairs (SECO) (2022). Swiss consumer survey. https://www.seco.admin.ch/seco/en/home/wirtschaftslage---wirtschaftspolitik/Wirtschaftslage/Konsumentenstimmung.html.

Swissdox (2022). Swissdox.ch - Medienbeobachtung. Accessed in July 2022, retrieved from https://swissdox.ch/.

Thurman, N. and Fletcher, R. (2019). Has digital distribution rejuvenated readership? revisiting the age demographics of newspaper consumption. *Journalism Studies*, 20(4):542–562.

Weber, M., D'Acunto, F., Gorodnichenko, Y., and Coibion, O. (2022). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. Technical report, National Bureau of Economic Research.

Appendix

A. Inflation News Indices

A.1 Description

To construct indices about inflation news, we focus on printed articles that write about the Swiss economy. For newspapers written in German, we select the *Tages Anzeiger, Berner Zeitung, Blick* and *20 Minuten*. For French-written newspapers, we select *Le Temps, 24 heures, Tribune de Genève, 20 minutes* and *Le Matin*. These newspapers are among the largest in terms of readership (Federal Statistical Office, 2022a). Note that we do not consider the *Neue Zürcher Zeitung* (NZZ). Swissdox classifies the NZZ as a national newspaper. As the goal of this paper is to analys regional newspaper reporting of inflation, we dismiss articles from the NZZ.

We analyse articles from January 1998 to April 2022. For all articles, we look for the words inflation, a synonym of inflation and price ("Inflation", "Teuerung" and "Preis" in German, "inflation", "renchérissement" and "prix" in French). Similar to GMZ, we search for words indicating an increase or decrease in the vicinity of inflation. The words that identify an increase or decrease are translated and described in table A1. We use similar words for both languages.

Table A1: Words indicating decrease or increase in Inflation

Indicator	Words
Increase	Increase∗, pop+upward, spike, augment∗, markup, boom∗, boost, growth, grow∗, increment∗, drive∗, high∗, soar∗, more expensive, accelerate∗
Decrease	Decrease∗, decline∗, drop∗, dampen∗, reduce∗, reduction, fall∗, dip∗, downward∗ , low∗, abate∗, shrink∗, go + down, plunge∗, attenuate∗

We count newsarticles that report an increase in inflation if inflation or its synonym appears within a five-word distance of an indicator for increase, and without an indicator of a decrease in a two word distance. Symetrically, we count newsarticles that report a decrease in inflation when inflation or its synonym appears within a five-word distance of an indicator for decrease and no indicator of an increase within a two word distance. In contrast to GMZ, we use a two word instead of a one-word distance. Especially in German, an increase (decrease) is often indicated with the combination of two words, i.e. "nehmen" + "zu" ("nehmen" + "ab").

With these articles at hand, we construct a "quantitative inflation news" measure. This measure is defined as the difference between the number of articles writing about an increase in inflation minus the number of articles that writing about a decrease in inflation. We then aggregate these measures on a monthly frequency.

To get an overview of the articles considered, Figure A1 shows a wordcloud for German and French-written articles. These wordclouds represent the keywords found in the newspaper articles weighted by the term frequency-inverse document frequency (tf-idf) score. The tf-idf score is based on how often a term appears in an article (term frequency) and how rare it is in the entire corpus (inverse document frequency).

The tf-idf score is defined as the product of the term-frequency and the inverse-document frequency. The term frequency is defined as $tf(w,a) = \frac{N_{w,a}}{\sum_{w' \in a, Nw',a}}$, where $N_{w,a}$ stands for the absolute frequency of word w in article a. $\sum_{w' \in a, Nw',a}$ is the total number of words in article a. The inverse document frequency is defined as $idf(w,a) = \log \frac{N}{|\{a \in A : w \in a\}|}$, where N is the total number of articles in the corpus and $|\{a \in A : w \in a\}|$ indicates the number of articles where the word w appears.

In Panel A1A, we plot the 50 words with the highest tf-idf score for newspapers written in German. The larger a word is plotted, the higher its tf-idf score. Reassuringly, all words displayed are economic-related terms. Furthermore, terms describing the Swiss National Bank (national bank, or SNB), related to the interest rate and the exchange rate are among the words with the highest tf-idf score.

In Panel A1B, we plot the wordcloud of French-written articles. Again, the Swiss National Bank (short SNB) and words related to the interest rate and national currency are important keywords. Overall, the wordcloud for French-written articles feature similar words. However, the weights attached to these words vary conditional on the newspapers. For example, the word "swiss_frank" has a lower score in the German-written newspapers compared to the French-written newspapers. Therefore, "swiss_frank" is less relevant in terms of the tf-idf score in the German-written newspapers compared to the French-written newspapers.

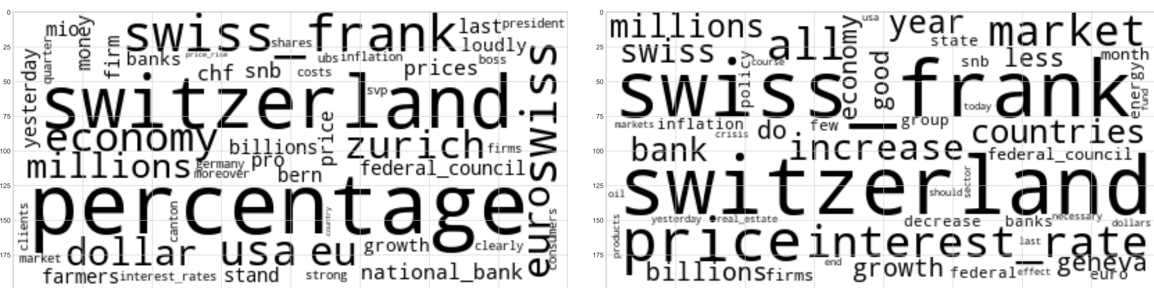

(A) Wordcloud German-written Articles

(B) Wordcloud French-written Articles

Figure A1: Wordclouds

Note: The Figure shows wordclouds built with all the articles considered in the main analysis for both newspapers written in German and newspapers written in French. The words are weighted by the term frequency-inverse document frequency (tf-idf). The tf-idf takes into account not only the frequency of each word, but also its relative importance in the context of the entire corpus. In both figures, words with higher tf-idf score are displayed with a larger font size.

A.2 Qualitative News Measure and GDP Growth

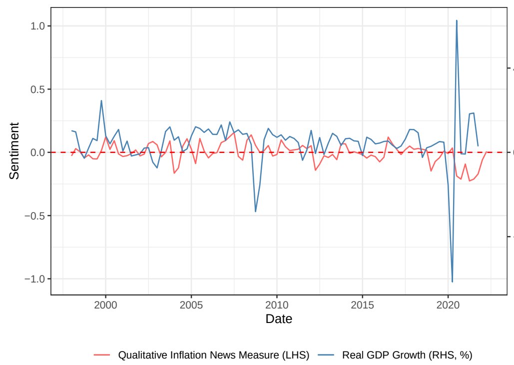

(A) Newspapers written in German

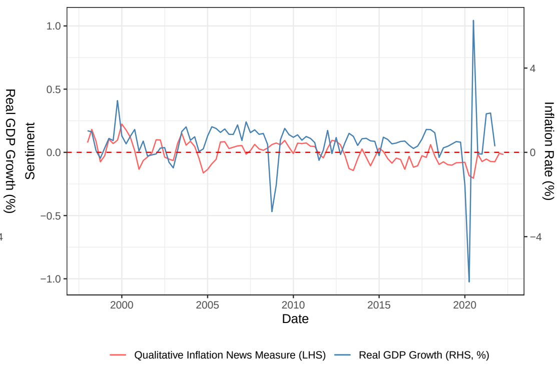

(B) Newspapers written in French

Figure A2: Qualitative Inflation News Measure and Real GDP Growth

Note: In the left panel, the figure plots the qualitative inflation news measure for newspapers written in German (left-hand scale) together with the real GDP growth rate (right-hand scale). The qualitative news measure is the average sentiment resulting from the prediction of the BERT model. In the right panel, we plot the qualitative news measure for newspapers written in French (left-hand scale) together with the real GDP growth rate (right-hand scale). The qualitative inflation news measure is standardized. Values above (below) the horizontal zero line indicate periods of a predominantly positive (negative) sentiment with respect to inflation. All data is averaged to quarterly frequency.

B. Inflation Expectations

In this section, we will discuss the probability approach of Carlson and Parkin (1975) in more detail. Originally, this approach, often called the Carlson-Parkin method (CP method henceforth), was developed for only three categorical answers with respect to price developments but can easily be extended to more categories. The method is widely used (see Nardo (2003) for a review), but has been criticized in its original form recently Lolić and Sorić (2018). Especially the strong assumption of unbiased inflation expectations has been questioned. For this reason, we will use the modified approach from Rosenblatt-Wisch and Scheufele (2015) which circumvents the assumption of unbiased inflation expectations.

To start with, the probability approach assumes that individuals form their inflation expectations from a subjective probability distribution $f_i(\pi_{i,t+4})$, characterized by mean $\mathbb{E}_t[\pi_{i,t+4}]$ and standard deviation $\sigma_t(\pi_{i,t+4})$. This probability distribution function is the same across all agents. Further, the method assumes that the individuals reply to the question whether prices will *go down* if $\mathbb{E}_t[\pi_{i,t+4}] \leq -\delta_{it}^L$, *stay the same* if $-\delta_{it}^L < \mathbb{E}_t[\pi_{i,t+4}] \leq \delta_{it}^U$, *will rise moderately* if $\delta_{it}^U < \mathbb{E}_t[\pi_{i,t+4}] \leq \lambda_{it}$ and *will rise strongly* if $\mathbb{E}_t[\pi_{i,t+4}] > \lambda_{it}$. Therefore, in between the range of $-\delta_{it}^L$ to δ_{it}^U, the individual does not notice that prices either increase or decrease.

Given that all individuals share the same probability distribution function, we can describe the aggregate probability distribution using the shares of the survey replies for each category.

$$
\begin{aligned}
&P(\mathbb{E}_t[\pi_{i,t+4}] \geq -\delta_{it}^L) = A_t \\
&P(\mathbb{E}_t[\pi_{i,t+4}] \geq \delta_{it}^U) - P(\mathbb{E}_t[\pi_{i,t+4}] > -\delta_{it}^L) = B_t \\
&P(\mathbb{E}_t[\pi_{i,t+4}] \geq \lambda_{it}) - P(\mathbb{E}_t[\pi_{i,t+4}] > \delta_{it}^U) = C_t
\end{aligned}
$$

We define the abscissae a_t, b_t, c_t of the distribution function to correspond to the cumulative probabilities of $A_t, A_t + B_t, A_t + B_t + C_t$, respectively. Following Rosenblatt-Wisch and Scheufele (2015), we choose the normal distribution as the distribution function, as they find that alternatives have only minor effects on the results.

Finally, assuming that the interval in-between individuals do not notice any differences in prices is symmetrical around 0, $-\delta_t^L = \delta_t^U$, we can define $\mathbb{E}_t[\pi_{i,t+4}]$ and $\sigma_t(\pi_{i,t+4})$ in terms of the parameter λ_t and the quantiles of the distribution.

$$
\mathbb{E}_t[\pi_{t+4}] = \frac{\lambda_t(a_t + b_t)}{(a_t + b_t - 2c_t)} \tag{B.1}
$$

$$
\sigma_t(\pi_{t+4}) = \frac{-2\lambda_t}{(a_t + b_t - 2c_t)} \tag{B.2}
$$

$$
\delta_t = \frac{\lambda_t(a_t - b_t)}{(a_t + b_t - 2c_t)} \tag{B.3}
$$

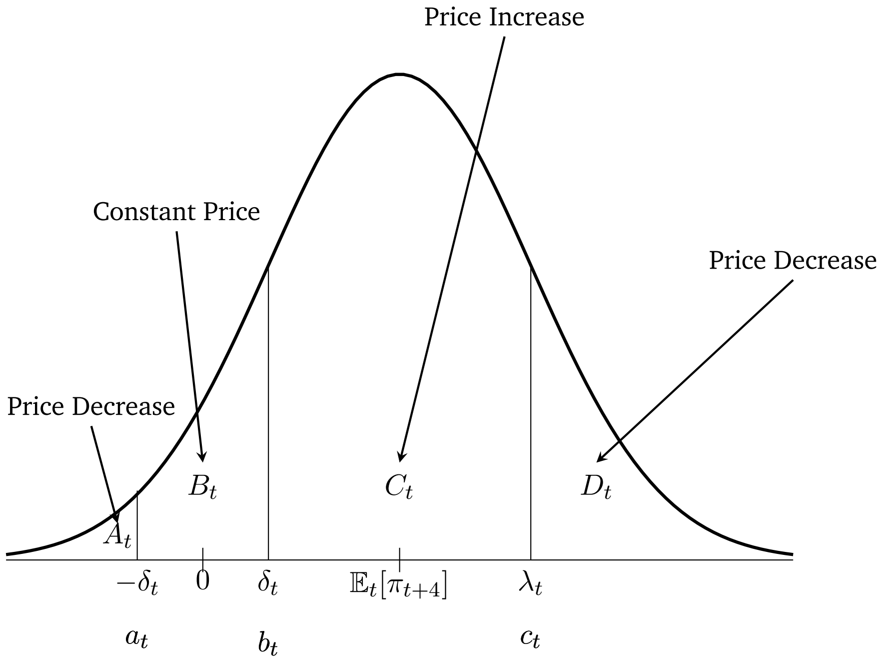

Figure B1: Illustration CP method - Joint Probability Distribution

Figure B1 provides an illustration.

There exist different methods to estimate the parameter λ_t. The CP method assumes that the parameter is time-invariant and inflation expectations are unbiased, such that $\mathbb{E}_t[\pi_{i,t+4}] = \pi_{i,t+4}$ and the parameter can be easily estimated using the observed inflation rate. These rigorous assumptions can be relaxed in several ways. First, Rosenblatt-Wisch and Scheufele (2015) propose a time-variant scaling parameter and second, they circumvent the assumption of unbiased inflation expectations by using information about the perceived inflation rate.

Rosenblatt-Wisch and Scheufele (2015) assume that individuals have, on average, a correct perception of prices (compare also Berk (1999)). This is arguably a less strong assumption than unbiased inflation expectations. Following the same logic as for inflation expectations using the survey shares for inflation perceptions, $\hat{\pi}_t$, we can write

$$\hat{\pi}_t = \frac{\lambda_t(a'_t + b'_t)}{(a'_t + b'_t - 2c'_t)}$$

with a'_t, b'_t, c'_t being the abscissae of the distribution function of perceived inflation. Now, assuming that $\hat{\pi}_t = \pi_t$ on average, we can estimate the parameter λ_t.

In this paper, we focus on the estimation of λ_t using a state-space model and present robustness checks using rolling regressions as an alternative. In Rosenblatt-Wisch and Scheufele (2015), both methods have led to reasonable results. First, we describe the state-space model and its estimation. Second, we briefly summarize the rolling regression method.

B.1 State-Space Model

The State-Space Model was originally proposed by Seitz (1988). It consists of one measurement equation (B.4) and a transition equation (B.5).

$$\pi_t = \lambda_t \frac{(a'_t + b'_t)}{(a'_t + b'_t - 2c'_t)} + u_t \tag{B.4}$$

$$\lambda_t = \lambda_{t-1} + v_t \tag{B.5}$$

with $\mathbb{V}ar(u_t) = (1-\gamma)\sigma^2$ and $\mathbb{V}ar(v_t) = \gamma\sigma^2$. This is a simple Kalman Filter set-up. To proceed, however, we need initial estimates of the variance parameters σ^2 and γ. We follow the approach of Cooley and Prescott (1976) using a constrained maximum likelihood function.

B.2 Rolling Regressions

Another approach to estimate λ_t is using rolling regressions. For each window, we run the regression

$$\pi_t = \lambda \frac{(a'_t + b'_t)}{(a'_t + b'_t - 2c'_t)} + u_t$$

. Using a window of 30 quarters, λ_t is defined as

$$\lambda_t^r = \frac{\sum_{k=t-w+1}^{t} (a'_k + b'_k)/(a'_k + b'_k - 2c'_k)\pi_k}{\sum_{k=t-w+1}^{t} ((a'_k + b'_k)/(a'_k + b'_k - 2c'_k))^2}$$

The choice of the window leads to similar results as in Rosenblatt-Wisch and Scheufele (2015). To calculate a_t, b_t and c_t, we use the mean shares the survey replies for each category with respect to price expectations and price perceptions. Figure B2 plots the qualitative survey replies for inflation expectations over the relevant sample period from the second quarter of 1997 up until the second quarter 2022 conditional on the region. In general, the shares over time follow a similar evolution. However, the share of households that expect inflation to increase (strongly or moderate) is, on average, higher for the French region.

Figure B3 plots the share for each category with respect to inflation perceptions. Similarly to inflation expectations, the share of households that perceive inflation to be high or moderate is higher compared to the German region.

There may be several reasons why we observe average differences in the shares of survey replies. First, it is possible that the different perceptions and expectations mirror different inflation rates in these regions. Second, it might be possible that expectations and perceptions might be influenced by shopping experience across the border. Third, the implicit rates for which households' perceive and expect inflation to belong to either category differ across regions.

For this reason, we plot in Figure B4 cantonal inflation indices.[13]. In Switzerland, cantonal inflation indices are only available for three cantons, namely Basel, Zürich and Geneva. These cities are also the three biggest in terms of population. Except for the very beginning of the sample period around 1998, inflation rates across the cantons are very similar and follow closely the national Swiss inflation rate.

Figure B5 plots the inflation rate for France and Germany. We show both the national inflation measures as well as the Eurostat's harmonized index of consumer prices (HICP). Again, the inflation rates are very similar across these two countries.

Therefore, the observed difference in the shares of perceived and expected inflation is most likely due to different thresholds in the rates at which households perceive (and expect) inflation to change.

B.3 Descriptive Results

Figure B6 plots the resulting inflation expectations using both methods, the Kalman Filter and the Rolling Regression approach for each region, respectively. Overall, the inflation expectations are similar with some exceptions at the very beginning of the sample and shortly after 2015.

In Figure B7 panel B7A, we plot the average inflation expectations $\mathbb{E}[\pi_{t+4}]$ derived from the qualitative survey conditional on each region together with the actual observed inflation rate π_{t+4}. Inflation expectations differ across regions but overall share a very similar trend in the French and German-speaking regions. Only after 2015, inflation expectations are higher for the German-speaking regions. Overall, inflation expectations seem well anchored, with exceptions during the financial crisis and most recently following the pandemic.

While households in the French-speaking part show higher shares of inflation perceptions and expectations in the categories of strong and slight increase, average inflation expectations are lower for households in the French-speaking part. This is due to the different thresholds at which these households perceive inflation to be decreasing, constant, slightly increasing or strongly increasing. As, by assumption, inflation perceptions are unbiased on average, this means that households in the French-speaking part perceive inflation to be slightly increasing or increasing at a lower rate than German households do. For this reason, adjusting inflation expectations for the scaling parameters results in lower inflation expectations for households

[13]Sources are linked in the references (Basel-Stadt, 2022; République et Canton de Genève, 2022; Stadt Zürich, 2022)

Figure B2: Qualitative Inflation Expectations

Note: The Figure shows the share of replies for the survey question about inflation expectations over the next 12 months conditional on the region. Note that from 2012 to 2013, there is no information about the region of the households available.

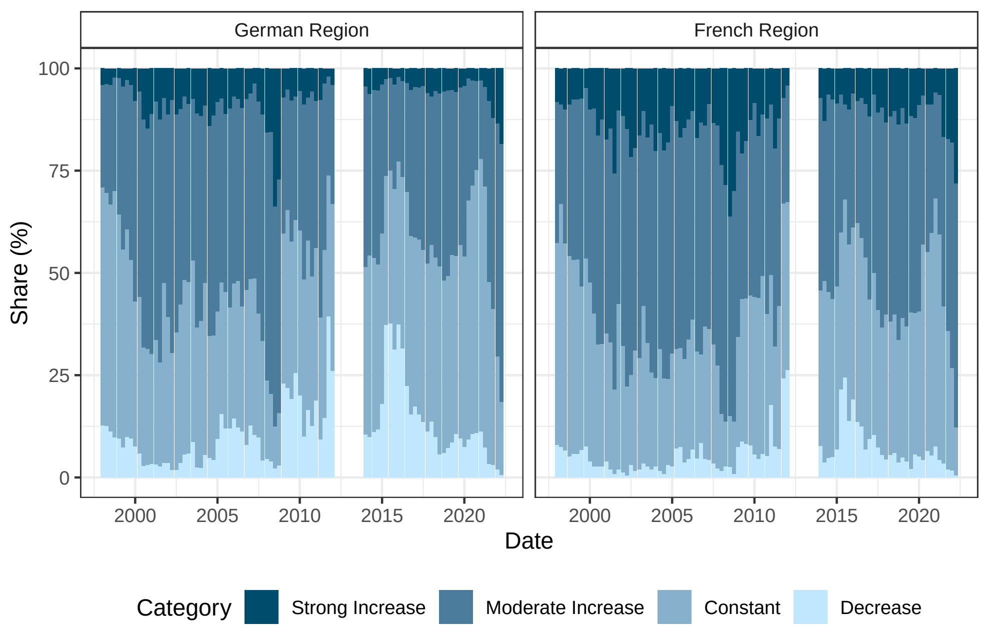

Figure B3: Qualitative Inflation Perception

Note: The Figure shows the share of replies for the survey question about inflation perceptions from the past 12 months conditional on the region. Note that from 2012 to 2013, there is no information about the region of the households available.

in the French-speaking part.

In panel B7B, we plot the standard deviation of the inflation expectations conditional on the

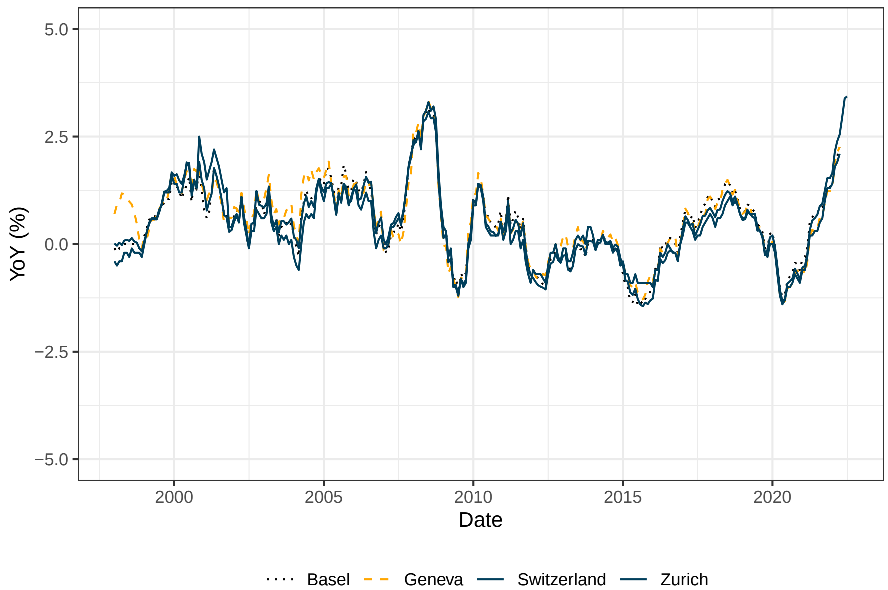

Figure B4: Cantonal Inflation Rates

Note: This Figure plots three different cantonal inflation indices, namely from Basel, Geneva and Zurich, together with the average inflation index of Switzerland.

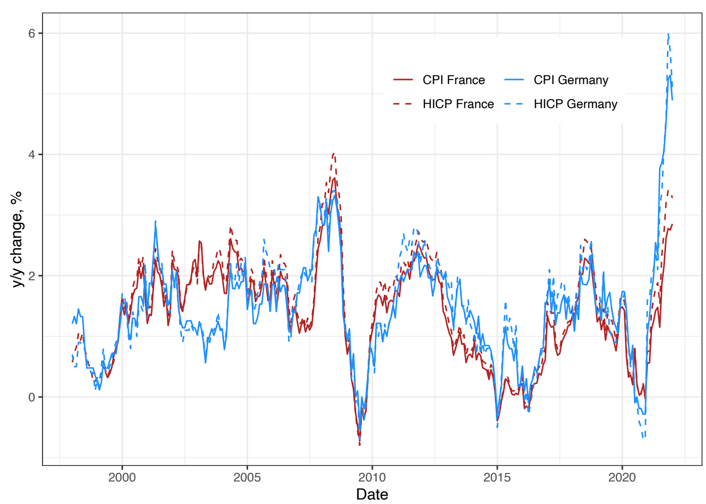

Figure B5: Inflation Rates for Germany and France

Note: The figure shows the consumer price indices and harmonized consumer price indices for Germany and France. Sources: Organization for Economic Co-operation and Development (2023a,c,b,d)

regions. In general, the inflation expectations in the French region are less volatile than in the German region. Similar to the mean inflation expectations, the standard deviation is higher during the financial crisis and increasing since the pandemic.

Table B1 summarizes the main moments of the observed inflation rate and households' inflation expectations over the entire sample period. The average observed inflation of 0.54 is quite close to expectations of the German-speaking region with 0.63 and French-speaking

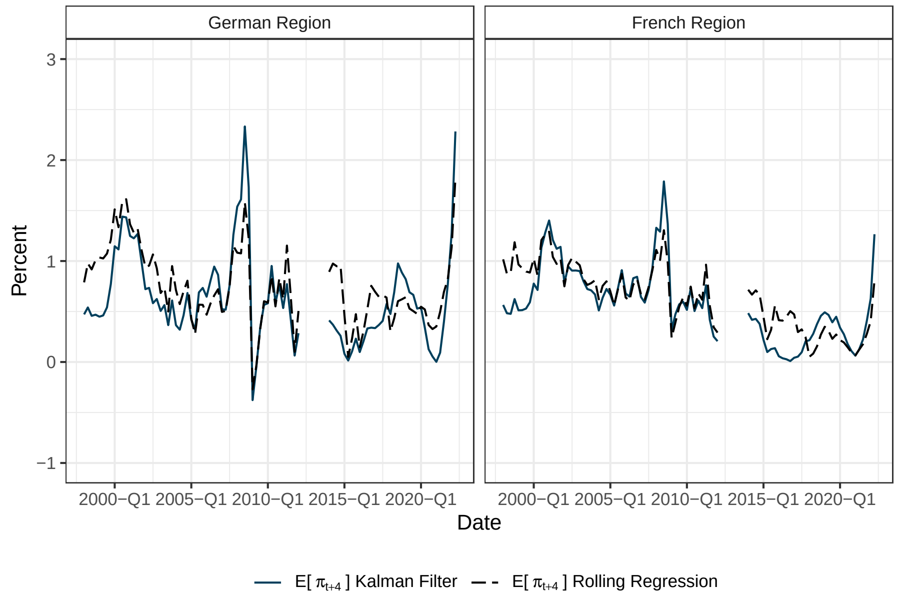

Figure B6: Inflation Expectations with Kalman Filter and Rolling Regressions

Note: The Figure shows the mean inflation expectations for households using two different estimation techniques for the parameter λ_t. We compare the results from the Kalman Filter estimation process described in B.1 with the Rolling Regression approach described in B.2.

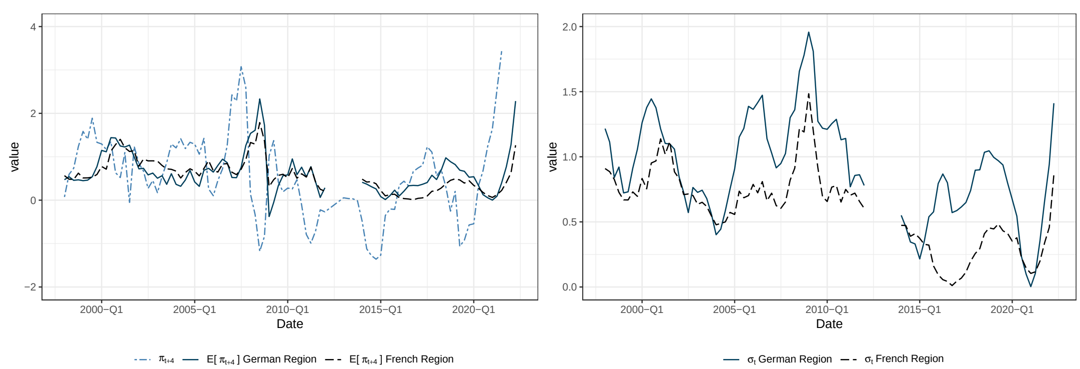

(A) Mean

(B) Standard Deviation

Figure B7: Mean and Standard Deviation of $\mathbb{E}[\pi_{t+4}]$ conditional on the Region

Note: The left panel shows the realized inflation rate (change comapred to the same month in the previous year) and the mean inflation expectations for changes in prices in the next 12 months for households in both regions, the German and French-speaking part (equation B.1). The right panel shows the standard deviation of the inflation expectations calculated as in equation B.2.

region with 0.58. While observed inflation reaches a maximum at 1.79 in the third quarter 2021, inflation expectations for both regions reach a local maximum in the third quarter of 2008, with 2.33 for the German region and 3.43 for the French-speaking region.

Table B1: Summary Table for Regional Inflation Expectations and Observed Inflation

Variable	Mean	Sd.	Min.	Max.
Observed Inflation	0.54	0.97	-1.36	3.43
Inflation Expectations German-speaking Region	0.63	0.47	-0.38	2.33
Inflation Expectations French-speaking Region	0.58	0.37	0.01	1.79

Note: The table shows summary statistics of observed inflation and the quantified inflation expectations derived from qualitative survey data for German and French-speaking households.

C. Additional Results Inflation Reporting of Newspapers

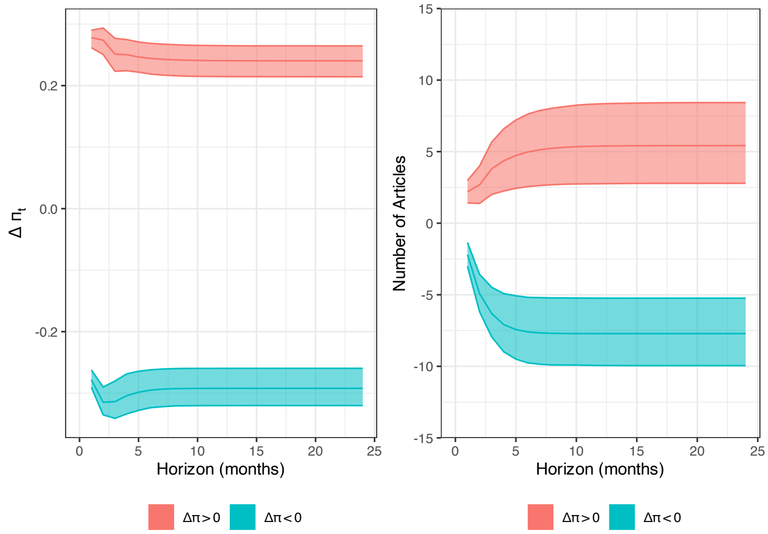

Figure C1: Results TSVAR for Quantitative News Measure in German Newspapers

Note: The Figure shows the cumulative impulse response functions of the TSVAR model described in section 3.1. The left-hand panel shows the cumulative sum of the responses of the change in the inflation rate to a positive (red) and negative (blue) inflation rate innovation, respectively. The right-hand side panel shows the cumulative sum of the responses of the quantitative inflation news measure (number of articles writing about an inflation increase minus number of articles writing about an inflation decrease) to the positive (inflation increases) and negative inflation rate innovation. The shadowed areas correspond to 68% confidence bands.

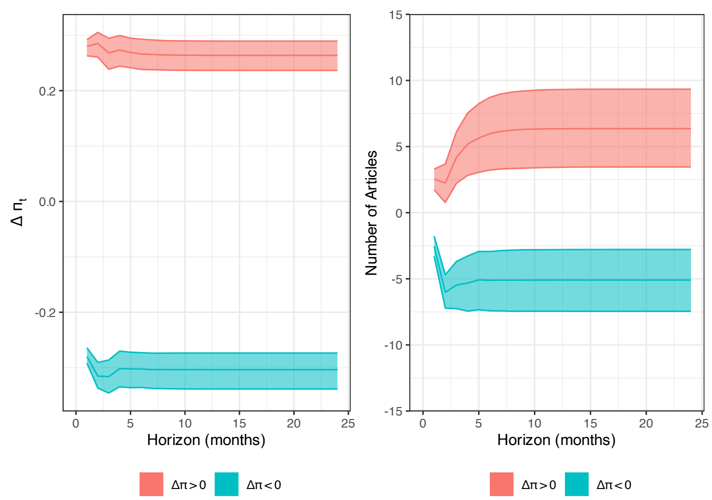

Figure C2: Results TSVAR for Quantitative News Measure in French Newspapers

Note: The Figure shows the cumulative impulse response functions of the TSVAR model described in section 3.1. The left-hand panel shows the cumulative sum of the responses of the change in the inflation rate to a positive (red) and negative (blue) inflation rate innovation, respectively. The right-hand side panel shows the cumulative sum of the responses of the quantitative inflation news measure (number of articles writing about an inflation increase minus number of articles writing about an inflation decrease) to the positive (inflation increases) and negative inflation rate innovation. The shadowed areas correspond to 68% confidence bands.

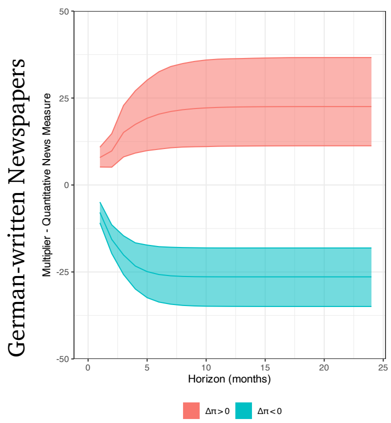

(A) Dynamic Multiplier

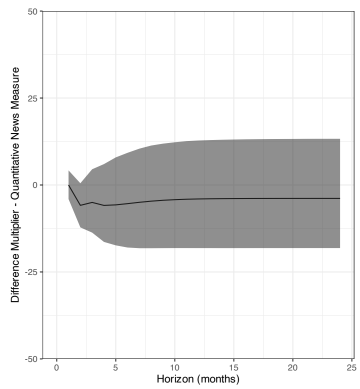

(B) Multiplier Difference

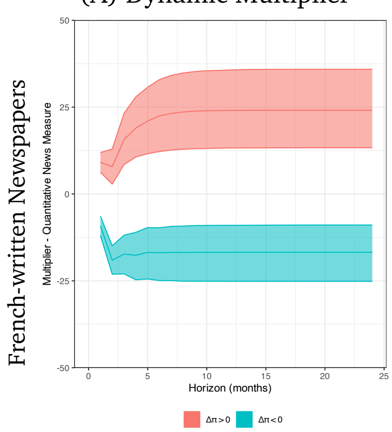

(C) Dynamic Multiplier

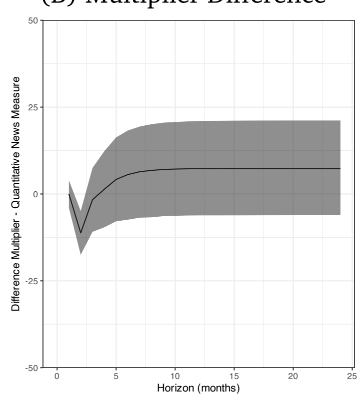

(D) Multiplier Difference

Figure C3: Dynamic Multipliers for Quantitative News Measure

Note: The Figure shows the dynamic multipliers of the TSVAR model described in section 3.1 for the qualitative news measure in Panel (a) and (c) and the multiplier difference in Panel (b) and (d). The dynamic media multiplier is a normalization of the impulse response function where, at every time t, the cumulative response of the qualitative news measure is normalized by the cumulative response of the inflation rate. The dynamic media multiplier can be interpreted as follows. At the end of the time horizon of our impulse response functions at 24 months, the dynamic media multiplier shows how many excess articles about an inflation increase compared to inflation decrease articles are generated by the news media in response to a 1 percentage point shock in inflation. The difference of the dynamic media multiplier in Panel (b) and (d) is calculated between states where changes in inflation are positive versus negative. A positive difference of the dynamic multiplier points towards a negativity bias, where newspapers would report more frequently about an inflation increase compared to a decrease. In Panels (a) and (b), the results for newspapers written in German is displayed, and in Panels (c) and (d) the results for the French-written newspapers. The shadowed areas correspond to 68% confidence bands.

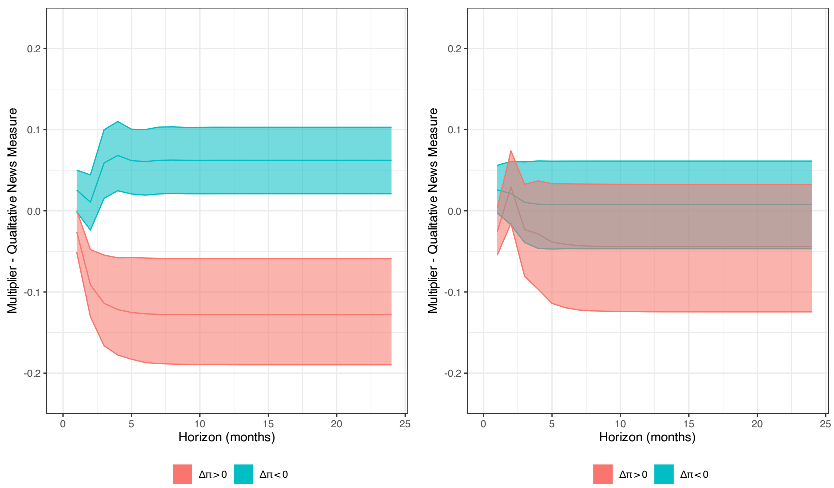

(A) German-written Newspapers (B) French-written Newspapers

Figure C4: Dynamic Multipliers for Qualtitative News Measure

Note: The Figure shows the dynamic multipliers of the TSVAR model described in section 3.1 for the qualitative news measure. The dynamic media multiplier is a normalization of the impulse response function where, at every time t, the cumulative response of the qualitative news measure is normalized by the cumulative response of the inflation rate. Panel (a) shows the results for the qualitative news measure derived from the German-written newspapers. Panel (b) shows the results for the French-written newspapers. The shadowed areas correspond to 68% confidence bands.

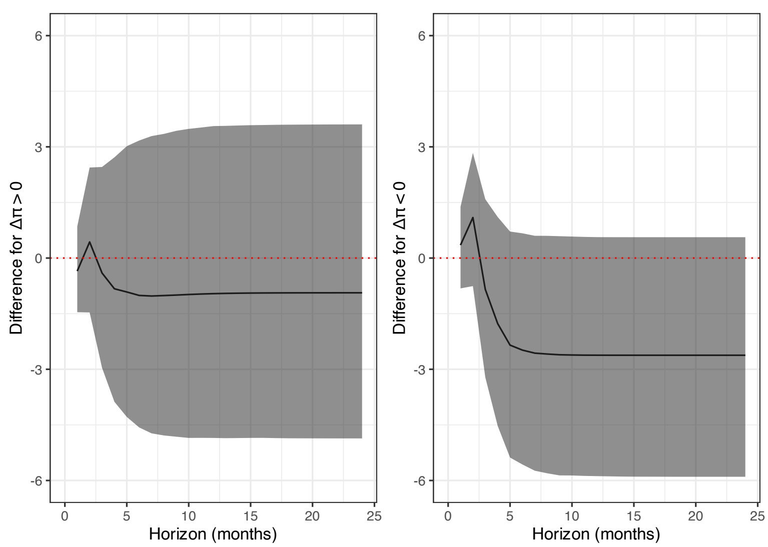

Figure C5: Differences in Quantitative News Measure Across Region

Note: On the left-hand side panel, the figure shows the difference in the quantitative news measure reaction to an unexpected inflation increase across region. Namely, we calculate the difference between the reaction of the quantitative news measure from the German newspaper and the French-written newspaper. On the right-hand side panel, the figure shows the difference in the quantitative news measure reaction to an unexpected inflation decrease across region. The shadowed areas correspond to 68% confidence bands.

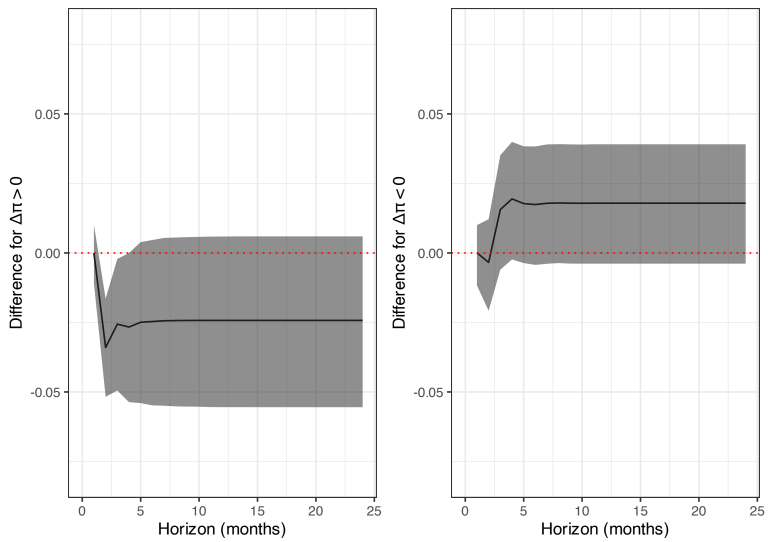

Figure C6: Differences in Qualitative News Measure Across Region

Note: On the left-hand side panel, the figure shows the difference in the qualitative sentiment measure reaction to an unexpected inflation increase across region. Namely, we calculate the difference between the reaction of the qualitative news measure from the German newspaper and the French-written newspaper. On the right-hand side panel, the figure shows the difference in the quantitative news measure reaction to an unexpected inflation decrease across region. The shadowed areas correspond to 68% confidence bands.

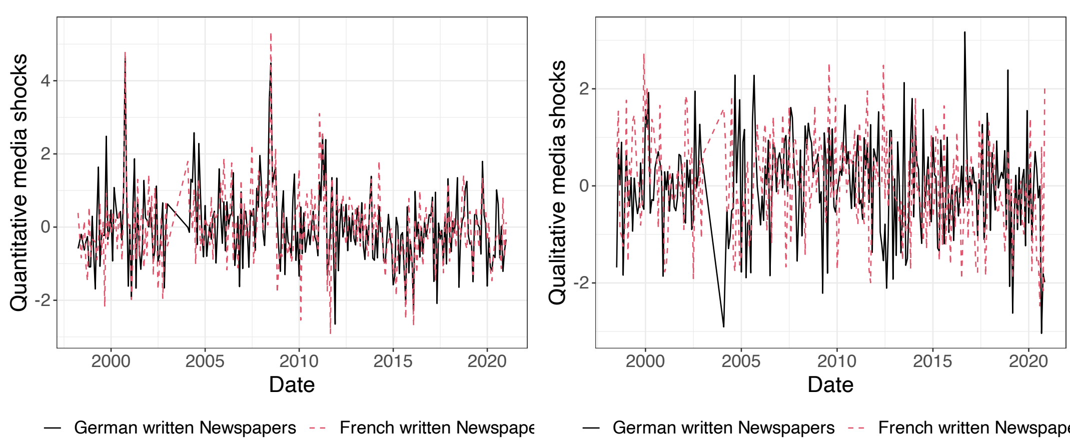

(A) Quantitative News Measure

(B) Qualitative News Measure

Figure C7: Media shocks

Note: The Figure shows the media shocks described in section 3.1 for the quantitative and qualitative news measure, conditional on the region for each month. Note that the gap in between the period of January 2003 to December 2003 is due to a reporting issue in the Swissdox database.

D. Robustness Checks

D.1 Inclusion Business Cycle Indicator

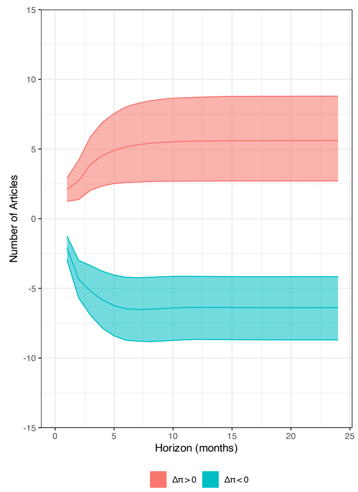

(A) DE: CIR

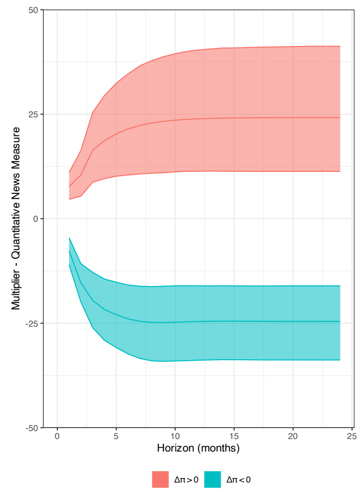

(B) DE: Normalized CIR

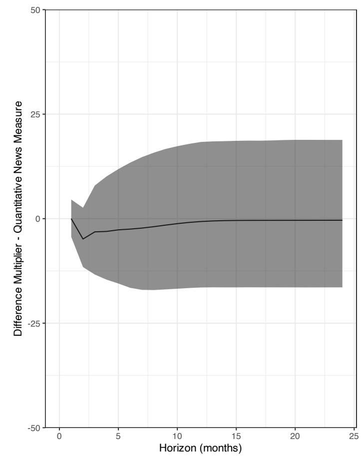

(C) DE: Difference Multiplier

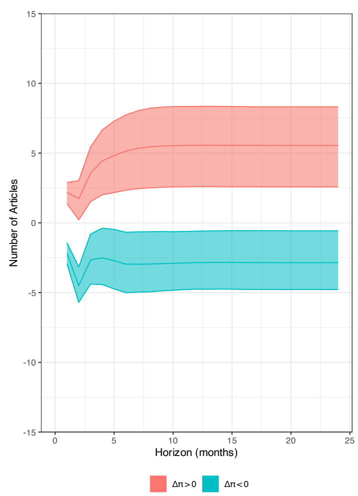

(D) FR: CIR

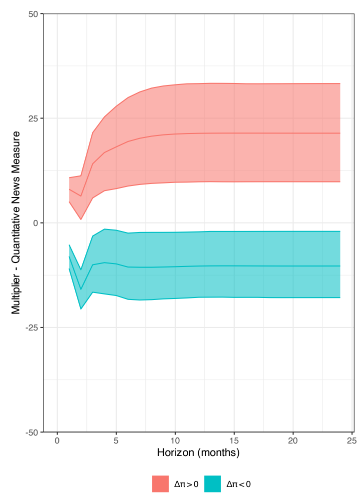

(E) FR: Normalized CIR

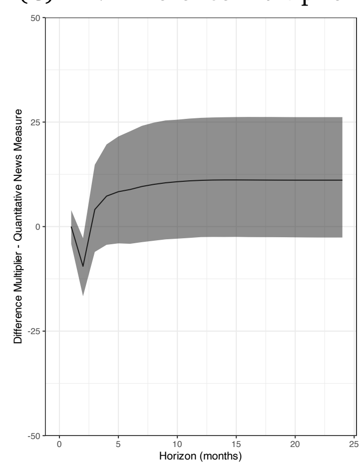

(F) FR: Difference Multiplier

Figure D1: Inclusion of Business Cycle Indicator - Quantitative News Measure

Note: The Figure shows the cumulative, normalized cumulative impulse response functions, as well as the difference between the normalized cumulative impulse response (dynamic multiplier) functions of the TSVAR model described in section 3.1 including a business cycle indicator in between the inflation rate and the quantitative news measure. It shows the response of the quantitative news measure to a one standard deviation shock in the inflation rate. CIR stands for Cumulative Impulse Response. All panels with subtitle "DE" show the results for the German-written newspapers, and those with subtitle "FR" the results for the French-written newspapers. The shadowed areas correspond to 68% confidence bands.

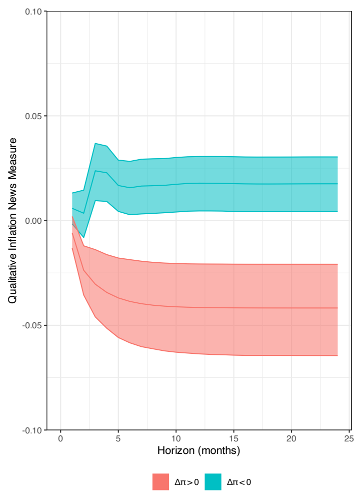

(A) DE: CIR

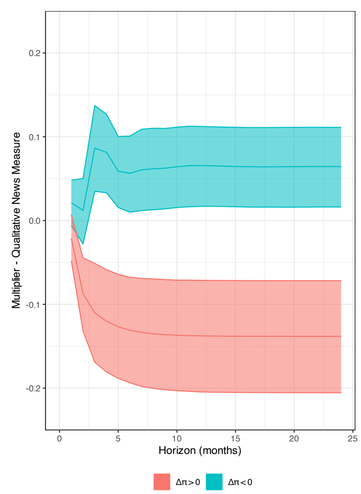

(B) DE: Normalized CIR

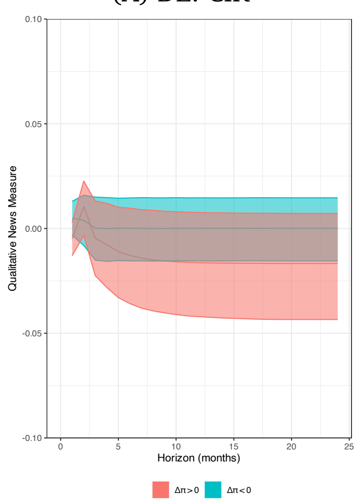

(C) FR: CIR

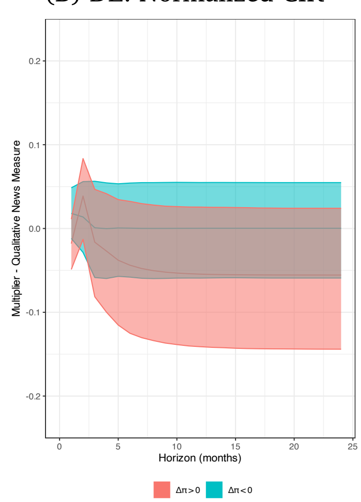

(D) FR: Normalized CIR

Figure D2: Inclusion of Business Cycle Indicator - Qualitative News Measure

Note: The Figure shows the cumulative and normalized cumulative impulse response functions of the TSVAR model described in section 3.1 including a business cycle indicator in between the inflation rate and the qualitative news measure. It shows the response of the qualitative news measure to a one standard deviation shock in the inflation rate. CIR stands for Cumulative Impulse Response. All panels with subtitle "DE" show the results for the German-written newspapers, and those with subtitle "FR" the results for the French-written newspapers. The shadowed areas correspond to 68% confidence bands.

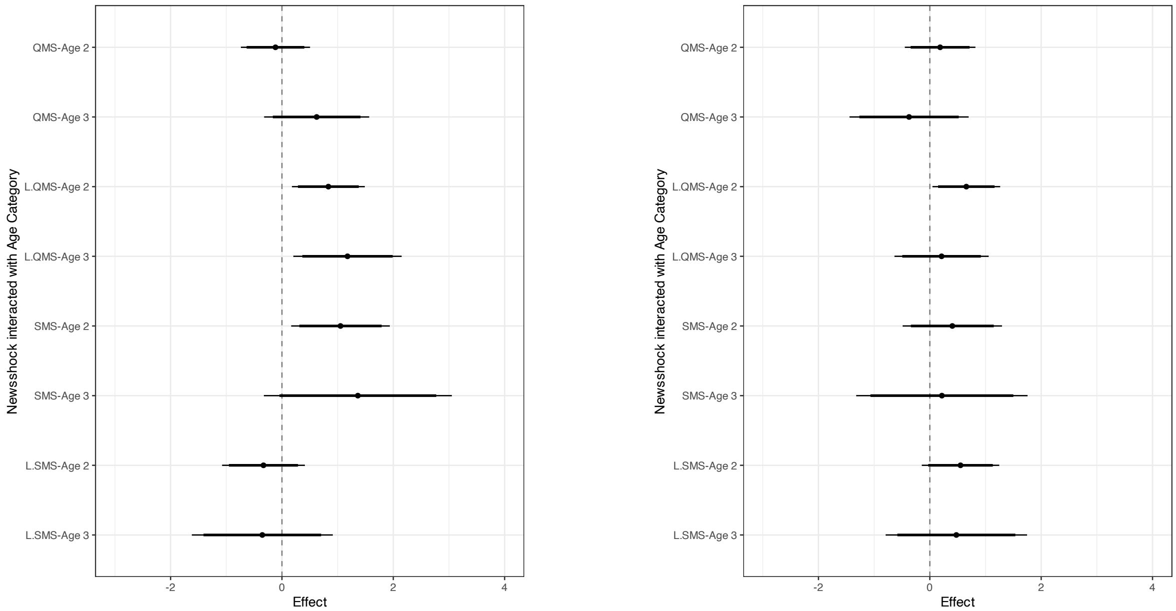

(A) Expectations

(B) Perceptions

Figure D3: Media shocks interacted with Age Categories

Note: The Figure shows the marginal effects of the media shock interacted with the Age Categories. The marginal effects are effects relative to the reference group, which are households aged between 15 and 40. The second age category are households aged between 40 and 60. The third age category are households aged 60 and above. Panel (a) shows the results for the interactions of the quantitative media shock (QMS), its lag (L.QMS) and the qualitative sentiment shock (SMS) and its lag (L.SMS) on the share of households expecting an inflation increase versus decrease. Panel (b) the same interaction effects but for inflation perceptions. The Figure shows 5% and 10% confidence intervals. The media shocks correspond to the model with the specification adding a business cycle indicator in between the inflation rate and the news measure, as described in the robustness section 6.

Table D1: Robustness Check: Estimation with Business Cycle Indicator

	Baseline		$\Delta\pi > 0$	$\Delta\pi \leq 0$	$\Delta\pi > 0$	$\Delta\pi \leq 0$	Region	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	reply$_e$	reply$_p$	reply$_e$	reply$_e$	reply$_p$	reply$_p$	reply$_e$	reply$_p$
Quantitative Media Shock$_t$	1.57**	1.08	1.30*	1.40	1.58	0.61	1.56**	2.49*
	(0.61)	(0.97)	(0.70)	(1.11)	(1.43)	(1.10)	(0.69)	(1.41)
Quantitative Media Shock$_{t-1}$	0.55	0.84	1.28*	-1.43	2.43**	-0.47	0.92	2.92***
	(0.66)	(0.86)	(0.69)	(1.44)	(0.99)	(1.06)	(0.66)	(0.77)
Qualitative Media Shock$_t$	0.19	1.09**	-0.86**	1.78	-0.12	3.28***	1.27*	0.64
	(0.47)	(0.47)	(0.37)	(1.13)	(0.46)	(0.78)	(0.73)	(0.47)
Qualitative Media Shock$_{t-1}$	0.09	0.37	0.08	-0.69	0.36	0.10	-0.33	0.44
	(0.48)	(0.49)	(0.50)	(1.36)	(0.59)	(1.03)	(0.69)	(0.68)
French-speaking Region=1 × Quantitative Media Shock$_t$							-0.32	-1.38***
							(0.46)	(0.51)
French-speaking Region=1 × Quantitative Media Shock$_{t-1}$							-1.37***	-1.01*
							(0.50)	(0.52)
French-speaking Region=1 × Qualitative Media Shock$_t$							-2.38**	-1.99*
							(1.09)	(1.11)
French-speaking Region=1 × Qualitative Media Shock$_{t-1}$							0.39	-1.22
							(1.13)	(0.91)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region$_r$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
HH$_{i,r,t}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,447	34,747	17,379	15,064	18,545	16,200	32,447	18,545
$\bar{y}$	85	86	89	81	89	83	85	89

Note: The table shows the main results from the model presented in section 5.1, estimated with a business cycle indicator in between the inflation rate and the news measure. Columns (1) and (2) show the unconditional effects of the quantitative and qualitative media shocks. Column (3) to (6) estimate the effects conditional on a positive or negative change in inflation. Columns (7) and (8) add the interaction effect of the region with the media shocks. The media shocks are standardized such that a one unit increase corresponds to a one standard-deviation increase in the media shock. $\bar{y}$ corresponds to the mean of the dependant variable. * p<0.10, ** p<0.05, *** p<0.010. Standard Errors are clustered at the date × region level.

D.2 Inclusion Stock Prices Growth

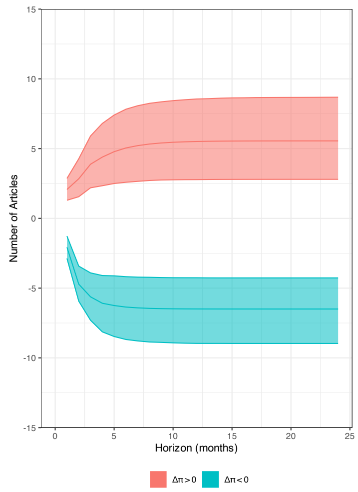

(A) DE: CIR

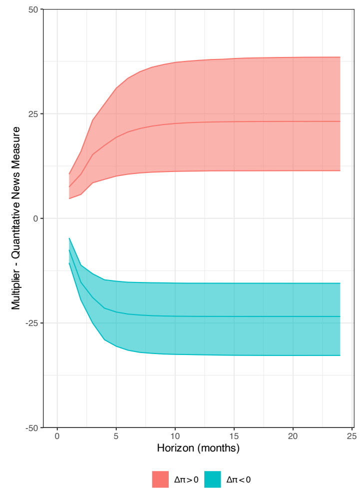

(B) DE: Normalized CIR

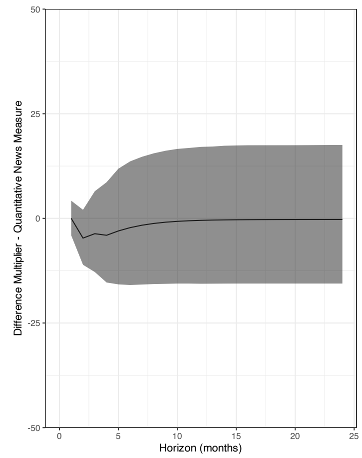

(C) DE: Difference Multiplier

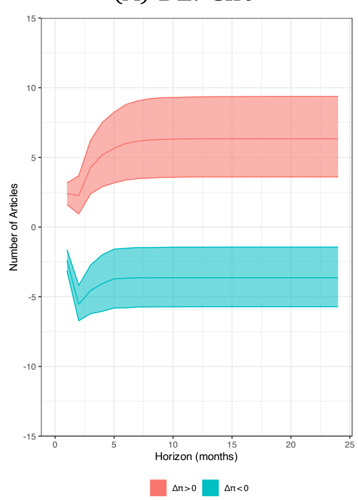

(D) FR: CIR

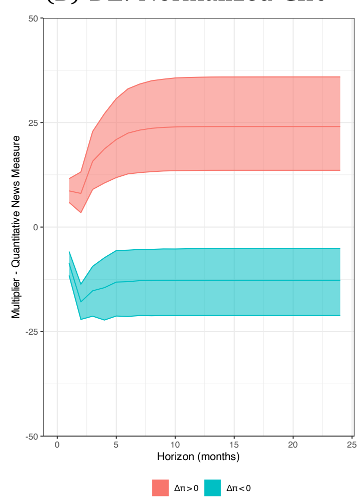

(E) FR: Normalized CIR

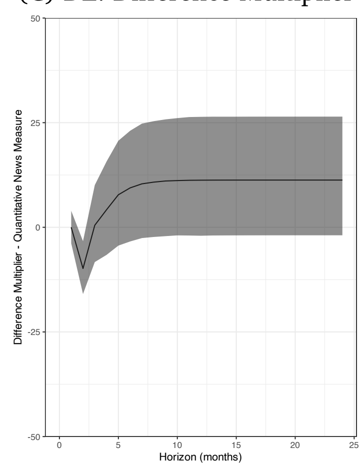

(F) FR: Difference Multiplier

Figure D4: Inclusion of Stock Prices Growth - Quantitative News Measure

Note: The Figure shows the cumulative, normalized cumulative impulse response functions, as well as the difference between the normalized cumulative impulse response (dynamic multiplier) functions of the TSVAR model described in section 3.1 including the growth rate of stocks as measured by the Swiss Market Index ordered last. It shows the response of the quantitative news measure to a one standard deviation shock in the inflation rate. CIR stands for Cumulative Impulse Response. All panels with subtitle "DE" show the results for the German-written newspapers, and those with subtitle "FR" the results for the French-written newspapers. The shadowed areas correspond to 68% confidence bands.

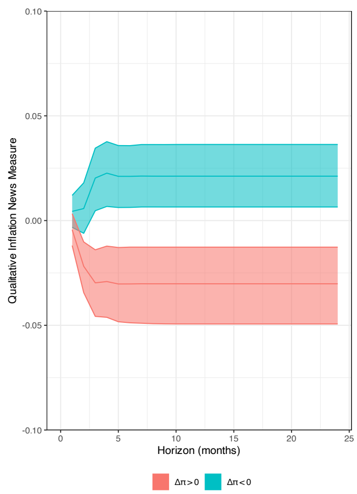

(A) DE: CIR

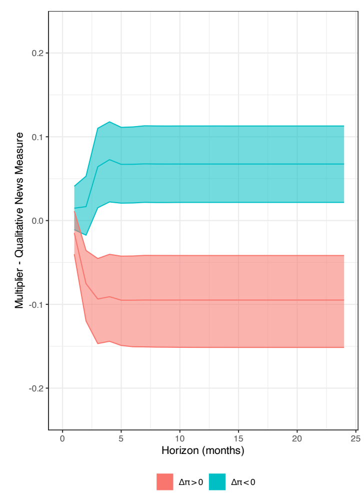

(B) DE: Normalized CIR

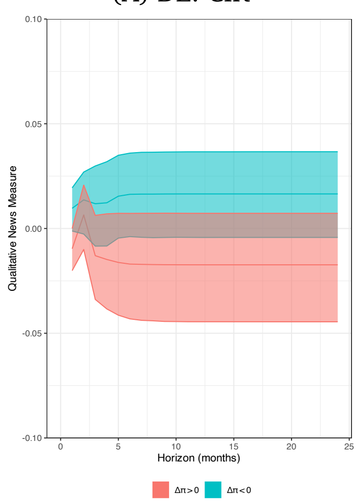

(C) FR: CIR

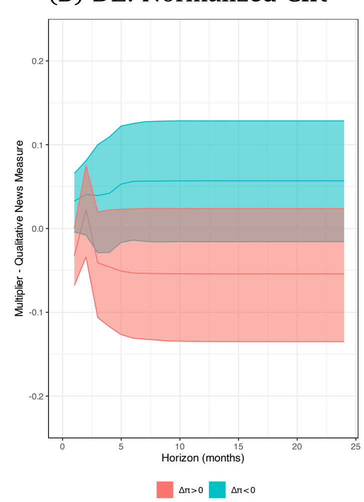

(D) FR: Normalized CIR

Figure D5: Inclusion of Stock Prices Growth - Qualitative News Measure

Note: The Figure shows the cumulative and the normalized cumulative impulse response functions of the TSVAR model described in section 3.1 including the growth rate of stocks as measured by the Swiss Market Index ordered last. It shows the response of the qualitative news measure to a one standard deviation shock in the inflation rate. CIR stands for Cumulative Impulse Response. All panels with subtitle "DE" show the results for the German-written newspapers, and those with subtitle "FR" the results for the French-written newspapers. The shadowed areas correspond to 68% confidence bands.

Table D2: Robustness Check: Estimation with Stock Prices Growth

	Baseline		$\Delta\pi > 0$	$\Delta\pi \leq 0$	$\Delta\pi > 0$	$\Delta\pi \leq 0$	Region	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	reply$_e$	reply$_p$	reply$_e$	reply$_e$	reply$_p$	reply$_p$	reply$_e$	reply$_p$
Quantitative Media Shock$_t$	1.29**	0.80	1.29*	0.97	1.44	0.32	1.27*	2.14
	(0.63)	(0.98)	(0.69)	(1.14)	(1.41)	(1.13)	(0.70)	(1.34)
Quantitative Media Shock$_{t-1}$	0.27	0.50	1.28*	-1.98	2.28**	-0.78	0.75	2.76***
	(0.68)	(0.88)	(0.72)	(1.45)	(1.03)	(1.11)	(0.68)	(0.79)
Qualitative Media Shock$_t$	0.19	1.09**	-0.85**	1.71	-0.10	3.29***	1.22	0.66
	(0.48)	(0.47)	(0.37)	(1.11)	(0.46)	(0.78)	(0.74)	(0.47)
Qualitative Media Shock$_{t-1}$	0.08	0.36	0.11	-0.80	0.44	0.01	-0.43	0.44
	(0.49)	(0.50)	(0.50)	(1.35)	(0.61)	(1.02)	(0.70)	(0.67)
French-speaking Region=1 × Quantitative Media Shock$_t$							-0.39	-1.34**
							(0.49)	(0.53)
French-speaking Region=1 × Quantitative Media Shock$_{t-1}$							-1.62***	-1.19**
							(0.55)	(0.57)
French-speaking Region=1 × Qualitative Media Shock$_t$							-2.37**	-2.02*
							(1.10)	(1.12)
French-speaking Region=1 × Qualitative Media Shock$_{t-1}$							0.57	-1.06
							(1.16)	(0.87)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region$_r$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
HH$_{i,r,t}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,447	34,747	17,379	15,064	18,545	16,200	32,447	18,545
$\bar{y}$	85	86	89	81	89	83	85	89

Note: The table shows the main results from the model presented in section 5.1, estimated with the growth rate of stocks as measured by the Swiss Market Index ordered last in the specification. Columns (1) and (2) show the unconditional effects of the quantitative and qualitative media shocks. Column (3) to (6) estimate the effects conditional on a positive or negative change in inflation. Columns (7) and (8) add the interaction effect of the region with the media shocks. The media shocks are standardized such that a one unit increase corresponds to a one standard-deviation increase in the media shock. $\bar{y}$ corresponds to the mean of the dependant variable. * p<0.10, ** p<0.05, *** p<0.010. Standard Errors are clustered at the date × region level.

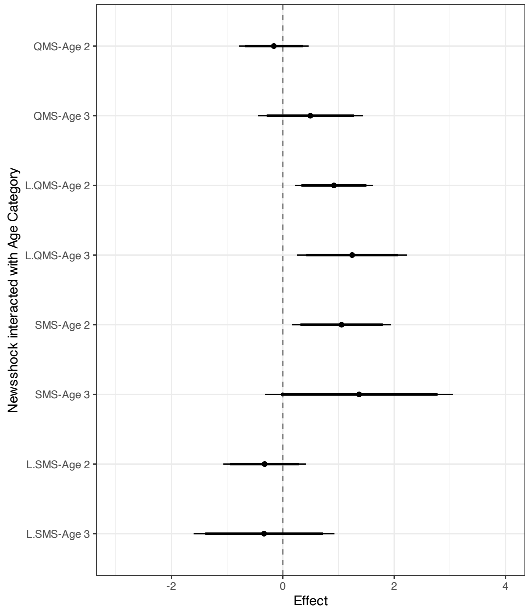

(A) Expectations

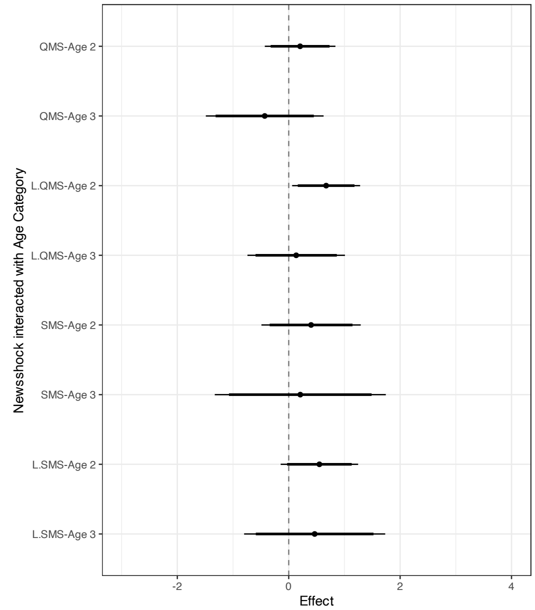

(B) Perceptions

Figure D6: Media shocks interacted with Age Categories

Note: The Figure shows the marginal effects of the media shock interacted with the Age Categories. The marginal effects are effects relative to the reference group, which are households aged between 15 and 40. The second age category are households aged between 40 and 60. The third age category are households aged 60 and above. Panel (a) shows the results for the interactions of the quantitative media shock (QMS), its lag (L.QMS) and the qualitative sentiment shock (SMS) and its lag (L.SMS) on the share of households expecting an inflation increase versus decrease. Panel (b) the same interaction effects but for inflation perceptions. The Figure shows 5% and 10% confidence intervals. The media shocks correspond to the model with the specification adding the growth rate from the Swiss Market Index ordered last as described in the robustness section 6.

D.3 Quantified Inflation Expectations

Table D3: Effect of Quantitative and Qualitative Media Shocks on Quantified Inflation Expectations

	Baseline	$\Delta\pi > 0$	$\Delta\pi \leq 0$	Region
	(1)	(2)	(3)	(4)
	$\mathbb{E}[\pi_{t+4}]$	$\mathbb{E}[\pi_{t+4}]$	$\mathbb{E}[\pi_{t+4}]$	$\mathbb{E}[\pi_{t+4}]$
Quantitative Media Shock$_t$	0.04	0.01	-0.00	0.02
	(0.25)	(0.03)	(0.04)	(0.03)
Quantitative Media Shock$_{t-1}$	-0.02	0.01	0.01	0.02
	(0.03)	(0.04)	(0.04)	(0.03)
Qualitative Media Shock$_t$	0.02	0.02	0.03	0.02
	(0.02)	(0.02)	(0.04)	(0.03)
Qualitative Media Shock$_{t-1}$	0.02	0.01	0.07*	0.03
	(0.02)	(0.02)	(0.04)	(0.03)
French-speaking Region=1 × Quantitative Media Shock$_t$				-0.04*
				(0.02)
French-speaking Region=1 × Quantitative Media Shock$_{t-1}$				-0.04
				(0.03)
French-speaking Region=1 × Qualitative Media Shock$_t$				-0.02
				(0.04)
French-speaking Region=1 × Qualitative Media Shock$_{t-1}$				0.01
				(0.05)
Date FE	Yes	Yes	Yes	Yes
Observations	154	80	74	154
$\bar{y}$	0.61	0.65	0.57	0.61

Note: The table shows the effects of quantitative and qualitative media shocks on the quantified inflation expectations. The quantification of qualitative survey data is described in detail in appendix B. * p<0.10, ** p<0.05, *** p<0.010. Standard Errors are clustered at the date × region level.