Essays on Economic Expectations and Monetary Union Heterogeneity

ESSAYS ON ECONOMIC EXPECTATIONS AND MONETARY UNION HETEROGENEITY

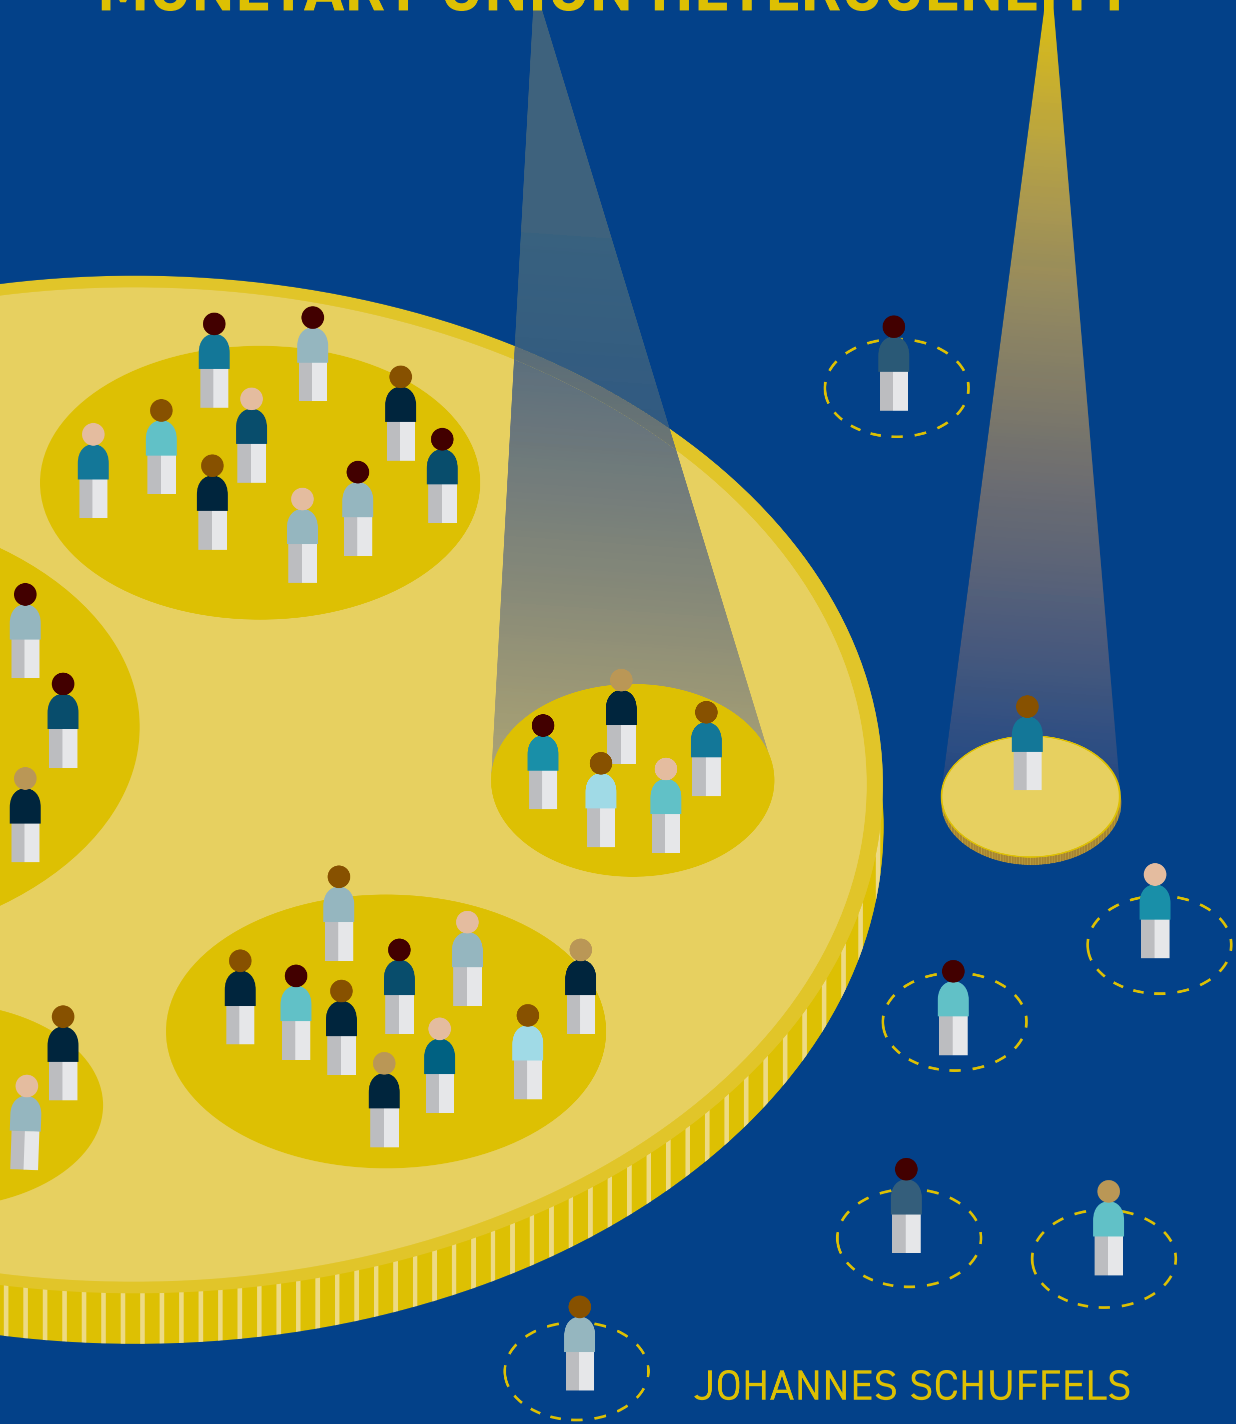

JOHANNES SCHUFFELS

Essays on Economic Expectations and Monetary Union Heterogeneity

Johannes Schuffels

2022

This book has been typeset by the author using LaTeX.

Cover design: Bregje Jaspers, 2022
Printing: ProefschriftMaken || www.proefschriftmaken.nl

Essays on Economic Expectations and Monetary Union Heterogeneity

DISSERTATION

to obtain the degree of Doctor at Maastricht University,
on the authority of the Rector Magnificus, Prof. dr. Pamela Habibovic,
in accordance with the decision of the Board of Deans,
to be defended in public
on Friday, 23 September 2022, at 13h00

by

Johannes Konrad Schuffels

Promotors

Prof. dr. Clemens Kool
Prof. dr. Tom van Veen

Co-promotor

Dr. Lenard Lieb

Assessment Committee

Prof. dr. Peter Schotman (chair)
Prof. dr. Jakob de Haan, Rijksuniversiteit Groningen
Dr. Maria Demertzis, Bruegel
Dr. Giulia Piccillo

Acknowledgments

The research contained in this thesis would not have been possible without my (co-)promotors Clemens Kool, Tom van Veen and Lenard Lieb. I am grateful to all three of you for trusting me with this position despite never having worked with me before. Clemens, thank you for rigorously challenging my work throughout the last 4 years. Learning how to respond to and anticipate these challenges has made me a better, more independent researcher. I am also thankful to you for supporting all of the decisions I took during the PhD even though I did not always follow your advice. Tom, you welcomed me at SBE in August 2017 and I immediately felt at home when I saw multiple pictures of Allianz Arena, my home town football club's stadium, on your office wall. Needless to say, we had at least as many chats about football as about research during my time in Maastricht. Thank you for always giving constructive advice on my research, even after retiring. Lenard, your active role as a co-author of the paper that forms Chapter 2 of this thesis has helped me learn how to put research ideas into practice. You were also instrumental in putting my ideas for Chapter 4 on the right tracks. You were always there for me to discuss anything from research or career choices to politics and football. Thank you for your support and friendship.

I would also sincerely like to thank Peter Schotman, Jakob de Haan, Maria Demertzis and Giulia Piccillo for assessing my thesis and giving me valuable comments as well as Joan Muysken and Olaf Sleijpen for their availability to be members of the opposition.

It is thanks to Luigi Moretti, my Master's thesis supervisor in Paris, that I decided to pursue a PhD in the first place. Thank you for your great classes and thesis supervision in Paris and your support in applying for PhD-positions.

Chapter 3 of this thesis would not have been possible without my co-authors Fiorella De Fiore and Marco Lombardi. You were very receptive of my attempts to start a research project together, invited me to Basel on two occasions and gave me the opportunity to present our research at the Bank for

International Settlements. I learned a lot from you and genuinely enjoy our cooperation.

In March 2020, I paused my PhD for five months for a traineeship at the European Commission. It is in no small part thanks to my traineeship supervisor, Stefan Zeugner, that I had the opportunity to start my career outside academia while, at the same time, finishing this thesis over the past two years.

The MILE department was my home at SBE for the past years. Thank you to Tania and Karsten, it was great fun organizing the workshop on globalization together. Thanks (again) to Giulia and also Kaj for being great colleagues. A special thanks has to go to Sylvia, Silvana and Marion for your constant help in all administrative matters.

Looking back, I am very grateful that the PhD has allowed me to make such great friendships along the way. First of all, Bart, we started our PhD on the same day and (if all goes well) will have defended our theses within 3 days. While this is of course coincidental, looking back at the past five years it is quite fitting. Despite working on seemingly very different topics at the start, realizing that we were experiencing very similar doubts and difficulties was incredibly important for me to be able to stay motivated. Halfway through, we noticed that even our research topics were quite similar in some senses and now we are also co-authors. We can truly say that we have gone through this together from beginning to end. Your constant curiosity leading you to come up with countless research questions over the years and your high standards as to what constitutes good research have always been exemplary for me. Brainstorming ideas with you was one of the most important steps for each research chapter of this thesis. But above all, you're just a very fine person and a great friend. Thanks for all of that.

Victor, I still remember how happy I was about the news that you would follow Max's and my footsteps from Paris to Maastricht. Since you arrived we became even better friends. Thanks to Isabella and you for always being there for me, many happy memories, your hospitality and our friendship. However, the group would have been far from complete without Mariana, Naina and

Max. Thank you for so many great moments together, the heated discussions about whatever was on our minds on a given day and many culinary delights (sorry, Max, the latter mainly applies to Mariana and Naina but I'm sure you would agree with me).

Thomas, we shared an office but more importantly, many of the joys and frustrations of doing a PhD. Our conference trip to Cyprus certainly belongs to the former category. Thanks for data artwork on our office walls, your introductory classes to Dutch politics (they are coming in very handy at the moment) and lots of laughter in A4.09. Luca, you were also part of the fourth floor community. Thanks for great trips to Cyprus and Oktoberfest and good times in Maastricht. Paun, with similar research interests and common supervisors we often felt like being in the same boat. I always enjoyed discussing economic expectations with you and I would like thank you for agreeing to be my paranymph. I would also like to thank many other friends who made these years in Maastricht very special for me. Thank you (in alphabetical order) to Adam, Alex, Benoit, Caterina, Christoph, Dewi, Elisa, Henrik, Inge, Inka, Kars, Kim, Martijn, Pascal, Sanne, Sean and Sergio.

Of course there was also a life outside of Maastricht and the university, especially since I left for Paris and Brussels. In this period you, Marta, have been one of the greatest sources of motivation for me to complete this thesis. Thank you so much for your advice, your support whenever I needed it and your patience when progress was much slower than we expected.

Lastly and most importantly, I want to thank my family. Oma, du hast immer mit viel Interesse meinen Fortschritt an dieser Arbeit begleitet und mir auch gesagt falls du ihn mal nicht ganz so groß fandest - womit du meistens Recht hattest. Mama, Papa, ihr habt viele der Voraussetzungen dafür geschaffen, dass ich diese Arbeit überhaupt beginnen konnte. Ich bin euch sehr dankbar für all eure Unterstützung und euer Vertrauen.

Contents

List of Figures

List of Tables

1
Introduction

The common motivation for the different lines of research in this thesis lies in the macroeconomic situation of Western economies in the mid-2010s. Figure 1.1 gives an overview of the development of some of the relevant macroeconomic variables over the past two decades in the United States and the Euro Area. Monetary policy reached historically accommodating levels by 2015 in both the United States and the Euro Area supporting the recovery in employment in both regions. Contrary to what many economists expected based on historical evidence, inflation did not respond and remained close to (in the case of the United States) or clearly below (in the Euro Area) target. As the fourth panel shows, households' expectations in the United States were much closer to reality: on average, they expected rather constant inflation rates over a one year horizon despite strong declines in unemployment and interest rates. In the Euro Area expectations were more volatile and elevated despite actual inflation rates that were below those in the US. While the swings in expected inflation among households in the Euro Area are fairly consistent with contemporaneous movements in unemployment, the strong increase in expectations since 2015 has not materialized in actual inflation. This apparent disconnect between monetary policy and unemployment on the one hand and inflation on the other as well as the mentioned differences in the behavior of expectations in the two regions opened up a range of research questions. The goal of this dissertation is to contribute to some of their answers.

One of the research agendas that emerged in response to the economic recovery after the global financial crisis of 2007/2008 focused on the economic expectations of the general public, in particular inflation expectations. This interest in economic expectations is in part driven by their renewed economic relevance in the macroeconomic context described above. With short-term nominal interest rates constrained by the zero lower bound, movements in the expected rate of inflation affect the real interest rate. When short-term nominal interest rates can no longer be reduced, inflation expectations among households and firms could therefore be an instrument for central banks to further reduce the real interest rate and kick start demand. To explore the viability of this strategy, some research focused on the identification of the effect of a change in inflation expectations on intertemporal substitution by households. Examples include Bachmann et al. (2015) who find a small but

Figure 1.1: Monetary policy, unemployment and inflation

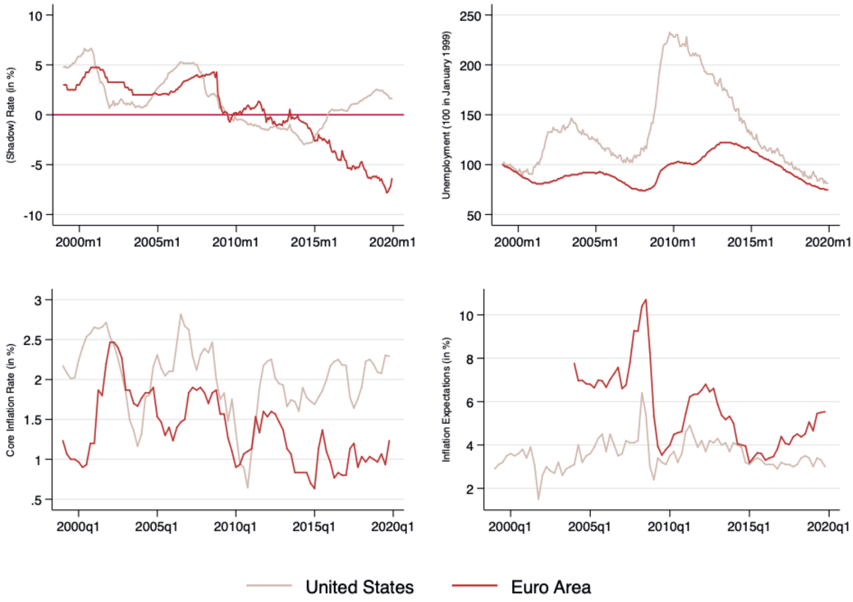

Note: The *top left* panel shows the evolution of the stance of monetary policy in the United States and the Euro Area. The series follow the official policy rates while they are above zero and take into account unconventional monetary policy when interest rates are below zero according to the methodology by Wu & Xia (2016). The *top right* panel shows unemployment indexed to 100 in January 1999 in both regions (Sources: BLS, Eurostat). The *bottom left* panel shows the core inflation rate (Sources: BLS, Eurostat). The *bottom right* panel shows expected inflation by households. The data on inflation expectations for the US comes from the Michigan Survey of Consumers and shows the average expected rate of inflation over the coming 12 months. For the Euro Area the data is taken from the European Commission's Business and Consumer Survey and measures the average expected rate of inflation over the coming 12 months starting in 2004.

negative association between inflation expectations and readiness to spend during the zero lower bound episode and an insignificant association outside the zero lower bound for the United States. By contrast, Duca et al. (2018) find for the Euro Area that there is a significantly *positive* association between changes in the expected rate of inflation and actual consumer spending. **Chapter 2** contributes to this literature by analyzing to what degree Dutch consumers' inflation expectations are linked to their purchasing decisions of durable goods, the goods category expected to be most affected by changes in the real interest rate. In particular, the chapter focuses on a channel that may moderate the reaction of consumption decisions to inflation expectations and has not received sufficient attention in the literature: the balance sheet of a household. The research shows that the degree to which households adapt their consumption in response to changes in their inflation expectations depends on their net financial position. The chapter gives some support to the argument made by Coibion et al. (2020) that monetary policy could use the steering of household inflation expectations as an additional tool when interest rates are constrained.

In order to use mechanisms like the one described above to affect, among others, consumption decisions of households, it is important to understand how expectations can be steered. It is therefore no surprise that the management of economic expectations became a centerpiece of central bank communication (Dincer & Eichengreen 2013). Before the start of the decade, highly attentive financial market participants were the target of such efforts. Over time, the communication strategies have been refined to reach all economic agents, such as businesses and households (Binder 2017). One component of the communication around the European Central Bank's (hereafter ECB) strategy review that was completed in 2021 can serve as an illustration of the degree to which communication strategies have shifted. A series of cartoons describing the components of the ECB's work and clearly targeted at the general public accompanied the documents outlining the strategy changes that were mostly intended for expert audiences. Two examples are shown in Figure 1.2. The left cartoon symbolizes, according to the ECB, the need for monetary policy makers to "reach out to wider audiences and make listening a regular feature of its communication" (Bank 2021). Similarly, the Federal Reserve has held so-called "Fed Listens" events in 2019 and 2020 to better understand the general public's economic and financial concerns (Federal Reserve 2020). The right panel of Figure 1.2 tries to illustrate the ECB's mandate of price stability around a symmetric inflation target in terms that are accessible to a broad audience, and is included as an example of communication *towards* citizens. In this context, academic research on the topic is still grappling with the question of how much attention the general public actually pays to monetary policy and how much it reacts to policy communication in the economic decision making process. **Chapter 3** is a contribution to this literature. The analysis exploits the fact that the national Survey of Consumer Expectations in the United States elicits responses around dates at which the Federal Reserve's monetary policy decision making body, the Federal Open Market Committee (hereafter FOMC), announces its decisions. This allows a comparison of the responses given right before the announcement to those given right after. Contrary to other studies on this topic using observational data the sample used in this chapter covers a large number of FOMC meetings and includes expectations on a wide range of economic variables. We find that most economic expectations are not affected by monetary policy decisions, even for highly financially literate respondents and around very salient decisions.

Chapters 4 and **5** discuss very different research questions than those before but share the same motivation that was described above. The relation between inflation and some measure of economic output that seems to no longer function the way historical data would suggest is formalized in most macroeconomic models as a variant of the Phillips Curve. The concept goes back to the empirical observation by Phillips (1958) that inflation and unemployment tend to be negatively correlated. Given the lack of such clearly observable correlations over the last decades, many potential explanations have

Figure 1.2: ECB communication surrounding strategy review

Source: https://www.ecb.europa.eu/home/search/review/html/workstreams.en.html, accessed on 18.09.2021, ©Miriam Wurster been proposed. These explanations range, among others, from inflation expectations being firmly anchored at the central bank's target (Boivin & Giannoni 2006) to falling unionization (Lombardi et al. 2020). **Chapter 4** shows, using a simple monetary union model, that if regions in a monetary union diverge in the degree to which the regional output gap affects inflation, a weaker correlation between inflation and the output gap emerges at the union level. As these reduced form analyses of the relationship between inflation and the output gap are common in monetary policy circles, the results carry policy relevance if decisions are based on this type of analysis. Fortunately, the econometric literature provides several remedies to take changes in sub-union heterogeneity into account to reliably estimate the Phillips Curve relationship.

Chapter 5 looks at the degree to which reduced form estimations of the Phillips Curve relationship in the Euro Area are affected by the mechanism described above and proceeds in two steps. First, the chapter provides evidence on the degree of heterogeneity in the effect of changes in unemployment[1] on inflation at the country-level. Equipped with this gauge of the existing heterogeneity, reduced-form estimates from regressions that do not control for the existing heterogeneity are compared to those that do. It turns out that estimates of the Phillips Curve relationship become slightly stronger when controlling for heterogeneity at the country level, particularly so in the period since 2010.

Finally, **chapter 6** summarizes the conclusions of the research contained in this thesis and reflects on potential directions for future research.

[1] Due to lacking data on the output gap at monthly or quarterly frequency, the unemployment rate is used instead, relying on the relatively stable relationship between the two variables first documented by Okun (1962).

2

Inflation Expectations and Consumer Spending: The Role of Household Balance Sheets[1]

[1] This chapter is based on Lieb, Lenard, & Johannes Schuffels. 2022. "Inflation Expectations and Consumer Spending: The Role of Household Balance Sheets." *Empirical Economics*. It makes use of data of the DNB Household Survey administered by CentERdata (Tilburg University, The Netherlands).

2.1 Introduction

Hypotheses on why inflation expectations can have an impact on consumption on the micro level are based on two arguments. First, inflation expectations impact the real interest rate and could therefore affect consumption through intertemporal substitution. Second, they affect expected real wealth and therefore consumption out of real wealth. In both cases the composition of a household's balance sheet can alter the size and direction of the effect of inflation expectations on spending. Attempts to gauge this interaction in the literature have been incomplete. Several authors estimated the impact of inflation expectations on consumption. These studies have often exploited some sort of natural experiment such as the zero lower bound or value-added tax increases to identify a causal relationship using cross-sectional (Bachmann et al. 2015; Ichiue & Nishiguchi 2015; D'Acunto et al. 2016) or panel data (Burke & Ozdagli 2013; Crump et al. 2015; Duca et al. 2018) without reaching consensus on the sign or size of the effect. However, no analysis has properly accounted for the potential role of the balance sheet as a moderator of the effect of price expectations on spending. In this chapter, we investigate empirically whether different components of a household's balance sheet interact with its inflation expectations in affecting realized consumer spending. To this end, we use panel data on household level balance sheets, inflation expectations and durable consumer spending from the Dutch Central Bank's (DNB) Household Survey.

While the use of micro level data to study the nexus between inflation expectations and consumer spending has allowed researchers to estimate cross-sectional effects, almost no attention has been paid to analyze the economic mechanisms behind these "general" effects. Changes in the real interest rate affect a household's optimal allocation of consumption over time. Differences in inflation expectations can lead to differences in the perceived real interest rate both over time and across households. Depending on their balance sheets, households might or might not be able to shift funds from savings to current spending or vice versa. Additionally, access to and costs of credit financed consumption might differ between households depending on the available collateral. We characterize these two channels through which inflation expectations can affect spending as real interest rate dependent. Another channel that motivates the research question of this chapter is a real wealth channel. Inflation expectations determine expected real wealth. In case of rising inflation expectations debtors expect increases in real wealth, while creditors expect falls in real wealth. The net nominal position of their balance sheet measures their exposure to price level changes. Empirical evidence suggests that consumption is sensitive to changes in wealth (Case et al. 2005; Mian et al. 2013). Consequently, inflation expectations and balance sheet positions might interact on the micro level. This could have macroeconomic effects if debtors have a higher propensity to consume than creditors. Here we refer to the growing heterogeneous agent literature that emphasizes the relevance of differential marginal propensities to consume of households with differing balance sheet compositions (Cloyne et al. 2020; Auclert 2019). Another reason is the inflation-hedging nature of certain assets: owners of real estate and stocks are relatively well protected against devaluation effects of inflation (Fama & Schwert 1977; Kim & In 2005) whereas financial liabilities are repaid in nominal terms. Accordingly, spending of net debtors is expected to be more sensitive to changes in expected inflation than for net owners of real estate and stocks.[2]

Our approach departs from the literature in important ways. First, we try to identify specific economic channels that determine the effect of inflation expectations on spending. The granular information on households' balance sheet in our data set allows us to test explicitly what role balance sheets play in moderating the effect of price expectations on durable spending.

Second, we analyze realized spending, rather than planned spending or attitudes towards spending. These two latter measures, often used in the literature, will likely overestimate a positive effect of inflation expectations on spending since households might be willing to consume but liquidity constraints impede them from doing so. Third, observing households over time allows us

[2] Other channels that are not affected by wealth have also been put forward: Wiederholt (2014) suggests that high inflation expectations could be a sign of policy uncertainty and thus depress spending. Cavallo et al. (2017) show that the existence of a relationship between inflation expectations and consumption can be explained by rational inattention: when the benefits of forming accurate expectations outweigh their costs - such as in episodes of high inflation - household spending behaviour is more sensitive to inflation expectations.

to better capture the intertemporal dimension of consumption decisions, which is particularly important if agents are forward looking and expectations play a crucial role.

Sufficient and accurate control for confounders in analyzes of large scale surveys poses problems. The DNB Household Survey contains a wide range of household characteristics. Including all characteristics that could potentially impact consumption behaviour is not feasible. Selecting controls only based on personal judgement or theory might lead to omission or unnecessary inclusion of some variables. Instead we apply a data-driven post-double variable selection procedure of the type introduced by Belloni et al. (2014a). With penalized regression techniques we only select those variables that impact the dependent variable and the independent variables of interest in the data. This limits the danger of omitted variable bias while ensuring a parsimonious specification. Moreover, the panel dimension of our data allows us to control for time-invariant confounders in general.

The results of the chapter give support to channels we classified as real interest rate and real wealth dependent. Financial investments amplify the effect of inflation expectations on spending which can be explained by the real interest rate channel. We also find that the positive relation between expected inflation and the probability of positive durable expenditures is amplified for households with lower net worth. The effect is stronger among a subsample of households with fixed interest mortgages. We interpret this result as evidence for the real wealth channel which depends on the *net nominal position* of the balance sheet combined with heterogeneities arising from the *composition* of the balance sheet.

The rest of the chapter is organized as follows. In section 2.2 we review the related literature. We discuss possible economic mechanisms that link consumption decisions, inflation expectations, and the balance sheet in section 2.3. The data is presented in section 2.4. In section 2.5 we present our econometric framework. Results are discussed in section 2.6. Section 2.7 concludes.

2.2 Related Literature

A number of influential contributions by Coibion, Gorodnichenko and co-authors (2012; 2015; 2017) have initiated a renewed discussion about the formation of inflation expectations and their macro and microeconomic effects. They provide substantial evidence that inflation expectations by consumers, businesses and even professionals and central bankers do not satisfy the conditions for full information rational expectations. Thus, consumers make systematic forecasting errors that, according to Coibion & Gorodnichenko (2015b), can help explain macro puzzles, such as the missing disinflation in the US after 2009. In this chapter we complement their work by investigating the channels through which consumers' inflation expectations affect microeconomic choices.

More closely related to our research question are previous studies that have used micro data to estimate the effects of inflation expectations on consumer spending. As stated above, no clear consensus has been reached on the direction or size of the effect. Bachmann et al. (2015) use repeated cross-sections of the Michigan Survey of Consumers to investigate the effect of inflation expectations of households on their "readiness to spend". The authors relate readiness to spend to a survey question on whether the current period is a good time to spend money on durable goods. They find that during the zero lower bound episode higher inflation expectations had slightly negative effects on the probability for households to have positive spending attitudes arguing that high inflation expectations might be correlated with increased economic uncertainty. The authors perform a number of regressions in search of heterogeneities in the relationship between inflation expectations and spending attitudes. For instance by including binary measures of home ownership and proxying an individual's debtor status with age. They do not specifically analyze wealth channels that moderate the spending response to inflation expectations. Ichiue & Nishiguchi (2015) approach the problem similarly, but with Japanese data and find strong positive effects of inflation expectations on planned spending. They argue that, after a long period of zero nominal interest rates, Japanese consumers have understood how inflation affects the real interest rate and therefore react. The authors do not further investigate the role of balance sheets. In contrast to both of these studies we construct a measure of realized spending and allow for a moderating role of balance sheet variables in the relation between expected inflation and spending.

A very different approach has been taken by D'Acunto et al. (2016). Their paper uses a value-added tax increase in January 2007 in Germany to estimate the effects of exogenous changes in inflation expectations. Compared to households in other European countries that did not experience the VAT increase, German households were substantially more likely to have positive attitudes towards spending in the months before the tax increase came into force. A limitation of this approach is that the price expectations of German households in November and December of 2006 contained considerably less uncertainty than those of households in other European countries. Households knew that a VAT increase will unambiguously increase prices of consumer products. They usually cannot form expectations with such certainty and precision. The effect of inflation expectations on consumption might differ substantially in times with less salient events or policy changes that nonetheless impact inflation.

The study most similar to ours is Burke & Ozdagli (2013). Using survey responses on expected inflation and realized spending on a wide range of products of a panel of American households between 2009 and 2012, they find much less clear results than the studies presented above. Households do not seem to increase their durable expenditures as a result of higher inflation expectations. In addition, they find evidence for effects on non-durable expenditures, driven by owners of real estate. Even though we analyze durable expenditures this finding justifies our strategy of carefully investigating potential interactions of expected inflation with balance sheet variables. Burke & Ozdagli (2013) can only observe binary measures of balance sheet variables, such as home ownership. Crump et al. (2015) estimate the subjective elasticity of intertemporal substitution based on survey responses on expected inflation and planned consumer spending of a panel of American households in the Survey of Consumer Expectations. They find that the elasticity of planned consumption to changes in expected inflation is around 0.5. While planned spending is a better proxy for spending than "readiness to spend", it isn't a realized measure neither. Based on a large panel of Eurozone households, Duca et al. (2018) find small positive effects of increased inflation expectations on households' "readiness to spend". While they control for household wealth, they do not examine the balance sheet channels we suggest.

2.3 Mechanisms

Next we discuss different mechanisms through which balance sheets could affect households' spending responses to changes in expected inflation. Potential candidates are real interest rate and real wealth changes that result from changed inflation expectations. In addition to balance sheet size and its net position, we also discuss how differences in its composition could moderate the spending response of inflation expectations.

Intertemporal Substitution

Consumers adapt their spending behaviour when relative prices change by substituting the more expensive for the cheaper good. Price changes over time also change the purchasing power of consumers' income in different periods which may affect their selected intertemporal consumption bundle. These standard substitution and income effects of relative price changes can be illustrated by the following basic set-up. Consider the following intertemporal budget constraint for a household with nominal income y_t, nominal interest rate i and consumption good c_t with price p_t in periods 1 and 2:

$$p_1 c_1 + \frac{p_2}{1+i} c_2 = y_1 + \frac{y_2}{1+i}$$

By normalising p_1 to 1 and defining $\pi^e = \frac{p_2 - p_1}{p_1}$ we can rewrite the previous equation as

$$c_1 + \frac{1+\pi^e}{1+i} c_2 = y_1 + \frac{y_2}{1+i}$$

An increase in π^e raises the expected future price of the consumption good relative to its current price and lowers the real interest rate. This triggers the standard substitution effect: consumers want to increase current spending relative to future spending since the price of the good is lower in the current period. In contrast, the direction of the income effect depends on whether the consumer is borrower or saver. The lower real interest rate benefits the borrower: by transferring income from period 2 to period 1, one can increase total consumption compared to a situation with higher real interest rates. Savers lose: the income they transfer from period 1 to period 2 earns less real interest, therefore total consumption falls. Even this very basic set-up predicts differential consumption responses for households based on their balance sheet position: debtors will increase their current consumption by more than savers if their expectations about future prices rise. The qualitative conclusion does not change if future income is indexed to inflation, only the degree to which consumption is transferred to the current period would be lower.

However, not all households face the same perceived borrowing conditions. Analogous to the argument made by Bernanke (1993) for firms, households with higher net worth are generally seen as more credit-worthy by banks and might face better borrowing conditions. Thus, even under constant economy-wide nominal interest levels the perceived borrowing conditions for households do not only depend on their inflation expectations. The same change in inflation expectations can lead to different household-specific perceived borrowing conditions if the balance sheet quality differs. Applying this idea to the relationship between inflation expectations and consumption is not new: Ichiue & Nishiguchi (2015) make the same point in their analysis, but cannot convincingly test it.

Real Wealth

An increase in expected inflation leads to a reduction in expected real wealth since the expected price level of the future period is now higher than before while nominal wealth has remained constant. For debtors the opposite is true: higher inflation will reduce the expected real value of debt and thus increase their expected net worth. The observation that changes in wealth have effects on consumption has been widely documented in the past using both macro and micro data (Case et al. 2005; Mian et al. 2013). The most appropriate measure for the exposure of a household's financial position to changes in the price level is its nominal net worth, i.e. assets minus liabilities.

Balance Sheet Composition

However, this view on the real wealth channel may be too simplistic. There are various reasons why differences in the composition of the balance sheet could lead to different consumption reactions of households with the same nominal net worth. First, there are differences in the sensitivity of various assets and liabilities to inflation. Real estate or financial investments can serve as a protection against inflation. Fama & Schwert (1977) have shown that returns on real estate protect fully against unanticipated as well as anticipated inflation. They regressed the expected nominal return of several assets on expected inflation. If the coefficient of expected inflation is equal to one, the nominal return compensates for losses in real returns on average. Thus, the expected real return does not change when inflation expectations change. More recent studies have confirmed the long-run inflation hedging nature of real estate and found mixed evidence for the short-run analysis conducted by Fama & Schwert (1977) (Anari & Kolari 2002; Hoesli et al. 2008). While Fama & Schwert (1977) cannot confirm the inflation hedging nature of stocks in the short term, later studies came to the conclusion that in the long-run stock investments have the same inflation hedging property as real estate (Schotman & Schweitzer 2000; Kim & In 2005). Households with a substantial part of their wealth invested in these asset classes might not regard higher future inflation as a threat to their future wealth since their investment strategy is designed to protect against such developments. Even if this protection is not perfect, it is superior to, say, for cash holdings. Households with cash holdings as their only assets have no way of protecting themselves against real losses due to inflation. Similarly, debt contracts usually specify a nominal amount that has to be repaid. Here, higher inflation expectations lead to an expected decrease in the real value of debt, i.e. increasing real wealth. To summarize, households who invested large parts of their wealth into real estate or financial investments are expected to exhibit less sensitivity to inflation expectations in their consumption decisions. Households with relatively large exposure to cash or debt may react more strongly since their expected real wealth necessarily changes in response to changing inflation expectations.

Composition effects could play a role on the liability side as well. While most liabilities are repaid in nominal terms, differences across liabilities arise with respect to the interest payment schemes. Specifics of mortgage contracts play an important role in the transmission of nominal interest rates to household behavior, especially consumption: Di Maggio et al. (2017) show that holders of adjustable rate mortgages respond significantly stronger to nominal interest rate shocks than those with fixed rate mortgages and without mortgages. These results have been confirmed in different settings. Cloyne et al. (2020) show that household balance sheet composition in the US and the UK alters the spending response to changes in the nominal interest rate, suggesting differing marginal propensities to consume between home owners with mortgages (high) and outright owners (low). Cumming & Hubert (2019) show a positive relation between the share of financially constrained (adjustable rate) mortgage holders and aggregate consumption responses to monetary policy shocks. While in the US and the UK interest on mortgages is predominantly paid at adjustable rates, interest in the Netherlands is predominantly paid at rates fixed for more than one year (83% of the total volume (DNB 2020)). We argue that households with these kind of mortgages are an interesting subsample to study the spending response to changes in inflation expectations on. The argument builds on a similar intuition as that applied by the authors cited above. Without nominal rigidities, changes in inflation expectations should not have real effects. The insensitivity of interest payments on fixed rate mortgages to nominal rates potentially increases the impact of changes in inflation expectations on real expected disposable income.[3] If the marginal propensity to consume for more constrained households is indeed higher, those fixed rate mortgage holders with lower net worth should exhibit a stronger response to changes in their inflation expectations. We test this hypothesis in subsection 2.6.3.

[3] This is the case under the assumption that real income stays constant

Any of the above channels imply that individuals with a different balance sheet composition (both concerning the relative sizes of assets and liabilities and the relative importance of specific classes of assets and liabilities) but identical changes in inflation expectations could exhibit differing spending responses. These considerations give rise to an econometric specification in which we allow for interactions between households' expected inflation and its different balance sheet components. Section 2.5 outlines how we aim to test the different mechanisms and what effects they would imply for our empirical analysis. By accounting for this interaction we depart from the previous literature on the topic. All of the aforementioned authors have stressed in their papers that wealth might play a role in the relationship between expected inflation and (durable) consumption. Our key contribution consists of testing this channel in a novel way.

2.4 Data

Our aim in this study is to explore the interaction between households' inflation expectations and their balance sheets in determining spending decisions. Information on all three variables needs to be at the household level and available for the same household over several years.

Contrary to previous studies, we set out to analyze realized consumer spending instead of attitudes to spending in general. However, specific survey answers on total (durable) expenditures might involve substantial measurement error. It is much easier to recall expenditures for specific durable goods since these items are seldom purchased and each individual purchase accounts for a substantial fraction of total spending of that period.

Additionally, our analysis requires balance sheet information on the household level. The literature on wealth effects on consumption concludes that different types of assets and liabilities might have different effects on consumer expenditures (Case et al. 2005). To provide a thorough account of the interaction we want to analyze individual balance sheet components as well as the net financial position of the households.

For the reasons mentioned above we make use of the DNB Household Survey (DHS) administered by CentERdata (Tilburg University, The Netherlands) and issued by the Dutch Central Bank (DNB). It includes households' self-reported balance sheets and their expected one year ahead inflation rate. Part of the self-reported balance sheet consists of vehicles owned by the household. We use this information to construct a variable of household vehicle expenditures (more details below). The DHS is an unbalanced panel of 12.439 households with annual observations between 1993 and 2018. More than one household member can respond to the survey. Since the balance sheets are aggregated at the household level, we primarily use responses to household member specific questions from the first member of the household. If the first member has not answered a specific question we use the response of the second member. This results in 52.055 household-year observations from which we construct our variables of interest.

We want to stress the unique fit of this data set for our purposes. To our knowledge, no previous study has made use of such extensive balance sheet information to analyze the effect of inflation expectations on realized consumer spending.

In the following, we give an overview of the different variables of interest and provide descriptive statistics.

Measuring Durable Consumption

In recent papers many authors concentrate on analysing the effects of inflation expectations on durable consumption (Burke & Ozdagli 2013; Bachmann et al. 2015; Ichiue & Nishiguchi 2015). We follow the literature in this respect. Durable consumption is the component of aggregate consumption most likely to be affected by variations in the real interest rate since it is more likely to be credit financed than expenditures on non-durable goods. Additionally, demand for non-durable consumption is less elastic to changes in macroeconomic conditions in general.

The DNB Household Survey does not include questions on expenditures on different classes of durable goods. However, households do report a large part of their assets. Among those are vehicles, such as cars, motorbikes and boats. For each of these items households report the purchasing price. We construct our expenditure variable by recording each time the purchasing price changes. For the extensive margin, the consumption variable takes the value 0 in case there is no change in the purchasing price and 1 in case there is a change. The fraction of households that have purchased a vehicle in a specific year is shown in Figure 2.1a.[4] For those households that did buy a car we construct a variable capturing the intensive margin of the purchase, i.e. the gross amount a household spent on vehicles. Figure 2.1b shows the mean, the 10th and 90th percentile of this variable's distribution over the sample period.

It is unclear whether we should expect the effects outlined above to materialize on the extensive or the intensive margin of a purchase. In theory, the mechanisms could play a role in both decisions a household has to make. When emphasising the extensive margin, we assume that households' tastes regarding durable goods are relatively fixed over time and the element of the decision that is subject to variations in expected inflation is the timing of the purchase. In a year in which a household has higher inflation expectations it might be more likely to buy the durable item it had already planned to acquire for longer. This reasoning is consistent with some results that emerged from the literature analysing the "hot potato" effect of inflation. The "hot potato" effect refers to the observation that consumers spend their money faster in times of high inflation. In a search based monetary theory model, Liu et al. (2011) find that inflation affects especially the extensive margin of the purchasing decision.

On the extensive margin, we observe 12,620 vehicle purchases throughout the entire sample period. In 39,435 household-year observations, no purchase has taken place. Figure 2.1c shows from how many household observations we can draw to construct the extensive margin variable. For roughly 30% of households we only observe the purchasing decision once. This means that

[4] The peak in 2009 in the extensive margin is due to a car scrapping scheme implemented by the Dutch government as a response to the crisis of 2008. No corresponding peak is observed on the intensive margin. This means households did not buy more expensive cars due to the scrapping scheme, there were simply more households that bought a car in that year. We use year-fixed effects to account for such effects.

these households participated in two consecutive waves of the survey, allowing us to evaluate whether the purchasing price of their vehicles changed. Figure 2.1d depicts the fraction of households with a certain number of vehicle purchases. For a majority of households we do not observe any purchase. Roughly 45% of households we observe between one and five purchases.

However, the sample that enters our regression analysis shrinks considerably since not all households answer all survey questions. Due to limited overlap with the variables capturing expected inflation, the remaining balance sheet variables, current and expected income, only 8663 observations from 3092 households enter our final sample. The application of the conditional logit model reduces our sample size further as it drops households for which the extensive margin variable does not change value. Therefore we are left with 4790 observations from 909 households.

On the intensive margin we would be limited to a much lower number of observations. In our preferred specification we would have to rely on a sample of 1476 observations from 1123 households. In a fixed-effects framework an average number of 1.37 observations per panel unit would not allow us to draw any meaningful conclusions. Therefore, we do not proceed with analysing the intensive margin further.

How much can vehicle expenditures tell us about durable consumption? To answer this question, we take a look at the aggregate durable and vehicle expenditures in the Netherlands. Figure 2.1e shows all subcategories of total durable consumption as defined by CBS, the Dutch statistical agency. Vehicle expenditures account for about 20 % of total durable consumption in the Netherlands across the whole sample period. They are the second biggest component of durable consumption after textiles and clothing. Additionally, as Figure 2.1f shows, they are highly correlated with total durable expenditures (correlation coefficient of 0.95 between 1995 and 2015).

Figure 2.1: Descriptive statistics (sources: DHS, CBS, own calculations)

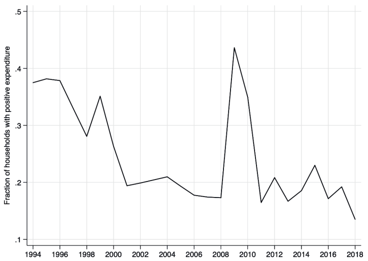

(a) Vehicle spending: extensive margin

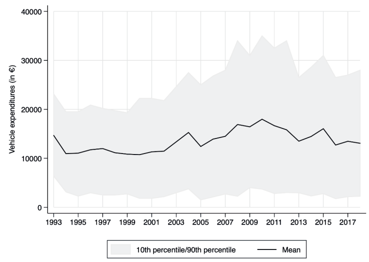

(b) Vehicle spending: intensive margin

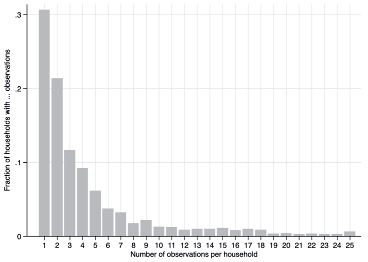

(c) Non-missing extensive margin observations

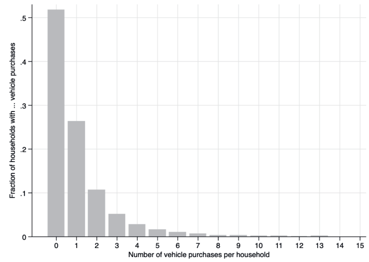

(d) Number of vehicle purchases per household

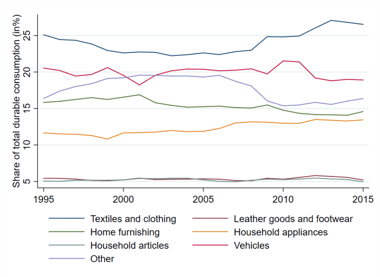

(e) Composition durable consumption

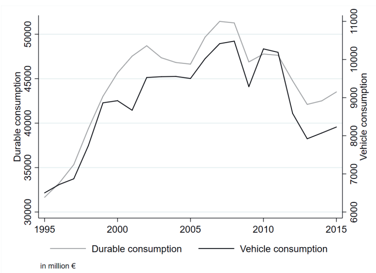

(f) Total durables and vehicle spending

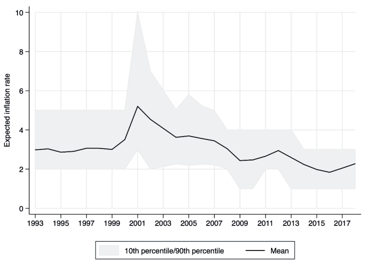

(g) Inflation expectations in Dutch Household Survey

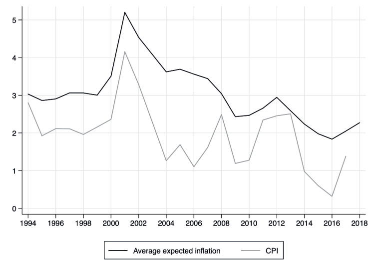

(h) Inflation expectations and CPI growth

Inflation Expectations

In the DHS households are asked the following question about their expectations for one year ahead inflation:

What is the most likely (consumer)prices increase over the next twelve months, do you think?

Since 2008 the possible answers are given between 1% to 10% in steps of one. Before, respondents were free to respond with any number they liked. Figure 2.1g shows the development of this variable over time. There is a clear peak after the introduction of the Euro. After that the downward trend in average expectations continues until well after 2008 and has stabilized close to but above 2% after that.

Figure 2.1h compares average expected inflation in the Netherlands with the realized CPI values. Expected inflation is structurally higher than realized inflation but trends are well anticipated by households. The latter observation is more relevant for our study since we are mainly interested in changes in inflation expectations. Secondly, this alleviates concerns that inflation expectations by (laymen) survey respondents are completely detached from actual inflation and instead measure expectations or perceptions of some other variables. However, the two series are very synchronized and tend to exhibit peaks and troughs in the same periods. One would rather expect the survey responses to lead realized inflation since respondents are asked what they expect inflation to be over the coming 12 months. Note that this might very well be the case. If respondents' expectations for the coming 12 months are elicited in January of a given year (and these expectations turn out to be correct), this figure would suggest some degree of synchronisation. The CPI value for that given year contains most of the 12 months that expectations were elicited for. Therefore, one cannot regard this figure as evidence that expectations merely reflect perceived inflation in a given period.

Balance Sheet

Table 2.1 shows the individual balance sheet components that households report as well as the aggregation level at which we include them in our models (in bold). Grouping of assets is largely determined by the liquidity of the balance sheet item. Among illiquid assets we differentiate between real estate and other assets to acknowledge the special role housing wealth could play. We group liabilities according to maturity. Mortgages and other longer term debt (referred to as loans) are aggregated separately. The net worth variable is constructed by subtracting liabilities from assets.

Instead of having to interpret our results in units of currency, we prefer to analyze percentage changes. The usual log-transformation is not well suited for our variables since many households do not possess some of the balance sheet variables. Their observations would be lost in case of a log-transformation. In the case of the net worth variable all negative net worth observations would be dropped as well. Instead, we perform an inverse hyperbolic sine transformation (ihs).[5] Table 2.2 gives descriptive statistics for all balance sheet variables that enter our regressions in the empirical analysis.

[5] This transformation has been widely used in empirical work on household wealth (Burbidge et al. 1988; Pence 2006). For values close to zero the transformation is approximately linear and it resembles a logarithmic shape for larger absolute values: $x^{ihs} = \log\left(x + \left(x^2 + 1\right)^{\frac{1}{2}}\right)$

Table 2.1: Balance sheet variables in the dataset and their aggregation in our analysis (in bold letters)

Assets		**Liabilities**	
Cash	Checking accounts Savings/Deposit accounts Deposit books	**Loans**	Private loans Loans from family/friends Study loans
Financial investments	Growth funds Mututal funds Bonds Stocks and shares	**Short term debt**	Extended lines of credit Finance debt Credit card debt
Real estate	Real estate, for own use Real estate, not for own use	**Mortgages**	Mortgages on real estate, for own use Mortgages on real estate, not for own use
Other illiquid assets	Employer-sponsored savings plans (ESSP) Savings certificates Single-premium annuity insurance policies (SPAI) Savings or endowment insurance policies Life insurances (as part of mortgages)		

Table 2.2: Descriptive statistics for balance sheet variables (in thousand €)

Variable	Panel	Mean	Median	Sd	Min	Max	Obs.
Net Worth	Overall	146.47	95.72	170.15	-255.19	1931.38	N = 8663
	Between			161.94	-255.19	1931.38	n = 3092
	Within			54.40	-504.74	797.69	T = 2.80
Cash Holdings	Overall	20.32	7.50	36.00	-11.34	455.45	N = 8663
	Between			31.60	-11.34	424.03	n = 3092
	Within			15.52	-159.50	258.45	T = 2.80
Financial Investments	Overall	7.01	0.00	23.59	0.00	333.32	N = 8663
	Between			22.23	0.00	333.32	n = 3092
	Within			9.36	-118.87	261.67	T = 2.80
Real Estate	Overall	177.08	172.00	161.42	0.00	1750.00	N = 8663
	Between			152.72	0.00	1373.44	n = 3092
	Within			51.03	-402.96	757.12	T = 2.80
Illiquid Assets	Overall	7.58	0.00	27.85	0.00	1704.21	N = 8663
	Between			34.89	0.00	1704.21	n = 3092
	Within			13.34	-176.27	425.35	T = 2.80
Loans	Overall	1.04	0.00	4.68	0.00	62.93	N = 8663
	Between			5.18	0.00	59.15	n = 3092
	Within			2.13	-29.96	54.98	T = 2.80
Short-Term Debt	Overall	0.72	0.00	2.97	0.00	45.75	N = 8663
	Between			2.71	0.00	29.37	n = 3092
	Within			1.67	-17.81	37.13	T = 2.80
Mortgages	Overall	63.75	22.00	87.35	0.00	566.00	N = 8663
	Between			81.57	0.00	545.00	n = 3092
	Within			34.89	-321.66	335.97	T = 2.80

Note: In the last column, N refers to the number of observations, n to the number of distinct households and T to the average number of periods a household is observed.

2.5 Empirical Approach

As pointed out in Section 2.3, there are several arguments why inflation expectations could matter for spending decisions and how wealth could alter size and direction of this relation. In this section we motivate our econometric approach in light of the transmission channels we aim to investigate. To that end, we run fixed effects linear probability models (LPM) as well as conditional logit (CL) regressions with the binary purchasing variable as dependent variable.

2.5.1 Specification

Our analysis consists of two baseline specifications. We estimate a fixed effects linear probability model as well as a conditional logit. Below we outline these two specifications. For the linear probability model we run the following regression:

$$Pr\left(c_{it}=1|E_{it-1}\left(\pi_t\right),W_{it-1},\boldsymbol{X}_{it-1},\alpha_i,\kappa_t\right)= \sigma E_{it-1}\left(\pi_t\right)+\delta E_{it-1}\left(\pi_t\right)\times W_{it-1}+\phi W_{it-1}+\boldsymbol{X}_{it-1}\boldsymbol{\theta}+\alpha_i+\kappa_t, \qquad (2.1)$$

where α_i and κ_t are household and year-fixed effects, $E_{it-1}(\pi_t)$ is household $i's$ expectation at time $t-1$ for the inflation rate at time t, W_{it-1} is the value of a particular balance sheet variable in $t-1$, and $\boldsymbol{X}_{it-1}$ is household i's set of other characteristics at time $t-1$.

In addition, we estimate the following conditional logit model:

$$\Pr\left(c_{it}=1|E_{it-1}\left(\pi_t\right),W_{it-1},\boldsymbol{X}_{it-1},\alpha_i,\kappa_t\right)= \lambda\left(\sigma E_{it-1}\left(\pi_t\right)+\delta E_{it-1}\left(\pi_t\right)\times W_{it-1}+\phi W_{it-1}+\boldsymbol{X}_{it-1}\boldsymbol{\theta}+\alpha_i+\kappa_t\right), \qquad (2.2)$$

where λ denotes the logistic function. The fixed effects logit model imposes the condition that $T>\sum_{t=1}^{T}c_{it}>0$, where T is the total number of periods that the household participated in the survey. This condition implies that only households whose expenditure variable takes on both possible values (0 and 1) are included in the estimation. We construct inference based on boostrapped standard errors.

Next we discuss how to interpret the models in equations (2.1) and (2.2) in light of the the mechanisms outlined in section 2.3. Two coefficients in the above regressions are of special interest: σ, the coefficient for expected inflation, and δ, the coefficient of the interaction term. δ measures in which direction and with what magnitude a specific balance sheet component scales the effect of inflation expectations on consumption. Conversely, when including a single balance sheet component, σ measures the effect of expected inflation on consumption if the household has no holdings of the balance sheet component. For instance, when including net worth as the balance sheet variable, σ measures the relation between inflation expenditures and spending if net worth would be zero. As we argued in Section 2.3, the real interest rate channel would suggest a positive effect of the interaction between expected inflation and household wealth, implying negative effects for any interaction between liabilities and expected inflation. In contrast, the real wealth channel would suggest a negative interaction effect between household wealth and expected inflation. However, many assets serve as hedges against inflation. The real wealth channel on its own would thus predict no significant interaction effect when financial investments or real estate holdings are interacted separately with expected inflation. Any interaction between liabilities and expected inflation is thus expected to have positive effects on the spending variable. The mechanisms that we discussed in Section 2.3 suggest opposite effects of the interaction between wealth and expected inflation. The coefficient of the interaction term is the average magnitude of the real interest and the real wealth channel. That is, if σ is significantly different from zero, one of the two effects dominates. However, this would not necessarily prove the absence of the other effect.

Table 2.3 gives an overview of the coefficients we would expect for the variables of interest in our regression if the channels could be measured separately. Thus, if the coefficients in our models align with the signs or magnitudes of these coefficients we could claim that the respective channel dominates over the other.

Table 2.3: Signs of coefficients consistent with the different potential channels for the interaction between the balance sheet variable and expected inflation (δ)

Channel	Coefficient	Expected Signs
Real Interest Rate	$\delta_{net\ worth}$	> 0
	δ_{assets}	> 0
	$\delta_{liabilities}$	< 0
Real Wealth	$\delta_{net\ worth}$	< 0
	$\delta_{inflation\ hedges}$	0
	$\delta_{liabilities}$	> 0

Timing of Households' Consumption Decision

We only include households that are observed in at least two waves of the survey, otherwise we cannot determine differences (or lack thereof) in their vehicles' purchasing prices. Since we construct the expenditure variable by comparing purchasing prices of vehicles and do not use specific questions on the subject, we do not observe the exact date of the purchase. In our regressions we relate the vehicle purchase that occurred between period $t-1$ and t to the balance sheet, inflation expectations and other characteristics observed in period $t-1$. Since households are asked about their expectations for the coming 12 months, we consider these 12 months as the current period in which the effect on spending should play out. Figure 2.2 shows which period's observation of each of the previously introduced variables is used in our analysis.

Selection of Controls

Consumers' purchasing decisions are driven by many factors. We attempt to isolate the role of inflation expectations and various balance sheet items. However, if we do not control for other key predictors, estimation of the coefficients of interest may be biased. While it is plausible to assume that current

Figure 2.2: Timing of the purchasing decision

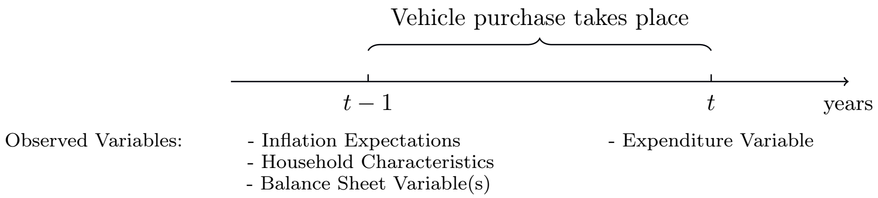

and expected income are relevant covariates in this context, the survey provides us with detailed information on individual household characteristics (e.g. attitudes toward saving and risk-taking, financial literacy, health, financial situation and expectations, etc.) and, hence, contain other possibly relevant predictors.

In order to identify relevant covariates, we use the "post-double-selection" method proposed by Belloni et al. (2014b). This involves a two-step LASSO regression, which in a first step selects covariates that predict the dependent variable, and in a second step selects variables predicting our independent variables of interest. The second step is necessary to control for the omitted variable bias. Note that selected controls may differ across regressions as we perform the "post-double-selection" for each regression separately. We always include current and expected income. Since we include individual fixed effects in all (LASSO) regressions we expect that most time-invariant household characteristics are controlled for, and only few (if any) additional controls are needed to estimate the impact of inflation expectations on car purchases.[6]

[6] Not including fixed effects results indeed in a number of additionally selected controls. Many of the selected covariates are often indeed time-invariant and make intuitive sense, for example "*Expected response to credit application*", "*Financial literacy*", or "*Car provided by employer*".

2.6 Results

For the exposition of the results of our analysis, we proceed in steps. First, we present the results from our baseline analysis in which we are mainly interested in the interaction terms between expected inflation and various balance sheet variables. We present estimates from fixed effects LPM and Logit regressions. In Table 2.4 results from the Logit regressions are marked as CL in the column title. Lastly, we analyze a subsample of households that have fixed interest rate mortgages.

As we argued in Section 2.3, the different balance sheet components are not expected to moderate the effect of inflation expectations on spending in the same fashion. The main reason are differences in their inflation-hedging potential. Certain assets like stocks or real estate protect the investor better against inflation than cash, for example. Additionally, we expect a difference between assets and liabilities in general. Debt is usually repaid in nominal terms, which makes its expected real value sensitive to expectations about inflation.

2.6.1 Balance Sheet Components

Table 2.4 presents the baseline results. For the regressions results shown in columns one and two, we included all single balance sheet components and their interactions with expected inflation. Collinearity is not an issue since net worth is not included and therefore free to move. The results do not depend much on the specification used, both the linear probability model (LPM) and the conditional logit (CL) give similar results. All but one balance sheet component do not significantly alter the relationship between inflation expectations and the probability to purchase a vehicle. For the interaction term between financial investments and expected inflation both the LPM and CL estimates are positive, the LPM estimate marginally above the 10% significance threshold, the CL estimate marginally below. The relation between expected inflation and the spending decision seems to be marginally different for households with within-household deviations from their average financial investment holdings compared to those at their average value. Households with higher than average financial investments exhibit a stronger positive reaction of expected inflation on their probability to spend. To quantify this relation, consider a household with inflation expectations 2%-points above their mean: a 10% increase in financial investments increases their predicted purchasing probability by around 3.8%-points. Compare that to a household that is 5%-points above their mean expected inflation: here, a 10% increase in financial investments increases the predicted purchasing probability by almost 10%-points.

Column 2 of Table 2.4 shows the results of the analogous conditional logit regression to the OLS regression in column 1. The results look qualitatively similar. The only balance sheet item that significantly alters the effect of inflation expectations on spending probabilities are financial investments. The estimated coefficient of 0.0226 corresponds to an odds ratio of roughly 1,023. An odds ratio larger than one mean that as the value of the interaction term increases, the odds of having positive vehicle expenditures in a given year rise.

A quantification of the fixed effects logit results in the same fashion as previously done for the linear probability model is not possible. Predicted probabilities can only be calculated by setting the fixed effects of all households to a uniform level and assuming different values for the explanatory variables. We want to stress that this is not an innocuous assumption. The reason why we chose to run fixed effects regressions is that we believe there are good reasons why time-invariant household heterogeneity should be controlled for in our analysis. By setting all fixed effects to zero we essentially assume this is not the case. The reason why we present our results in this way nonetheless is to illustrate how the estimated interaction effect would play out absent any other heterogeneity and to quantify our results in a meaningful way. The predicted probabilities are not to be interpreted as such literally. Including fixed effects would certainly alter them. Figure 2.3a shows the predicted probabilities of positive vehicle expenditures for different values of expected inflation and financial investments.[7] Each panel displays the predicted probability of posi-

[7] The predicted probabilities are obtained in the following way: for all combinations of a given grid of values for expected inflation (1 to 10 in intervals of 1) and the ihs-transformed net worth variable (fixed at the shown percentiles of the net worth distribution in 2018) the plot shows the average predicted probability across the sample

tive expenditures on the vertical axis and expected inflation on the horizontal axis for values of financial investments corresponding to the 50^{th}, 75^{th}, 90^{th}, 95^{th} and 99^{th} percentile of the distribution in 2018. The 50^{th} percentile corresponds to financial investments of 0 €. The vast majority of households does not invest in financial markets. For low levels of expected inflation predicted probabilities of positive expenditures don't differ much across households with different levels of investments. At the other end of the distribution of inflation expectations, the point estimates of the predicted probability are virtually unchanged for households without financial investments while those at the top of the investment distribution have a markedly higher probability to spend, reaching almost 50% for those in the 99th percentile of investments. Note that, as we pointed out above, this does not mean households like this necessarily have a probability to purchase a vehicle of 50% in a year. Unobserved, time-invariant household heterogeneity is not taken into account here. Additionally, the confidence interval becomes very wide for high levels of inflation expectations. Since both expected inflation and financial investments do not significantly affect the probability of positive expenditures on their own, the significant interaction term is not enough to produce predicted probabilities significantly different from zero.[8]

This result is in line with the real interest channel presented in Section 2.3. A falling perceived real interest rate increases incentives to substitute future spending for current spending. Only households with either sufficient collateral or sufficient internal finance are able to act on their increased willingness to

(not the predicted probability at the mean of the remaining covariates). Net worth in the regression was measured using positive values only but re-transformed to negative numbers for negative net worth for better readability. Each observation is treated as if the given values in the grids were the observed values for expected inflation and net worth. Then each household's predicted probability is computed based on the grid values and the remaining observed covariate values. The resulting probability in the graph is the average predicted probability for each combination across households. Additionally, the fixed effect for each household is set to 0.

8 One may argue that the significant interaction between financial investment and inflation expectations points to an endogeneity problem. If holders of risky assets would form more accurate expectations about future price level changes, the interpretation of the results above would be altered. To investigate that issue we compare inflation expectations of households that are holding risky assets with those of households that are not. Year-by-year KS tests show that the distributions of inflation expectations rarely differ between holders of risky assets and the rest of the sample.

spend. However, as the predicted probability plot shows, this moderating effect does not seem to be large enough to affect the outcome in an economically meaningful way.

2.6.2 Net Worth

We continue our analysis by taking a different perspective on the role that individual balance sheet components play. The rationale for analysing components individually is that they differ in terms of their return or real value sensitivity to inflation. At the same time, no component on its own is an appropriate measure of household wealth. Therefore, we now analyze whether net household wealth modifies the relation between expected inflation and the probability to spend. Columns 3 and 4 of Table 2.4 provide the baseline results of this analysis. We apply the same strategy as above by interacting the expected inflation rate of each household with their net worth to explain the following period's spending decision. The net worth variable is transformed from levels using the inverse hyperbolic sine transformation that accommodates zero and negative values while mimicking a log-linearisation (see Section 2.4).

In none of the three specifications in which we include net worth (columns 3, 4 and 5 in table 2.4, we find strong evidence in favor of a moderating effect of net worth on spending. However, all point estimates are negative and of similar magnitude. This means that for households with a net worth that is below their household specific mean, the predicted probability to spend increases. Figure 2.3b illustrates the results in the same fashion as previously done for the financial investments. We use net worth values corresponding to six different percentiles of the net worth distribution to compute the probability of positive expenditures that the estimation results predict. We can clearly see that at the lower end of the net worth distribution, i.e. households with negative net worth, there is a stark difference in the point estimates of the predicted probability of positive expenditure between low and high levels of expected inflation. The negative point estimates we found are driven by those households at the lower end of the net worth distribution. However, due to its insignificance and the imprecise estimation of the coefficients for expected

Figure 2.3: Predicted probability plots: financial investments and net worth

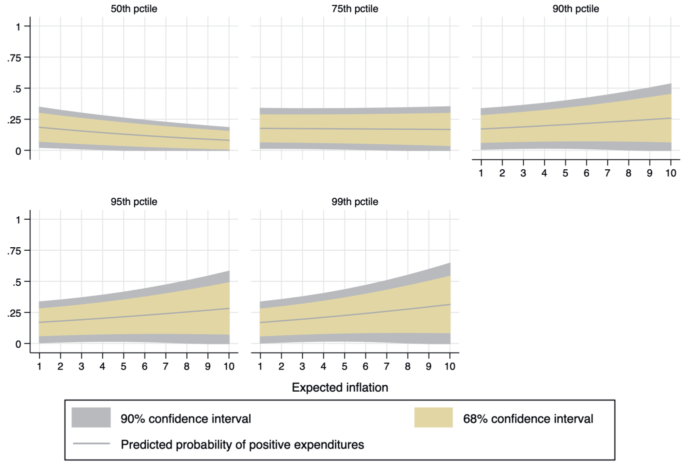

(a) Plot of the predicted probability of positive vehicle expenditures for given percentiles of the financial investments distribution (based on estimates of column (2), Table 2.4)

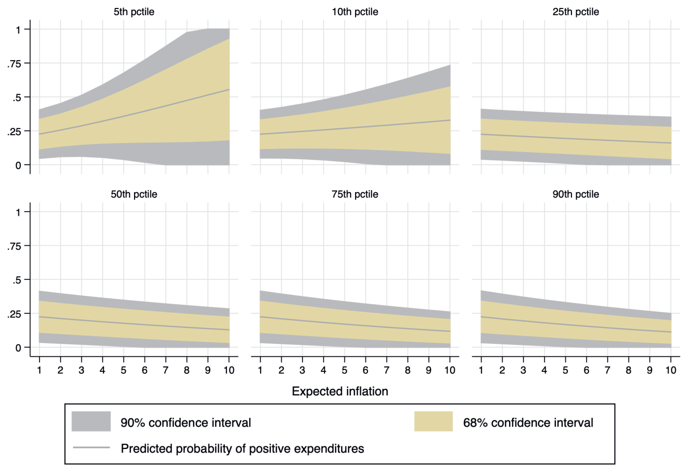

(b) Plot of the predicted probability of positive vehicle expenditures for households for given percentiles of the net worth distribution (based on estimates of column (5), Table 2.4) inflation and net worth, we cannot make strong statements about the robustness of this result. As the figure shows, for households with high expected inflation and low net worth, the 90% confidence interval includes all possible probabilities.

In Column 5 of Table 2.4 we present the results of a specification in which we include only the two balance sheet measures whose interactions with expected inflation we emphasized above: net worth and financial investments. This exercise supports the findings from above. Financial investments amplify the spending response to expected inflation while net worth has an (insignificant) dampening effect. This shows how the net nominal exposure to inflation (measured by net worth) *and* balance sheet composition (in this case, financial investments) can alter the spending response. While the former effect would support the relevance of a real wealth channel were it stronger, the latter is in line with the intertemporal substitution channel.

2.6.3 Fixed Interest Rate Mortgage Holders

For our research question fixed interest rate mortgage holders are an interesting case. An important part of their expenses is directly tied to the nominal interest rate. In our analysis thus far we have not been able to perfectly control for the nominal interest rate. Time fixed effects take out variation in spending decisions due to movements of the economy-wide nominal interest level. Controlling for net worth can also capture household specific movements in the nominal interest rate by acting as a measure of available collateral or the risk that a household will not be able to service its debt. Especially the latter is an imperfect measure though. The payment of interest on mortgages at fixed rates introduces an insensitivity of a large part of disposable income to business cycles. At the end of 2018 mortgages worth roughly 30 billion € were outstanding in the Netherlands, making up roughly 4% of GDP. Mortgages corresponding to about 83% of the total volume have interest rates that are fixed for more than one year (DNB 2020). With constant real income, changes in inflation expectations therefore have a direct effect on expected real disposable income. If less wealthy households have a higher propensity to consume, those households in our sample that are more financially constrained

Figure 2.4: Predicted probability plot: fixed rate mortgages (based on estimates of column (5), Table 2.5)

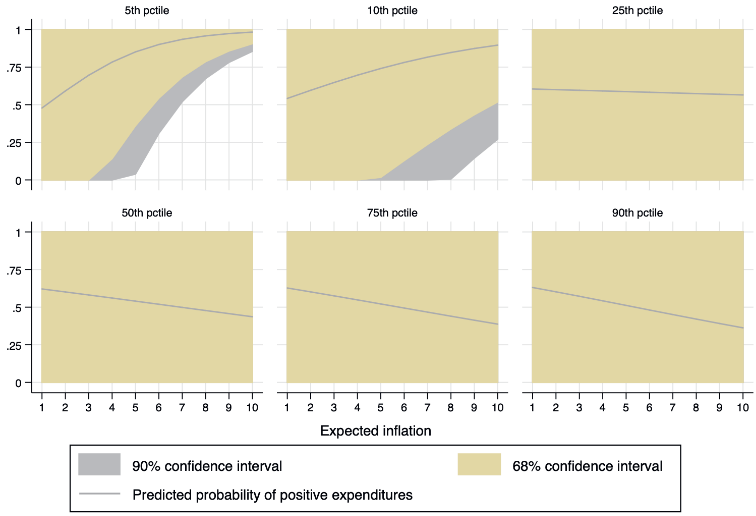

Table 2.4: Baseline results. Binary dependent variable of purchasing choice.

	(1) Components LPM	(2) Components CL	(3) NW LPM	(4) NW CL	(5) NW - Financial CL
Expected inflation (L1)	-0.0213	-0.160	-0.00112	0.0290	0.00623
	(0.0145)	(0.127)	(0.0104)	(0.0878)	(0.0817)
Expected inflation (L1) × Cash holdings (L1)	-0.000166	0.000587			
	(0.00135)	(0.0104)			
Expected inflation (L1) × Financial investments (L1)	0.00189	0.0156*			0.0187**
	(0.00123)	(0.00835)			(0.00857)
Expected inflation (L1) × Real estate (L1)	0.00109	0.00735			
	(0.00110)	(0.00841)			
Expected inflation (L1) × Illiquid assets (L1)	0.000188	0.00338			
	(0.00110)	(0.00853)			
Expected inflation (L1) × Loans (L1)	0.00281	0.0142			
	(0.00221)	(0.0170)			
Expected inflation (L1) × Short-term debt (L1)	0.000116	0.00163			
	(0.00165)	(0.0129)			
Expected inflation (L1) × Mortgages (L1)	-0.000911	-0.00703			
	(0.00111)	(0.00868)			
Expected inflation (L1) × Net worth (L1)			-0.000944	-0.00852	-0.0112
			(0.000854)	(0.00722)	(0.00695)
Net worth (L1)			0.000847	0.00495	0.0107
			(0.00262)	(0.0234)	(0.0216)
Net income (L1)	0.00202	0.00277	0.00222	0.00493	0.00447
	(0.00362)	(0.0275)	(0.00366)	(0.0295)	(0.0273)
Expected income (L1)	0.00293	0.00986	0.00286	0.0109	0.0119
	(0.00603)	(0.0412)	(0.00567)	(0.0442)	(0.0428)
Cash holdings (L1)	-0.000529	-0.00255			
	(0.00434)	(0.0334)			
Financial investments (L1)	-0.00228	-0.0245			-0.0322
	(0.00423)	(0.0278)			(0.0280)
Real estate (L1)	-0.00175	-0.0261			
	(0.00391)	(0.0299)			
Illiquid assets (L1)	-0.00133	-0.0181			
	(0.00332)	(0.0236)			
Mortgages (L1)	0.00385	0.0347			
	(0.00346)	(0.0270)			
Loans (L1)	-0.00534	-0.0285			
	(0.00732)	(0.0554)			
Short-term debt (L1)	0.0000559	0.00187			
	(0.00556)	(0.0409)			
Constant	0.488***		0.468***		
	(0.0901)		(0.0820)		
Household FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Lasso-selected controls	None	None	None	None	None
Observations	8663	4790	8663	4790	4790
Households	3092	909	3092	909	909
R^2	0.148		0.147		

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(i.e. those with a lower net worth) should exhibit a stronger spending response to expected inflation.

We apply this specification to the sub-sample of households with fixed interest rate mortgages. In our sample around 90% of households that report a mortgage as part of their balance sheet have a fixed interest rate mortgage. Unfortunately, the number of households with a variable interest rate mortgage is too low to perform the same analysis. We therefore resort to a sub-sample analysis instead of interacting all variables of interest with the mortgage's interest rate policy. Table 2.5 shows the results of this exercise. Apart from interacting the household's net worth with expected inflation we control for the household's net income as well as its expected income for the following period. The Lasso post double variable selection procedure did not select any additional control variables. A comparison with the results to those in column 4 of Table 2.4 reveals that the observed behaviour from the full sample is much stronger in the sub-sample of households with fixed interest rate mortgages. The coefficients on expected inflation, net worth and their interaction are all larger in absolute value and have a p-value below 0.1.

Figure 2.4 shows the predicted probabilities for different values of net worth under the assumption that the fixed effects are equal to zero (we refer to the previous section for a critical discussion of this assumption). Absent time-invariant household heterogeneity, the figure visualizes the mechanics of the interaction between expected inflation and net worth. Low net worth households with fixed interest rate mortgages react more strongly to higher inflation expectations than those with a higher net worth. This result holds when including an interaction term between expected inflation and the amount of outstanding mortgages the household has in its balance sheet. This interaction term is insignificant and its inclusion barely changes the values of the other coefficients of interest. In column 5 of Table 2.5 we include net worth and financial investments as balance sheet variables, the two measures that turned out to significantly affect the spending response in the whole sample. Among fixed rate mortgage holders the coefficient on financial investments is roughly the same as before, but not significant anymore.

These results show that while individual components of a household's balance sheet, such as fixed interest mortgages, matter for their consumption decisions the net nominal position determines the strength of this relation.[9]

How can these highly indebted households finance a vehicle purchase? Descriptive statistics can shed some light on this question. First, due to these households' likely limited access to external finance, we should expect them to buy less expensive vehicles. This is indeed the case: for households with negative net worth, the average purchasing price is only half that of the rest of the sample. Additionally, even though these households are net debtors, over 90 % of them have positive cash balances. This suggests that they do have internal finance available to make a car purchase. Another frequently applied method of payment for cars is to include the old car in the payment for the new one, in which case even less cash would be necessary.

Table 2.5: Subsample: households with fixed interest rate mortgages

	(1) NW CL	(2) Mortgages CL	(3) NW - Mortgages CL	(4) NW - Financial CL	(5) All CL
Expected inflation (L1)	0.180	0.0217	0.713	0.146	0.628
	(0.132)	(0.836)	(0.933)	(0.141)	(0.870)
Expected inflation (L1) × Net worth (L1)	-0.0219**		-0.0235**	-0.0242**	-0.0256**
	(0.0108)		(0.0113)	(0.0112)	(0.0125)
Expected inflation (L1) × Mortgages (L1)		-0.00676	-0.0427		-0.0386
		(0.0681)	(0.0736)		(0.0692)
Expected inflation (L1) × Financial investments (L1)				0.0206	0.0205*
				(0.0126)	(0.0118)
Net worth (L1)	0.0440	-0.0130	0.0480	0.0495	0.0531
	(0.0344)	(0.0147)	(0.0338)	(0.0348)	(0.0358)
Mortgages (L1)		0.0480	0.146		0.143
		(0.211)	(0.241)		(0.227)
Financial investments (L1)				-0.0378	-0.0377
				(0.0414)	(0.0399)
Net income (L1)	0.0491	0.0494	0.0477	0.0502	0.0487
	(0.184)	(0.209)	(0.180)	(0.187)	(0.160)
Expected income (L1)	-0.0509	-0.0513	-0.0528	-0.0494	-0.0509
	(0.0751)	(0.0689)	(0.0763)	(0.0785)	(0.0788)
Household FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Lasso-selected controls	None	None	None	None	None
Observations	2277	2277	2277	2277	2277
Households	458	458	458	458	458

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

9 A similar endogeneity problem as discussed in section 2.6.1 may bias our results. However, we do not find evidence that fixed interest rate mortgage holders are any better in predicting inflation. Year-by-year KS-tests show that expected inflation does not differ significantly between fixed mortgages holders and the rest of the sample.

2.7 Conclusion

In this chapter we provide evidence of a balance sheet channel through which inflation expectations affect durable consumer spending. We use a household survey that contains uniquely detailed balance sheet information as well as a large range of other household characteristics including inflation expectations. We discuss different hypotheses why balance sheets could potentially mediate the spending response to expected inflation. Our results suggest a mediating role of the real wealth channel: the positive response of the probability to spend when inflation expectations increase is stronger for households with lower average net worth. This effect is stronger for households with fixed interest rate mortgages.

We relate our findings to the growing literature on the consequences of agent heterogeneity for the transmission of monetary policy, in particular to Cloyne et al. (2020). They show that mortgage holders react particularly strongly to interest rate shocks in their spending choices. We show that a similar pattern is observable for changes in expected inflation.

We find differential effects of inflation expectations across the wealth distribution: households with high amounts of debt and substantially overestimated inflation expectations seem to commit costly mistakes if inflation does not live up to their expectations (which it did not throughout our sample). Here, our study connects well to Vellekoop & Wiederholt (2017). These authors show that households with higher inflation expectations have lower net worth and are less likely to own non-liquid assets, such as bonds, stocks or real estate. The remaining, inflation-sensitive balance sheet components have much higher relative importance than for households with lower inflation expectations. One conclusion for policy is therefore to improve the accuracy of households' inflation expectations. Recent research has shown that this can be done in two ways. More financially literate individuals tend to be better at forecasting inflation (Bruine de Bruin et al. 2010).

At the same time, central banks themselves can contribute to better formation of expectations. Coibion et al. (2019b) show that providing survey respondents with details about FOMC meetings - be it only the decision or the entire minutes - substantially improves the accuracy of their inflation forecasts. Better central bank communication could thus play an important role in helping households avoid costly mistakes in their economic decision making.

3

Are Households Indifferent to Monetary Policy Announcements?[1]

[1] This chapter is based on De Fiore, Fiorella, Marco Jacopo Lombardi, & Johannes Schuffels. 2021. "Are households indifferent to monetary policy announcements?" *BIS Working Papers 956*. The analysis makes use of the Survey of Consumer Expectations. Source: Survey of Consumer Expectations, © 2013-2020 Federal Reserve Bank of New York (FRBNY). The SCE data are available without charge at http://www.newyorkfed.org/microeconomics/sce and may be used subject to license terms posted there. FRBNY disclaims any responsibility for this analysis and interpretation of Survey of Consumer Expectations data.

3.1 Introduction

The effects of monetary policy depend critically on the public getting the message about what policy will do months or years in the future.

Janet Yellen, April 04, 2013 (Yellen 2013)

Central banks have taken measures to become more transparent and improve their communication with the public over the past decades (Dincer & Eichengreen 2013). More specifically, their communication strategies have been refined with the aim of reaching all economic agents – most notably households – rather than just a few highly attentive financial market participants (Binder 2017). Recent changes in major central banks' monetary policy frameworks have further emphasised the role played by the management of expectations. The shift to average inflation targeting at the Federal Reserve System and the emphasis on enhanced forward guidance at the European Central Bank signal confidence in the ability of monetary policy to affect agents' inflation expectations.

While the responsiveness of financial markets participants to monetary policy has been extensively documented (see, e.g. Gürkaynak et al. (2004), Bernanke & Kuttner (2005), Gagnon et al. (2011), D'Amico & Farka (2011), Andrade & Ferroni (2021), Swanson (2017)), the evidence on households' reaction is more scant, not least due to data limitations.

In this chapter, we contribute to filling this gap by analysing the response of US households' expectations to monetary policy decisions of the Federal Open Market Committee (FOMC henceforth). We use a sample of roughly 35,000 responses to the NY Fed Survey of Consumer Expectations between June 2013 and March 2019 to identify causal effects of monetary policy announcements. We do so by comparing responses given in the days right before FOMC meetings to those given right after. This identification strategy relies on the fact that respondents are assigned randomly to the day of the month on which they are invited to answer the survey. Our guiding principle in the analysis of the effects of monetary policy announcements on expectations is consistency with some basic economic relationships. Do respondents anticipate a negative relation between interest rates and output or unemployment, as predicted by a standard Euler equation? Do potential expected changes in aggregate demand and labor market conditions then feed through into respondents' inflation expectations, as predicted by the Phillips curve?

To characterise monetary policy announcements, we use several measures with a different degree of complexity. In principle, the multidimensional nature of monetary policy since the Great Financial Crisis suggests that a single quantitative measure of monetary policy surprises might not be sufficient to adequately characterise the decisions taken in a FOMC meeting. Yet less sophisticated households may lack the capacity (and willingness) to decipher fine details of monetary policy announcements and may only be able to digest simpler information. We therefore consider different measures of monetary policy actions in growing order of complexity and sophistication. First, we consider a crude dummy variable that takes unit value if there is a change in the policy instrument. Second, we look at changes in a shadow policy rate, taken as a summary measure of conventional and unconventional monetary policy. Finally, we consider monetary policy shocks extracted using high-frequency identification techniques (see Swanson (2021)) that portray more comprehensively the various facets of unconventional monetary policy: a Federal Fund Rate factor, a Large Scale Asset Purchase factor and a Forward Guidance factor.

We first test whether and to what extent households change their expectations about the aggregate economy and their personal financial situation, following the communication of FOMC's actions. Second, we study expectations about household consumption and personal job prospects. Expectations are policy-relevant insofar as they translate into consumption, labor and spending decisions. While we cannot observe such decisions, expectations about personal unemployment and spending plans can give us an indication of whether households *plan* to react to a monetary policy change.

The degree to which monetary policy affects both macroeconomic and personal financial expectations might differ across respondents. There is an extensive literature on personal characteristics that determine the understanding of, interest in, and reaction to news about the economy (see, e.g. Bruine de Bruin et al. (2010)). Financial literacy stands out as one of the most important determinants. We therefore test whether there are subgroups of the population that thanks to their economic and financial literacy, have an easier time deciphering monetary policy news and their impact on personal finances.

Our results point to a robust reaction of expectations about the interest rate on saving accounts to FOMC decisions in the sample as a whole. Respondents with higher financial literacy or exposure to financial decisions tend to respond more strongly. On the contrary, we find no effects of monetary policy announcements on inflation expectations or on expectations of personal financial conditions. We also document that households react particularly strongly in the first few days following the announcement of the decision, while the effect dissipates within two to three weeks of the decision. The results are also illustrated using one of the most salient episodes of monetary policy making that occurred during the sample period – the "Taper Tantrum" between June and December 2013.

This chapter relates to an experimental literature that documents the responsiveness of household economic expectations to specific types of central bank communication (Coibion et al. (2019b) and Coibion et al. (2020)). An important difference is that in this literature central banks' communication is forced upon the respondents – that is, households are given content to read before answering the questions, rather than having to fetch the information themselves, as in the real world. While these experiments provide very important contributions by isolating the mechanisms through which news about the macroeconomy affect household expectations, the strong impact found in experimental studies may be amplified by the clear-cut and easy-to-interpret information provision in those experiments.

Our methodology can assess whether the signal provided by real world monetary policy announcements is sufficiently strong to reach possibly inattentive households. Our contribution belongs to a different strand of literature that evaluates announcement effects using an event study approach (Lewis et al. (2019), Lamla & Vinogradov (2019), Claus & Nguyen (2020))). We add to this literature along three main dimensions. First, our long sample covering 47 FOMC meetings includes the period when policy rates were close to the zero lower bound and the Fed implemented unconventional policy measures. Second, we use a wide range of elicited expectations, on the macroeconomy as well as on personal finances, to understand whether household expectations adhere to simple yet meaningful economic relationships. Third, we exploit household characteristics to test whether numerical and financial literacy facilitate the transmission of monetary policy through expectations.

In the next section we present the related literature in more detail. Section 3.3 presents the survey data and discusses how we measure monetary policy decisions. Section 3.4 lays out our identification and estimation strategy. Our baseline results are presented in section 3.5. In section 3.6 we investigate whether exposure to financial decision making and financial and numerical literacy help households grasp the effects of monetary policy. Section 3.7 concludes.

3.2 Related Literature

Several recent papers study announcement effects of monetary policy decisions relying on event studies.

The closest paper to this chapter is Lamla & Vinogradov (2019) which surveys a random sample of the US population two days before and after 12 FOMC meetings between 2015 and 2018, applying a very similar methodology as we do. They find that announcements have no effects on respondents' inflation nor on interest rate expectations. One main difference relative to our approach is that they only have limited variation in treatment intensities across FOMC meetings. In 7 out of the 12 meetings they consider, interest rates had been increased, while during the remaining 5 meetings no changes occurred. Another key difference is that monetary policy changes are only measured using the federal funds rate. The effects of forward guidance or large scale asset purchases therefore cannot be measured. By opting for high frequency identified financial market shocks as treatment variable, we can differentiate between conventional and unconventional measures and document which one is more apt to affect household expectations. Our analysis also benefits from a larger sample of FOMC meetings. We cover 47 FOMC meetings which capture interest rate reductions as well as changes in unconventional monetary policies. The richer data and methodology have material implications for the results. We find that household interest rate expectations do react to FOMC announcements, despite the remaining expectations being largely insensitive. Moreover, we show that households with higher numerical and financial literacy revise interest rate expectations more strongly.

Claus & Nguyen (2020) take a different approach to identify monetary policy shocks on economic expectations of Australian consumers. They find that consumers react to monetary policy both in their macroeconomic expectations as well as their personal financial decisions. Inflation expectations seem to be well anchored and do not react instantaneously to monetary policy shocks. Despite tackling a very similar research question, their identification strategy differs quite substantially from ours. The authors identify unobserved news shocks driven by monetary policy changes in a latent factor model through co-movements in the second moments of elicited expectations on the day in which a monetary policy announcement occurred. We see our contribution as complementary to their work due to the differences in the identification strategy. The authors choose to identify shocks to respondents' information sets through heteroskedasticity in the expectations of the respondents themselves. They point out that this may be superior to measures obtained from financial markets as the expectation formation mechanisms may differ between consumers and financial markets. If that is the case, shocks identified on financial markets may simply be uncorrelated to reactions of consumer expectations. However, we argue that the "off the shelf" monetary policy measures we apply prove to be an adequate gauge of treatment intensity. Our FOMC-meeting specific analysis in section 3.5.3, which is agnostic concerning the expectation formation of consumers and does not make use of measures obtained from financial markets, confirms our baseline results.

Lewis et al. (2019) use daily Gallup consumer surveys in the United States to assess the impact of monetary policy news on households' consumer sentiment. The consumer sentiment variable is an aggregate measure of responses to two questions about household view of economic conditions at the time of the response, over an undefined future horizon. The authors then estimate the impact of monetary policy shocks on this time series using local projections. They find that an upward shock to the federal funds rate has significantly negative effects on consumer confidence in the days after FOMC meetings. They find no evidence for effects of Forward Guidance or Asset Purchases and do not test other, more naive, measures such as mere policy changes.

Bottone & Rosolia (2019) use Italian firm managers' survey responses around ECB Governing Council meetings to estimate the impact of monetary policy decisions on these managers' inflation expectations and their expectations about their own firms' prices. The dual focus – macroeconomic variables as well as firm-specific variables – is analogous to ours, yet on a different set of respondents. Monetary policy changes are measured using high-frequency financial market variables. They find significantly negative effects of those shocks on managers' inflation expectations while no effects are detected for the managers' own firms' price expectations.

Due to our wide range of survey questions covering both macroeconomic as well as personal financial expectations, we also relate to the literature on economic understanding among the general population. Dräger et al. (2016)

evaluate the consistency of survey answers to the Michigan Survey of Consumers with economic theory. In their sample, about half of the responses are consistent with the Taylor Rule and roughly one third are consistent with the Phillips Curve. A closely linked study by Andre et al. (2019) experimentally analyzes economic models of the general public. Following a hypothetical change in the Federal Funds Rate, the authors show that a substantial portion of individuals deviates from experts' predictions when forecasting the reaction of inflation to an interest rate shock, while predicting qualitatively similar changes in unemployment. Higher financial literacy increases consistency of responses with those given by experts. Another related experimental study is Roth & Wohlfart (2020). The authors randomly treat respondents in an online experiment with expert opinions on the likelihood of a recession. They find that negative macroeconomic expectations translate into higher personal job loss expectations as well as lower consumption growth expectations.

Overall, experimental studies tend to find strong effects of macroeconomic news on a range of household expectations - both regarding the macroeconomy and personal finances. On the contrary, our observational study based on the impact of FOMC announcements on household expectations suggests that real-world news provide weaker and more difficult to interpret signals relative to experiments in which information is explicitly provided.

The choice of personal characteristics that we condition on is guided by the literature on financial literacy. Lusardi & Mitchell (2014) document stark differences in financial literacy across demographic groups in developed countries and its impact on decision making. Individuals with lower levels of education tend to give more incorrect answers to survey questions eliciting financial literacy. Younger and older respondents perform worse, as well as women. Bruine de Bruin et al. (2010) show that individuals with lower financial literacy tend to have higher and less accurate inflation expectations than those with high financial literacy. In the experimental analysis of individuals' reactions to economic news Andre et al. (2019) find that higher financial literacy is correlated with reactions that are more similar to those of experts. We therefore investigate whether more financially literate respondents according to two complementary measures exhibit different reactions to monetary policy news compared to those with lower financial literacy.

3.3 Data

3.3.1 Survey of Consumer Expectations

The Survey of Consumer Expectations (SCE henceforth) is a monthly online survey conducted by the Federal Reserve Bank of New York, eliciting economic expectations among the U.S. population. The questions cover a wide range of macroeconomic as well personal financial expectations. To complete the core survey module respondents usually take 15 minutes. Participation in the survey is capped at 12 months, after which a respondent ceases to be surveyed. Outgoing respondents are being replaced on a rolling basis and new respondents are selected based on a stratified sampling procedure aiming to maintain a representative sample of the population in terms of its demographic and socioeconomic composition. In total, between 1200 and 1400 respondents are surveyed each month since June 2013. The sample available at the time of our analysis goes up to March 2019. Armantier et al. (2017) provide a comprehensive overview of the survey design. In the following, we will give more information about the sample composition and descriptive statistics on the outcome variables of our analysis.

3.3.2 Expectation Variables

Our analysis studies the reaction of economic expectations of SCE respondents to monetary policy announcements. Table 3.1 contains information about the outcome variables employed in our analysis. The survey makes use of two different approaches to the measurement of economic expectations. Some variables, namely those on interest rates on savings accounts, aggregate unemployment and stock market expectations, are elicited by asking respondents about the probability they assign to an increase in the respective variable over the 12 months following the survey response. The other macroeconomic and personal financial variables (except unemployment) are elicited in terms of their expected growth rate over a specified time horizon. Expectations about personal unemployment in the 12 months following the survey response are instead elicited by asking for the expected probability of that event. For the exact wording of each question we refer to the second column of Table 3.1.

Binder (2019) presents evidence that survey responses in the Survey of Consumer Expectations exhibit patterns consistent with panel conditioning. This problem occurs when survey respondents progressively change their behavior because of participation in the survey. Specifically, she finds that respondents in their early rounds of participation consistently revise their inflation expectations downwards, irrespective of actual inflation dynamics. Despite no evidence of non-random assignment of respondents to specific parts of the month according to the survey description, control and treatment groups exhibit differing average panel tenure levels if the full sample is used. We therefore exclude responses that are given before the seventh round of participation of each respondent. For more details on our approach on this issue we refer to Appendix A.

Table 3.2 gives descriptive statistics about each of the variables in the previous table, for the subsample of responses given after having participated at least seven times. The maximum number of observations possible per respondent is therefore five.

3.3.3 Measurement of Monetary Policy Announcements

Our selection of the monetary policy measures that characterise an announcement was guided by two counteracting considerations. First, monetary policy has become more multidimensional since the Great Financial Crisis. One single, quantitative measure like the Fed funds rate might therefore not be sufficient to adequately characterise the decisions taken in an FOMC meeting. Secondly, the capacity and willingness of agents to devote the necessary attention to monetary policy in order to understand its multidimensionality might be limited. This calls for a unified and easy-to-interpret measure. We

Table 3.1: Overview of economic and financial expectations

Variable Name	Survey Question	Time Coverage	Answer Range
Interest Rate 12m	What do you think is the percent chance that 12 months from now the average interest rate on saving accounts will be higher than it is now?	2013/03-2019/03	0-100%
Unemployment 12m	What do you think is the percent chance that 12 months from now the unemployment rate in the U.S. will be higher than it is now?	2013/06-2019/03	0-100%
Stock Market 12m	What do you think is the percent chance that 12 months from now, on average, stock prices in the U.S. stock market will be higher than they are now?	2013/06-2019/03	0-100%
Inflation Rate 12m	What do you expect the rate of inflation/deflation to be over the next 12 months?	2013/06-2019/03	$\mathbb{R}$
Inflation Rate 36m	What do you expect the rate of inflation/deflation to be between 24 and 36 months from now?	2013/06-2019/03	$\mathbb{R}$
House Price Inflation 12m	By about what percent do you expect the average home price to increase/decrease over the next 12 months?	2013/06-2019/03	$\mathbb{R}$
House Price Inflation 36m	By about what percent do you expect the average home price to increase/decrease between 24 and 36 months from now?	2013/06-2019/03	$\mathbb{R}$
Lose Job 12m	What do you think is the percent chance that you will lose your main/current job during the next 12 months?	2013/06-2019/03	0-100%
Household Spending 12m	By about what percent do you expect your total household spending to increase/decrease?	2013/06-2019/03	$\mathbb{R}$
Household Income 12m	Over the next 12 months, what do you expect will happen to the total income of all members of your household (including you), from all sources before taxes and deductions?	2013/06-2019/03	$\mathbb{R}$

therefore apply multiple measures that each weigh these two considerations differently.

The most naive measure of monetary policy decisions is a dummy variable that simply indicates whether the Federal Funds Rate (in non-ZLB periods) or the shadow rate (following the methodology of Lombardi & Zhu (2018) during the ZLB-period) increased due to the decisions taken at an FOMC meeting. It takes the value 0 in case of a constant or declining rate, and 1 otherwise. A more refined measure is the change in the shadow rate itself as it allows for varying treatment intensities.

While the change in the shadow rate in principle incorporates different dimensions of monetary policy, it neither disentangles the different dimensions nor does it reveal to what extent a policy change has been anticipated.

Table 3.2: Descriptive Satistics for Economic and Financial Expectations

Variable	Panel	Mean	Median	Sd	Min	Max	Observations
Interest Rate 12m	Overall	33.03	30.00	26.47	0.00	100.00	N = 36755
	Between			22.67	0.00	100.00	n = 7974
	Within			13.69	-41.97	116.37	T = 4.61
Unemployment Rate 12m	Overall	35.27	31.00	22.83	0.00	100.00	N = 36754
	Between			19.38	0.00	100.00	n = 7975
	Within			12.84	-48.07	116.27	T = 4.61
Stock Market 12m	Overall	40.25	45.00	23.47	0.00	100.00	N = 36629
	Between			20.11	0.00	100.00	n = 7958
	Within			12.60	-43.08	123.58	T = 4.60
Inflation Rate 12m	Overall	4.22	3.00	4.21	-5.00	25.00	N = 34750
	Between			4.10	-5.00	25.00	n = 7818
	Within			2.08	-15.78	24.88	T = 4.44
Inflation Rate 36m	Overall	4.25	3.00	4.30	-9.00	25.00	N = 34613
	Between			4.17	-5.00	25.00	n = 7807
	Within			2.18	-14.35	22.45	T = 4.43
House Price Inflation 12m	Overall	5.04	5.00	4.52	-10.00	20.00	N = 34902
	Between			4.17	-10.00	20.00	n = 7870
	Within			2.43	-14.12	22.44	T = 4.43
House Price Inflation 36m	Overall	4.88	4.50	4.56	-10.00	20.00	N = 34556
	Between			4.24	-10.00	20.00	n = 7838
	Within			2.48	-14.29	20.68	T = 4.41
Prob. to lose job	Overall	14.07	6.00	19.54	0.00	100.00	N = 21104
	Between			18.22	0.00	100.00	n = 4939
	Within			9.92	-69.26	97.40	T = 4.27
Spending 12m	Overall	3.63	3.00	6.15	-20.00	25.00	N = 34412
	Between			5.43	-20.00	25.00	n = 7847
	Within			3.80	-23.37	28.63	T = 4.39
Household Income 12m	Overall	3.62	2.00	6.04	-20.00	35.00	N = 33997
	Between			5.55	-20.00	30.00	n = 7788
	Within			3.44	-21.38	33.62	T = 4.37

High-frequency identified financial market surprises can deliver both a multi-dimensional view of monetary policy as well as a quantification of policy changes that were unexpected by financial market participants. The measures reflect the multidimensionality of monetary policy as they exploit the fact that the same FOMC decision can have different effects at different points along the yield curve. We therefore extend the analysis by considering the three monetary policy factors proposed by Swanson (2017). In a nutshell, the Federal Funds Rate Factor measures surprises at very short maturities, the Forward Guidance Factor at intermediate maturities and the Large Scale Asset Purchases Factor at longer maturities. Using these factors as monetary policy measures allows us to measure the degree to which different policies implemented by the Federal Reserve, e.g. forward guidance or asset purchases, have differential effects on agents' expectations. The three factors are estimated from the first three principal components of the asset price response in a 30-minute window around FOMC meetings. Due to the short time frame around the announcement of the decision, any change in these asset prices can be interpreted as a reaction to the decision taken at an FOMC meeting.

Jarociński & Karadi (2020) have shown that the same monetary policy surprise can have significantly different effects on the economy depending on the information that the central bank reveals with the decision. They differentiate between two situations: if stock markets rise after a tightening of monetary policy, the central bank has revealed that its information about the outlook of the economy is more positive than previously expected. The authors label this channel the "central bank information shock". In case stock markets fall in response to a policy tightening, as conventional monetary transmission would predict, a traditional monetary policy shock has occurred. In case of a tightening of monetary policy, the information shock is shown to have expansionary effects while the monetary policy shock has contractionary effects on the economy. In a fourth regression for each outcome variable we also include these shocks.

Table 3.3 lists the specific values of all seven monetary policy measures for each FOMC meeting that is part of our analysis. Additionally, it shows the exact observation count for each meeting. In our regressions, all continu-

Table 3.3: List of FOMC meetings

FOMC Meeting Day	Tigtening Y/N	Δ Shadow Rate	FFR	FG	LSAP	MP	CBI	Obs. Before	Obs. After
19 Jun 2013	1	0.19	0.16	1.28	1.96	0.01	0.01	154	123
31 Jul 2013	1	0.00	0.09	0.08	-0.23	-0.02	0.04	178	302
18 Sep 2013	0	-0.06	0.08	-1.34	-2.55	-0.06	0.06	314	236
30 Oct 2013	0	-0.34	0.10	0.08	0.33	0.03	-0.01	259	437
18 Dec 2013	1	0.10	0.21	0.02	0.63	-0.04	0.06	311	189
29 Jan 2014	1	0.16	0.22	-0.04	-0.24	0.04	-0.02	183	393
19 Mar 2014	1	0.26	0.06	1.04	0.57	0.03	0.01	315	219
30 Apr 2014	1	0.19	0.15	0.12	0.04	0.01	0.01	341	356
18 Jun 2014	1	0.22	0.09	0.41	-0.16	-0.01	0.03	323	319
30 Jul 2014	1	0.17	0.15	-0.09	-0.23	-0.02	0.03	457	483
17 Sep 2014	1	0.34	0.07	0.75	0.16	-0.00	0.03	423	316
29 Oct 2014	1	0.04	0.09	0.88	-0.01	0.04	-0.00	418	417
17 Dec 2014	1	0.49	0.29	-1.54	0.50	-0.07	0.03	317	334
28 Jan 2015	0	-0.17	0.16	-0.14	-0.14	0.02	-0.01	396	420
18 Mar 2015	1	0.12	0.19	-2.42	-0.77	-0.12	0.06	310	339
29 Apr 2015	0	-0.08	0.20	0.31	0.87	0.00	0.01	404	387
17 Jun 2015	1	0.04	0.09	-0.65	0.14	-0.04	0.01	342	353
29 Jul 2015	1	0.07	0.06	0.48	0.20	-0.01	0.01	370	335
17 Sep 2015	0	-0.12	-0.53	-1.53	-0.64	-0.04	-0.02	323	315
28 Oct 2015	0	-0.02	0.11	1.80	-0.05	0.07	-0.01	390	351
16 Dec 2015	1	0.11	0.31	-0.02	-0.54	0.01	0.02	353	273
27 Jan 2016	1	0.13	0.01	-0.46	-0.06	0.02	-0.03	392	361
16 Mar 2016	0	0.00	-0.11	-1.81	0.04	-0.06	-0.02	333	275
27 Apr 2016	1	0.01	0.10	0.33	-0.25	-0.00	0.03	370	364
15 Jun 2016	1	0.01	0.04	-0.78	0.19	-0.01	-0.02	336	319
27 Jul 2016	1	0.01	0.09	0.16	-0.32	-0.00	0.02	368	326
21 Sep 2016	1	0.01	-0.39	-0.18	-0.47	-0.03	0.01	333	343
02 Nov 2016	1	0.01	0.12	0.18	-0.05	0.00	0.01	340	366
14 Dec 2016	1	0.13	0.03	1.39	0.23	0.04	0.02	327	376
01 Feb 2017	1	0.12	0.13	-0.38	0.13	-0.02	0.01	401	444
15 Mar 2017	1	0.13	0.25	-1.31	0.03	-0.03	-0.00	401	376
03 May 2017	1	0.12	0.19	0.40	-0.00	0.01	0.02	322	364
14 Jun 2017	1	0.13	0.32	0.35	0.01	0.02	0.01	330	355
26 Jul 2017	1	0.11	0.10	-0.21	-0.21	-0.00	0.00	404	339
20 Sep 2017	0	0.00	0.05	1.17	-0.12	0.04	0.02	370	332
01 Nov 2017	1	0.01	0.14	0.14	0.02	0.00	0.01	373	354
13 Dec 2017	1	0.14	0.20	-0.21	-0.17	-0.00	0.01	378	296
31 Jan 2018	1	0.11	0.18	0.25	0.16	0.02	0.01	376	445
21 Mar 2018	1	0.10	0.12	0.11	0.37	-0.02	0.03	380	367
02 May 2018	1	0.19	0.16	-0.19	-0.10	-0.01	0.01	302	383
13 Jun 2018	1	0.12	0.02	0.84	0.10	0.03	0.01	370	349
01 Aug 2018	1	0.09	0.19	-0.05	-0.06	0.01	0.01	332	402
26 Sep 2018	1	0.04	0.31	-0.19	0.04	-0.01	0.02	343	386
08 Nov 2018	1	0.25	0.13	0.27	-0.05	0.01	0.02	330	379
19 Dec 2018	1	0.07	0.50	-0.04	-0.48	0.07	-0.03	323	357
30 Jan 2019	1	0.13	0.13	-0.67	0.08	-0.06	0.04	383	424
20 Mar 2019	1	0.01	0.36	-1.22	-0.18	-0.04	0.02	408	310

Note: Meeting Day indicates day of press conference. Change in shadow rate and shock variables are rounded to two decimal points. FFR: Federal Funds Rate Factor; FG: Forward Guicance Factor; LSAP: Large Scale Asset Purchases Factor (all Swanson (2017)). MP: Monetary Policy Shock; CBI: Central Bank Information Shock (all Jarociński & Karadi (2020)). Obs. Before reports number of responses up to (and including) 21 days before FOMC Meeting. Obs. After reports all responses up to (and including) 21 days after the meeting.

ous monetary policy measures – the change in the shadow rate, the Swanson factors and the Jarociński & Karadi shocks – are standardised with mean 0 and variance 1, while the table lists the non-standardised values. Standardis-

Figure 3.1: Identification Timeline

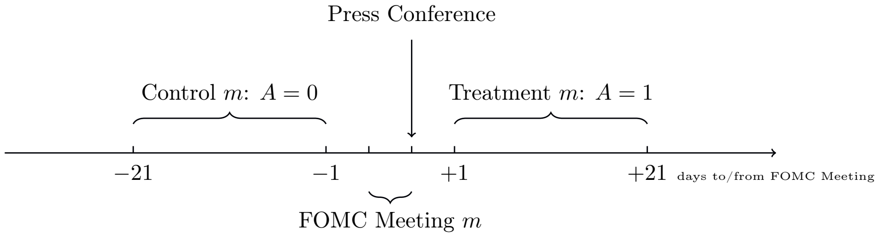

ing enables us to compare the magnitude of the coefficients. A comparison of the effects of the different shocks is useful to better understand the degree to which the multidimensionality of monetary policy is understood and processed by the general population.

3.4 Estimation Strategy and Identification

We estimate the treatment effect of monetary policy announcements by comparing the expectations of survey respondents right before the FOMC meeting with those given right after. This identification method for announcement effects of monetary policy is borrowed from event studies on financial market responses (see among others e.g. Gürkaynak et al. (2004) and Swanson (2017)) and has recently been applied to household and firm survey data at a lower frequency (see among others e.g. Bottone & Rosolia (2019) and Lamla & Vinogradov (2019)). Figure 3.1 shows the exact timing of our analysis for an exemplary FOMC announcement m. We use a symmetric time window around FOMC meetings. Each cohort is split into control and treatment group based on whether a survey response has been completed before or after a given FOMC meeting, measured by the treatment dummy variable A that takes on the value 0 before the announcement and 1 afterwards. We exclude any responses that have been filed on the days of a meeting as we cannot observe whether the survey module was completed before or after the meeting. Figure 3.2 shows the average number of responses to the question on interest rate expectations over the symmetric 42-day window around FOMC meetings. The number of

Figure 3.2: Average number of responses per day relative to FOMC meeting

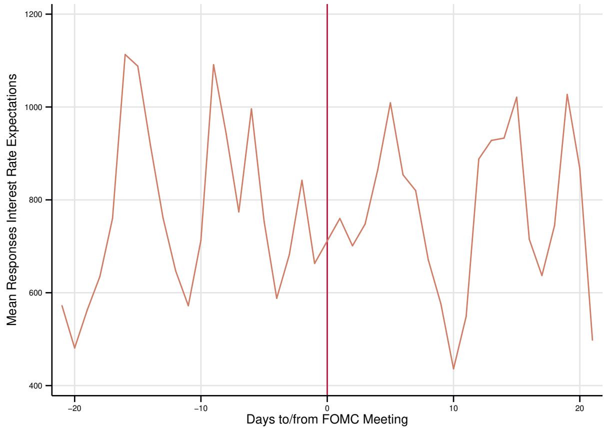

responses for the other questions follow virtually the same pattern. As the overview of the SCE explains, households are assigned randomly to one of the three batches in which the survey module is sent out in their first month of participation. Afterwards, respondent allocation is done with the aim of ensuring equal spacing between each monthly module, thereby implicitly preserving the initial random assignment of respondents to batches.

In our baseline analysis we run four separate regressions of each expectation variable from Table 3.1 on the treatment indicator interacted with i) a dummy variable indicating whether monetary policy was tightened at an FOMC meeting, ii) the change in the shadow rate according to Lombardi & Zhu (2018) between two FOMC meetings, iii) the three factors characterizing the surprise component of an FOMC decision according to Swanson (2017), and iv) the two factors disentangling surprises according to the information conveyed by the central bank, as proposed by Jarociński & Karadi (2020).

More precisely, we estimate the following baseline regression specification:

$$Y_{im} = \theta_m + \beta \times A_{im} \times \mathbf{s'_m} + \alpha_i + \epsilon_{im} \qquad (3.1)$$

where Y_{im} is the expectation response elicited from respondent i before or after FOMC meeting m, θ_m is a cohort-specific constant, β is a $1 \times k$ row vector, A_{im} is a dummy variable taking on the value 1 if a response of individual i was elicited *after* FOMC meeting m and $\mathbf{s_t'}$ is a $k \times 1$ column vector containing a constant and $k - 1$ monetary policy measures.[2] Individual fixed effects are denoted by α_i and ϵ_{im} is an idiosyncratic error term. Individual fixed effects control for time-invariant factors that might impact the level of expectations. Cohort-fixed effects are meant to control for all information that is common for those that answer before and after the FOMC meeting. Controlling for the common information is crucial for the assumption that the only relevant information treatment is provided by the monetary policy news generated by a given FOMC meeting. Standard errors are clustered at the respondent-level.

For the heterogeneity results in section 3.6 we interact the term $\beta \times A_{im} \times \mathbf{s_m'}$ with the levels of the factor variable measuring household characteristics that may determine the reaction to monetary policy news. We thus obtain a separate treatment effect for the base category and the remaining levels.[3]

Identification of causal announcement effects relies on the assumption that the only difference between the information sets of those in the control group and those in the treatment group is the content of the FOMC meeting. Our results are robust to a shortening of the time frame within which responses are included. The choice for 21 days before and after an FOMC meeting was made in order to maximize the number of observations per meeting.

2 When applying the first two measures, $k = 2$: a constant and the dummy variable or the change in the shadow rate. For the Swanson-Factors, $k = 4$: a constant and the three factors. For the Jarociński & Karadi-Shocks, $k = 3$: a constant and the two shocks.

3 Numerical literacy and the role in financial decision making are only elicited once for each respondent, the variables are therefore time-invariant at the respondent level. Both heterogeneity variables are factor variables.

3.5 Baseline Results

3.5.1 Macroeconomic Expectations

Table 3.11 reports the regression results for the following three outcome variables: expectations on interest rates on savings accounts for 12 months ahead, unemployment expectations for 12 months ahead and stock market expectations for 12 months ahead. For all three variables the respondent is asked to fill in the estimated probability that the variable will be higher 12 months after the survey response. It is therefore impossible to draw quantitative conclusions about the marginal effect of a monetary policy change on the expected *level* of the variable in the future. However, as the continuous monetary policy measures are normalised, we can compare the magnitudes of the coefficients and draw conclusions about the relative impact of the different measures.

Columns 1 to 3 in Table 3.11 report the effects of the post-FOMC meeting dummy depending on the direction of the monetary policy change. The treatment dummy is interacted with a dummy variable that takes the value 1 if the shadow rate increased between two FOMC-meetings. Therefore, the effect of the post-FOMC dummy alone can be interpreted as the treatment effect in case the shadow rate has remained unchanged or decreased. The results show that in case monetary policy is not tightened, the estimated probability of rising interest rates on savings over the 12 months following the survey response decreases significantly. On average it falls by about 1.2 percentage points, the coefficient is significant at the 5%-level. In case of a tightening of monetary policy, the size of the Post-FOMC coefficient increases to around 0.4. A joint test for significance of the sum of two coefficients yields an F-statistic of about 3, with a p-value below 0.1. As columns 2 and 3 show, we do not observe any significant effect on the expectations about rising unemployment or stock market prices.

Next, we move on to a more refined measure of monetary policy – quantitative changes in the shadow rate proposed by Lombardi & Zhu (2018). The pattern we observe for the tightening dummy variable holds for quantitative changes in the shadow rate as well. A one standard deviation increase of the shadow rate leads to an increase of the average probability that the interest rate on savings will rise in the 12 months following the response of about 0.5 percentage points. Coefficients of the shadow rate interacted with the treatment dummy in both the unemployment and stock market expectations regressions are close to zero and insignificant.

By using high-frequency identified financial market surprises of monetary policy decisions we try to disentangle the effect of different dimensions of monetary policy on expectations of the general public. The identification procedure of Swanson (2017) tries to capture differential effects of monetary policy decisions along the yield curve, the Federal Funds Rate Factor (FFR Factor) capturing the short end, the Forward Guidance Factor (FG Factor) the medium term and the Large Scale Asset Purchases Factor (LSAP Factor) the long end. Overall, the results from the previous two measures are confirmed: only interest rate expectations are robustly affected.

For an FFR Factor that is one standard deviation above its mean, expectations of higher interest rates significantly increase by about 0.6 percentage points on average (see Column 7, Table 3.11). At the mean interest rate expectation in the underlying sample, the effect corresponds to a 1.7% increase in the probability of rising interest rates. To put the magnitude of the underlying monetary policy shock into perspective, an example of an FFR Factor roughly one standard deviation above its mean is the FOMC announcement on December 14$^{\text{th}}$, 2014. The language used regarding the timing of monetary policy normalization changed between two press statements. The October statement read that the Federal Funds Rate Target could be raised a "considerable time following the end of the asset purchase programme" (Federal Reserve 2014b). The December statement indicated that "the Committee judges that it can be patient in beginning to normalize the stance of monetary policy" (Federal Reserve 2014a). This was perceived by financial markets as an indication that policy normalization could come sooner than previously expected. As the expected probability of rising rates within a year at that time was at 28% among the respondents on average, the effect measured by the regression in Table 3.11 implies an almost 2% increase in the expected probability on average.

The LSAP Factor also explains some of the variation in the interest rate expectations after FOMC meetings – in a similar magnitude as the FFR Factor.

Excluding the Taper Tantrum episode, that we discuss separately in subsection 3.5.3, this effect becomes insignificant. Variation in the FG Factor is not a powerful determinant of variations in interest rate expectations. Regarding unemployment, the effect of the LSAP Factor is marginally significant and negative. A standard deviation surprise tightening of LSAPs leads to a decrease of expectations of rising unemployment of about 0.4 percentage points. However, as for interest rate expectations, this effect is entirely driven by the Taper Tantrum episode, which is discussed in detail later. Stock market expectations are not significantly affected by the Swanson (2017) shocks, neither including nor excluding the Taper Tantrum.

The last set of results in Table 3.11 concerns the effects of high-frequency identified shocks proposed by Jarociński & Karadi (2020). Consistent with the previous set of results, interest rate expectations react similarly as to the FFR and LSAP Factors. Expectations about unemployment do not react. Contrary to the Swanson (2017) shocks, expectations about rising stock market prices react positively to both the monetary policy shock as well as the information shock. Conditional on being treated by the FOMC decision, the probability of higher stock market prices 12 months after the survey response rises by 0.361 percentage points for a one standard deviation monetary policy shock above its mean. The reaction to a central bank information shock is higher, at 0.486 percentage points for a one standard deviation increase.

Next, we turn to the results for the variable measuring expectations about attainment of the Fed's price stability mandate presented in Table 3.12. Expected inflation is elicited in a different format than interest rate, unemployment and stock market expectations. Respondents are asked to report their expected inflation rate over the 12 months following the response and for the period between 24 and 36 months following the response. One year ahead expected inflation does not react to either an easing or a tightening of monetary policy, as column 1 in Table 3.12 shows. The expected inflation rate 3 years ahead does not react either. The coefficients are very similar for both variables and close to zero. The same is the case for the the announcement effects of the change in the shadow rate. Neither one year nor three year ahead inflation expectations are significantly affected by a change in the shadow rate. The effects of high frequency identified shocks confirm this pattern. The coefficients for all three Swanson (2017) Factors are insignificant and close to zero. The shocks proposed by Jarociński & Karadi (2020) are also largely consistent with the previous null-results. The information shock tends to reduce inflation expectations both over 12 and 36 months and to a similar magnitude. Both effects are insignificant. The monetary policy shock has a marginally significant negative effect on 36-month ahead inflation expectations. The corresponding coefficient on 12-month ahead expectations is also negative but insignificant.

The last set of results on macroeconomic expectations concerns house prices. Respondents are asked to predict the growth of house prices nationally over the 12 months following the survey and over the period between 24 and 36 months following the survey. Table 3.13 reports the results. House price expectations do not react significantly to the dummy indicator and the change in the shadow rate. They are affected by the Swanson (2017)-shocks. A one standard deviation increase in the FFR Factor decreases expected house price growth over the 12 months following the survey by about 0.07 percentage points, the coefficient is somewhat closer to zero for 36-month ahead expectations. At the mean house price growth expectation (about 5.1%), the effect on 12 months ahead expectations corresponds to an expected 1% decline of the house price growth rate in response to a one standard deviation tightening of the FFR Factor. The LSAP Factor exerts a significant and counter intuitive effect on the 12 month ahead house price growth expectations. They increase by about 0.1 percentage points when the LSAP Factor is one standard deviation above its mean. The LSAP coefficient on house price expectations between 24 and 36 months ahead turns insignificant. However, these effects remain unconfirmed by any other measure of monetary policy as none of the two variables respond significantly to the Jarociński & Karadi (2020)-Shocks.

The effects on interest rate expectations are stronger than the results obtained by Lamla & Vinogradov (2019). The authors find no effect on interest rate and inflation expectations after 12 FOMC meetings between 2015 and 2018. However, when limiting our sample to those 12 meetings the treatment effect according to all measures we apply is below that of the complete sample. The coefficient of the FFR Factor falls to 0.439 but the joint treatment effect of a standard deviation FFR increase above its mean remains significant. Regarding inflation expectations our results are consistent with Lamla & Vinogradov (2019) irrespective of the sub sample as our and their analysis finds no effects that are significantly different from zero.

Our first key result is:

Result 1: *Only expectations about the interest rate on savings accounts are robustly affected by the various measures of monetary policy changes. For any measure of policy tightening/easing the expected probability of rising interest rates increases/falls. No other macroeconomic expectations move in response to FOMC announcements across different monetary policy measures.*

Overall, this set of results suggests limited consistency of respondents' expectations with the basic relationships suggested by mainstream macroeconomic models. Households expect FOMC decisions to affect future nominal interest rates but do not expect interest rates to further transmit to inflation, employment and output.

Robustness to Estimation Window Length

We pointed out before that our baseline sample stretches over 21 days before and after each FOMC meeting. Over the course of six weeks it is possible that other news that are relevant for the future course of economic variables are released. We therefore provide further evidence on the robustness to the length of control and treatment periods of the results on interest rate expectations that we obtain above.

Table 3.7 shows the effects of the three factors proposed by Swanson (2017) for different sample lengths. It shows that the effect presented in table 3.11 has the lowest magnitude of all possible sample lengths. The effect of the FFR factor reaches its peak during a symmetric time window of 5 days around FOMC meetings, with a coefficient of 1.885. For longer sample sizes, the effect shrinks quite smoothly. The table suggests that the announcement effect

Table 3.4: Baseline regression results: macroeconomic expectations

	(1) Interest 12m	(2) Unempl. 12m	(3) Stocks 12m	(4) Interest 12m	(5) Unempl. 12m	(6) Stocks 12m	(7) Interest 12m	(8) Unempl. 12m	(9) Stocks 12m	(10) Interest 12m	(11) Unempl. 12m	(12) Stocks 12m
Post-FOMC	-1.191**	0.454	0.0500	0.0803	0.0330	-0.103	0.115	0.00679	-0.108	0.0836	0.0324	-0.109
	(0.463)	(0.415)	(0.422)	(0.188)	(0.180)	(0.172)	(0.188)	(0.181)	(0.172)	(0.188)	(0.180)	(0.172)
Post-FOMC × Tightening	1.547***	-0.508	-0.187									
	(0.502)	(0.455)	(0.461)									
Post-FOMC × Δ Shadow Rate				0.472**	-0.000798	-0.0552						
				(0.195)	(0.184)	(0.176)						
Post-FOMC × Federal Funds Rate Factor							0.558***	-0.143	0.202			
							(0.187)	(0.174)	(0.172)			
Post-FOMC × Forward Guidance Factor							0.215	0.195	0.251			
							(0.210)	(0.192)	(0.191)			
Post-FOMC × LSAP Factor							0.506**	-0.412*	0.0303			
							(0.239)	(0.228)	(0.236)			
Post-FOMC × Monetary Policy Shock										0.440**	0.0371	0.361*
										(0.222)	(0.203)	(0.201)
Post-FOMC × Information Shock										0.482**	0.0827	0.486**
										(0.226)	(0.212)	(0.207)
Constant	33.15***	35.12***	40.36***	33.16***	35.12***	40.36***	33.15***	35.12***	40.36***	33.15***	35.12***	40.36***
	(0.0942)	(0.0903)	(0.0862)	(0.0943)	(0.0903)	(0.0863)	(0.0943)	(0.0903)	(0.0862)	(0.0943)	(0.0903)	(0.0863)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FOMC Meeting FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	31662	31661	31525	31662	31661	31525	31662	31661	31525	31662	31661	31525
Respondents	6910	6910	6880	6910	6910	6880	6910	6910	6880	6910	6910	6880
R^2	0.740	0.683	0.717	0.740	0.683	0.717	0.740	0.683	0.717	0.740	0.683	0.717

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.5: Regression results: baseline macroeconomic expectations

	(1) Inflation 12m	(2) Inflation 36m	(3) Inflation 12m	(4) Inflation 36m	(5) Inflation 12m	(6) Inflation 36m	(7) Inflation 12m	(8) Inflation 36m
Post-FOMC	-0.0857	-0.0701	-0.0376	-0.0625**	-0.0360	-0.0621**	-0.0373	-0.0617**
	(0.0680)	(0.0745)	(0.0294)	(0.0308)	(0.0295)	(0.0308)	(0.0294)	(0.0308)
Post-FOMC × Tightening	0.0579	0.00887						
	(0.0756)	(0.0814)						
Post-FOMC × Δ Shadow Rate			-0.00633	-0.0113				
			(0.0310)	(0.0338)				
Post-FOMC × Federal Funds Rate Factor					0.00105	-0.0102		
					(0.0303)	(0.0305)		
Post-FOMC × Forward Guidance Factor					-0.0333	-0.0471		
					(0.0319)	(0.0343)		
Post-FOMC × LSAP Factor					0.00990	-0.0261		
					(0.0382)	(0.0428)		
Post-FOMC × Monetary Policy Shock							-0.0279	-0.0594
							(0.0339)	(0.0362)
Post-FOMC × Information Shock							-0.0455	-0.0590
							(0.0340)	(0.0371)
Constant	4.223***	4.268***	4.223***	4.268***	4.223***	4.268***	4.223***	4.268***
	(0.0148)	(0.0155)	(0.0148)	(0.0155)	(0.0148)	(0.0155)	(0.0148)	(0.0155)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FOMC Meeting FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	29831	29652	29831	29652	29831	29652	29831	29652
Respondents	6677	6649	6677	6649	6677	6649	6677	6649
R^2	0.752	0.740	0.752	0.740	0.752	0.740	0.752	0.740

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.6: Baseline regression results: macroeconomic expectations

	(1) House Prices 12m	(2) House Prices 36m	(3) House Prices 12m	(4) House Prices 36m	(5) House Prices 12m	(6) House Prices 36m	(7) House Prices 12m	(8) House Prices 36m
Post-FOMC	-0.118	0.0109	-0.114***	-0.0882**	-0.109***	-0.0856**	-0.112***	-0.0865**
	(0.0818)	(0.0869)	(0.0344)	(0.0357)	(0.0345)	(0.0358)	(0.0344)	(0.0357)
Post-FOMC × Tightening	0.00535	-0.119						
	(0.0905)	(0.0947)						
Post-FOMC × Δ Shadow Rate			0.0273	0.0213				
			(0.0357)	(0.0375)				
Post-FOMC × Federal Funds Rate Factor					-0.0698**	-0.0596*		
					(0.0353)	(0.0351)		
Post-FOMC × Forward Guidance Factor					-0.0450	-0.0280		
					(0.0379)	(0.0378)		
Post-FOMC × LSAP Factor					0.121**	0.0625		
					(0.0469)	(0.0470)		
Post-FOMC × Monetary Policy Shock							-0.0592	-0.0455
							(0.0403)	(0.0401)
Post-FOMC × Information Shock							-0.0181	0.00542
							(0.0405)	(0.0417)
Constant	5.111***	4.938***	5.112***	4.937***	5.111***	4.937***	5.112***	4.938***
	(0.0173)	(0.0180)	(0.0173)	(0.0180)	(0.0173)	(0.0180)	(0.0173)	(0.0180)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FOMC Meeting FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	29987	29616	29987	29616	29987	29616	29987	29616
Respondents	6711	6654	6711	6654	6711	6654	6711	6654
R^2	0.708	0.696	0.708	0.696	0.708	0.696	0.708	0.696

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.7: FOMC announcement effects on 12-month ahead interest rate expectations for varying control and treatment window lengths

	(1) 4 Days	(2) 5 Days	(3) 6 Days	(4) 7 Days	(5) 8 Days	(6) 9 Days	(7) 13 Days	(8) 17 Days	(9) 21 Days
Post-FOMC	1.545*	0.777	0.314	0.169	0.379	0.326	0.233	0.238	0.115
	(0.923)	(0.701)	(0.546)	(0.477)	(0.406)	(0.357)	(0.263)	(0.213)	(0.188)
Post-FOMC × Federal Funds Rate Factor	1.198	1.885***	1.528***	1.162***	0.952**	1.033***	0.857***	0.590***	0.558***
	(0.846)	(0.665)	(0.500)	(0.439)	(0.386)	(0.355)	(0.274)	(0.208)	(0.187)
Post-FOMC × Forward Guidance Factor	0.379	0.435	0.745	0.493	0.477	0.375	0.582*	0.399	0.215
	(0.899)	(0.680)	(0.545)	(0.486)	(0.438)	(0.400)	(0.310)	(0.243)	(0.210)
Post-FOMC × LSAP Factor	-0.367	-0.599	-0.352	-0.223	-0.0930	-0.133	-0.0812	0.640**	0.506**
	(1.395)	(1.102)	(0.844)	(0.739)	(0.638)	(0.601)	(0.453)	(0.261)	(0.239)
Constant	33.07***	33.23***	33.02***	32.84***	32.95***	33.14***	33.31***	33.01***	33.15***
	(0.468)	(0.363)	(0.276)	(0.242)	(0.203)	(0.174)	(0.130)	(0.103)	(0.0943)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FOMC Meeting FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2772	4503	6587	8343	10265	12303	18605	26142	31662
Respondents	1330	2104	2934	3571	4224	4819	6019	6664	6910
R^2	0.824	0.824	0.816	0.805	0.801	0.795	0.774	0.754	0.740

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Note: Each column shows the results of estimating the regression in equation 3.1 on a symmetric window around the FOMC meeting. The respective length of treatment and control window is given in the column title (e.g. column 1 reports results for a symmetric window of 4 days before and after each FOMC meeting, giving a sample length of 8 days in total).

decays substantially.

However, in each of the columns in table 3.7, both control and treatment group vary. It is therefore also instructive to decompose the result of the baseline regression into sub-periods of the treatment period of 21 days while keeping the control window constant. Figure 3.3 shows the evolution of the marginal effect of the FFR factor throughout the treatment window and confirms the observation that the announcement effect may not be long-lived and stems from an immediate reaction in expectations within the first few days after an FOMC meeting.

3.5.2 Expectations About Personal Finances

We now turn to the effects of monetary policy announcements on the personal financial expectations of SCE respondents. Table 3.14 shows the results for three variables capturing personal financial expectations of individuals: the probability to lose one's job in the following 12 months, and the expectations about personal spending and overall household income for the following 12

Figure 3.3: Marginal effect of FFR factor on 12-month ahead interest rate expectations over the course of the treatment window and corresponding 95% confidence interval

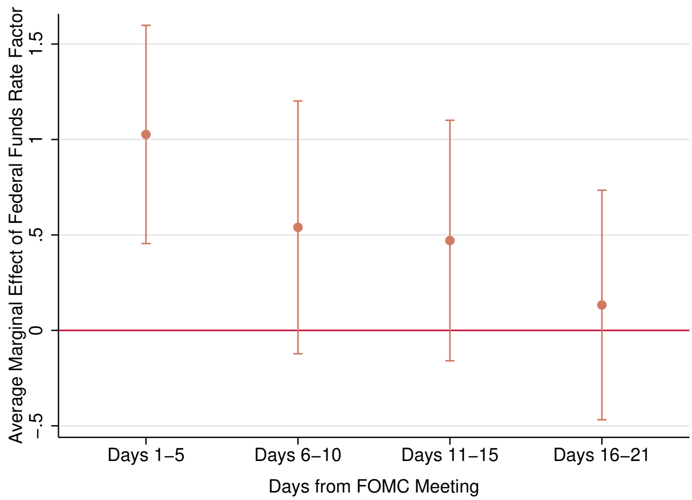

months. We apply the same baseline regression model as for the previous set of results.

In general, personal financial expectations hardly react to any of the monetary policy measures we consider. We will point out some of the stronger reactions, that are nonetheless insignificant. In case of a tightening, the treatment effect on personal job loss expectations is slightly above zero. For the remaining variables, reactions are small and insignificant throughout. We observe no effects of monetary policy announcements on spending plans of households. No measure of monetary policy leads to any significant change in the expectations about the overall household income of the respondent. The clear absence of effects on personal financial expectations stands in contrast to the experimental results obtained by Roth & Wohlfart (2020). Their results suggest that information about recession probabilities significantly impacts respondents' personal job loss and consumption growth expectations. We see this contrast as suggestive evidence that pieces of real world news about the macroeconomy are not as clear cut and straightforward to interpret as experimental interventions, where respondents are artificially fed with relevant information. Therefore our second key result is the following:

Result 2: *Personal financial expectations of respondents are unaffected by monetary policy announcements. Specifically, survey participants do not expect to adjust their spending behavior in response to monetary policy changes. (Surprise) easing or tightening of monetary policy also have no effect on expected household income.*

3.5.3 The "Taper Tantrum" and its Effects on Household Expectations

To better illustrate our findings from above, we look at a particularly relevant - also in terms of its media coverage - episode of monetary policy making in isolation, the so-called "Taper Tantrum" of 2013. The "Taper Tantrum" followed a series of communications made by the Federal Reserve in 2013 attempting to prepare the public and financial markets for a reduction in the pace of asset purchases. During the press conference following the FOMC meeting on June 18^{th} and 19^{th} 2013, then Chairman Ben Bernanke announced that conditional on further positive economic data in the months ahead, asset purchases could be reduced later in the year and halted over the course of 2014 (Bernanke 2013). The announcement surprised financial markets and received significant media attention. The Dow Jones Industrial Average fell by around 570 points or 4% between June 18^{th} and 20^{th} (Prial 2013). After the meeting on September 18^{th}, Bernanke reassured markets that asset purchases would continue (Park 2013a) before announcing their final scaling down as of January 2014 at the press conference on December 18^{th} 2013 (Park 2013b).

The large swings on financial markets in response to these various announcements are also reflected in the high-frequency identified surprises presented in Table 3.3. Both the Forward Guidance Factor and the Long Term Asset Purchases Factor move strongly in response to the meetings of June and September 2013. After establishing the effects of these monetary policy measures and others over the whole sample period, we zoom in on this sub-sample

Table 3.8: Baseline regression results: personal financial expectations

	(1) Lose Job 12m	(2) Spending 12m	(3) Income 12m	(4) Lose Job 12m	(5) Spending 12m	(6) Income 12m	(7) Lose Job 12m	(8) Spending 12m	(9) Income 12m	(10) Lose Job 12m	(11) Spending 12m	(12) Income 12m
Post-FOMC	-0.646	-0.0920	-0.135	-0.114	-0.0249	-0.0513	-0.108	-0.0244	-0.0555	-0.108	-0.0264	-0.0505
	(0.455)	(0.127)	(0.117)	(0.184)	(0.0544)	(0.0491)	(0.182)	(0.0545)	(0.0494)	(0.183)	(0.0544)	(0.0492)
Post-FOMC × Tightening	0.641	0.0810	0.102									
	(0.498)	(0.141)	(0.131)									
Post-FOMC × Δ Shadow Rate				0.00229	-0.00546	0.0296						
				(0.186)	(0.0569)	(0.0518)						
Post-FOMC × Federal Funds Rate Factor							0.126	-0.000534	-0.0312			
							(0.198)	(0.0536)	(0.0480)			
Post-FOMC × Forward Guidance Factor							-0.148	0.0368	0.0121			
							(0.191)	(0.0586)	(0.0533)			
Post-FOMC × LSAP Factor							-0.161	0.0524	-0.102			
							(0.245)	(0.0741)	(0.0649)			
Post-FOMC × Monetary Policy Shock										-0.168	0.0551	-0.00371
										(0.208)	(0.0613)	(0.0561)
Post-FOMC × Information Shock										0.121	-0.0825	0.0408
										(0.214)	(0.0610)	(0.0564)
Constant	13.98***	3.647***	3.637***	13.98***	3.647***	3.637***	13.98***	3.647***	3.638***	13.98***	3.646***	3.638***
	(0.0910)	(0.0273)	(0.0247)	(0.0910)	(0.0273)	(0.0247)	(0.0908)	(0.0273)	(0.0247)	(0.0911)	(0.0273)	(0.0247)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FOMC Meeting FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	17929	29452	29108	17929	29452	29108	17929	29452	29108	17929	29452	29108
Respondents	4114	6668	6590	4114	6668	6590	4114	6668	6590	4114	6668	6590
R^2	0.740	0.611	0.667	0.740	0.611	0.667	0.740	0.611	0.667	0.740	0.611	0.667

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

to provide a narrative-based illustration of our findings. Table 3.9 presents the results of regressions of the main macroeconomic and personal expectations on meeting-specific treatment indicators. The coefficients show the mean difference between expectations of those surveyed before the respective FOMC meeting and those surveyed in the days following the meeting after controlling for respondent and FOMC meeting fixed effects.

The press conference on June 19$^{\text{th}}$ led to substantial reactions among households. We estimate that this announcement significantly increased the expected probability of increasing interest rates over the 12 months following the announcement by 9.6 points. At the average probability respondents attached to rising interest rates *before* the meeting (30.9%) this corresponds to a jump in the expected probability of about 30%. Unemployment expectations were also affected significantly and increased by 7.8 points in response to the announcement. This corresponds to a 25% jump in the probability of increasing unemployment over the 12 months following the announcement. Notably, respondents do not react in their expectations about the stock markets or the inflation rate. Expectations about the respondents' personal probability to lose their job as well as their own spending or income were not affected.

The meeting on September 18$^{\text{th}}$ had less strong effects on interest and unemployment expectations. By announcing a delay of the tapering of asset purchases, respondents' interest rate expectations dropped by 3.8 points. Despite the announcement of an easier monetary policy, expectations of rising unemployment continued to increase. As mentioned in section 3.5.1, this announcement causes the effect of the LSAP Factor on unemployment expectations we observe in the baseline results. In the context of this episode, the sign of this effect makes sense. The easing of large scale asset purchases as measured by the LSAP Factor consisted of delaying the scaling down of purchases and was perceived to reveal a more pessimistic outlook of the economy by the Fed than previously assumed. The movement of unemployment expectations due to this specific FOMC meeting is therefore consistent with a central bank information channel (Jarociński & Karadi 2020).

This interpretation of the communication released after the FOMC meeting also explains the sign of the coefficient of LSAP purchases on unemployment

Table 3.9: Treatment Effects of Selected FOMC Meetings in 2013

	(1) Interest 12m	(2) Unempl. 12m	(3) Stocks 12m	(4) Inflation 12m	(5) Inflation 36m	(6) Lose Job 12m	(7) Spending 12m	(8) Income 12m
Post-FOMC × Meeting June 19th 2013	9.624***	7.845**	-0.313	-0.0540	-0.998	-1.837	-1.263	-1.148
	(2.742)	(3.447)	(3.766)	(0.478)	(0.616)	(1.877)	(1.139)	(0.973)
Post-FOMC × Meeting September 18th 2013	-3.828**	2.801*	-1.462	0.0205	0.272	0.236	-1.060**	-0.350
	(1.621)	(1.582)	(1.602)	(0.258)	(0.293)	(1.877)	(0.482)	(0.378)
Post-FOMC × Meeting December 18th 2013	4.812***	2.531	2.573	0.179	-0.0671	1.709	-0.702	0.301
	(1.765)	(1.929)	(1.626)	(0.282)	(0.344)	(1.323)	(0.536)	(0.438)
Constant	31.23***	35.59***	40.53***	4.879***	4.909***	14.53***	4.976***	2.268***
	(0.500)	(0.531)	(0.437)	(0.0734)	(0.0925)	(0.426)	(0.156)	(0.124)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FOMC Meeting FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	884	882	828	826	822	488	828	803
Respondents	345	345	326	327	325	194	324	315
R^2	0.807	0.750	0.748	0.827	0.782	0.816	0.719	0.797

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

that we observed in the baseline results. After the September meeting, expectations about personal spending growth dropped by about 1 percentage point. As the FOMC finally did announce a specific time table for the tapering of their asset purchases at the December meeting, interest rate expectations were sent upwards again, this time increasing by 4.8 percentage points, while no other expectations were significantly affected.

These results confirm that, while households are attentive to monetary policy, the range of expectations that are affected are rather limited. Not even the most publicised monetary policy decisions in our sample trigger substantial reactions in inflation expectations among survey respondents. We do, however, measure effects on unemployment expectations and personal spending growth that the rest of the FOMC meetings during our sample do not trigger. In contrast, the effects on interest rate expectations that we document in section 3.5.1, are observable throughout the sample. When excluding the year 2013 from our analysis, the effect of the FFR factor remains significant and of similar magnitude as the one presented in Table 3.11. The full set of results when excluding the year 2013 can be found in Appendix C.

3.6 Financial and Numerical Literacy

Next, we investigate whether more financially literate respondents exhibit different reactions to monetary policy news. First-time respondents of the SCE are asked to answer seven questions eliciting their numerical and financial literacy. These questions measure respondents' understanding of compounding interest, probability, risk and numerical reasoning (for the specific questions, see Appendix B). The survey's administrators use the answers given to these questions during the initial participation in the SCE and categorise respondents as highly numerically literate if they answer at least four of these questions correctly. We use the same classification. Since the questions do not only elicit numerical but also financial literacy we use this variable as a combined measure of numerical and financial literacy.

As a complementary measure we use respondents' exposure to financial decision making. Respondents are asked who in their household is largely responsible for financial decisions and based on this answer assigned to one of three categories: decisions are taken only or mostly by their partner, they are shared equally or taken only or mostly by themselves. We find this additional measure important since more responsibility in financial matters likely translates into more exposure to financial topics and potentially news about the economy – even for respondents with correct answers to the seven numerical and financial literacy questions. As the results of Binder (2019) suggest, exposure to financial decision making is important for the expectation formation of highly financially literate respondents. Table 3.10 shows how the two measures relate in our sample and underlines their complementary nature. Even among respondents whose partners handle most or all of financial decisions of the household, 65% score highly in the numerical and financial literacy questions. Nevertheless, respondents with more responsibility in the financial matters of the household tend to answer more numerical and financial literacy questions correctly. Between the respondents that share the responsibility equally with their partner and those that are mostly responsible themselves, we see no substantial difference in numerical and financial literacy scores. These descriptive statistics show that the two measures do not necessarily measure the same characteristic.

Table 3.10: Role in Financial Decision Making and Numerical/Financial Literacy (percentage of row total in parentheses)

	Numerical & Financial Literacy		
Who takes financial decisions?	Low	High	Total
Only/mostly spouse	304 (34.2%)	586 (65.8%)	890
Shared equally	3,231 (24.4%)	10,027 (75.6%)	13,258
Only/mostly respondent	2,965 (26.0%)	8,453 (74.0%)	11,418
Total	6,500 (25.4%)	19,066 (74.6%)	25,566

For ease of exposition we will present the results of this section graphically. Each figure presents the marginal effect of treatment on the outcome variable for a one standard deviation increase in the monetary policy measure, for each level of the respective factor variable. In the case of the dummy variable indicating a tightening of monetary policy, the figure presents the marginal effect of treatment in case the dummy variable takes the values 0 and 1 separately. Each point estimate is accompanied by the corresponding 95% confidence interval. In the following, we discuss the significance of the marginal effects as well as whether these marginal effects significantly differ across groups. For the latter purpose we report the results of F-tests of equality of marginal effects in the text.

Figure 3.4 shows these results for the effect of the role in the financial decision making process in the household on interest rate expectations. We plot the effects for the following four monetary policy measures for which we found the strongest effects in section 3.5.1: the dummy variable indicating a policy tightening, the shadow rate, the FFR Factor and the LSAP Factor.

Panel 3.4a shows the marginal effects of the tightening dummy across the three groups. The only sub group for which we find significant effects in case of easing/no change and tightening in the expected directions are those that are themselves responsible for financial decision making. The marginal effects for those respondents are also significantly different from each other (F-statistic of about 12, p-value of less than 0.01). Those whose partner is responsible for most or all financial decisions show no significant reaction to easing/no change of monetary policy and a significantly positive reaction in interest rate expectations to a tightening of policy. The point estimate is even larger than that for those carrying the main responsibility for financial decisions.

The respondents' reaction to a change in the shadow rate (Figure 3.4b) confirms that those with most exposure to financial decision making react the strongest to changes in monetary policy. The marginal effects of the three sub groups are all significantly different from each other at least at the 10% significance level.

For the effects interacted with the FFR Factor this is clearly not the case. All three groups have roughly similar point estimates, while only the marginal effects for the two groups with more exposure to financial decisions are significantly different from zero. Those two groups also exhibit similar reactions to the LSAP Factor. When excluding the year 2013, the reactions of all three groups to changes in the LSAP factor are closer to zero and insignificant, while the effects of the other three measures remain the same. The breakdown by financial decision making role also does not reveal any previously hidden reactions of any of the three sub groups on other variables considered in our baseline regressions.

The pattern becomes even clearer when considering the effects according to the respondents' level of numerical and financial literacy. Figure 3.5a confirms the finding from above. Only those with high literacy react significantly in both directions to the dummy variable - lowered expectations in case of easing/no change and increased interest rate expectations in case of a tightening. Unsurprisingly, the effects of a change in the shadow rate are consistent with this result. While the FFR Factor caused similar reactions across the three groups considered above, the breakdown by numerical and financial literacy reveals differential effects. Only the group of highly literate respondents reacts significantly to the policy surprise. However, the difference between the point estimates of the announcement effects for each group is not significantly different from zero (F-statistic of about 1.4, p-value of about 0.2). The reaction to the LSAP Factor is insignificant for both groups and remains so after excluding the year 2013.

In contrast, we find no robustly measured effects on other expectation variables that are different from zero. Figure 3.6 shows the announcement effects of the previously analyzed monetary policy measures on inflation expectations by numerical literacy category. Inflation Expectations of respondents react to none of the policy measures, just as in our baseline specifications. This means that the insignificant baseline effects on inflation expectations and personal financial conditions do not mask any substantial heterogeneities in announcement effects determined by numerical and financial literacy. Our third and last key result is therefore:

Result 3: *We find evidence that respondents with higher financial and numerical literacy react more strongly to monetary policy announcements in their interest rate expectations. Even respondents with high numerical and financial literacy or strong exposure to financial decision making show no substantially different reactions to other macroeconomic expectations nor predictions about their personal finances.*

3.7 Conclusion

In this chapter, we document causal effects of monetary policy announcements on household expectations about the economy in the United States. We compare responses to the Survey of Consumer Expectations given in the days before an FOMC meeting to those given afterwards and find that FOMC decisions robustly affect interest rate expectations of surveyed individuals. These effects are detectable using a diverse range of monetary policy measures. However, no other macroeconomic expectations are affected by FOMC decisions in the population as a whole. Additionally, the effect on interest rate expectations decays over the horizon of three weeks after an FOMC meeting.

The second key result that carries special importance for monetary policy making is the lack of announcement effects on personal financial expectations, such as spending or income expectations. We find that monetary policy, even in cases when interest rate expectations are affected strongly, is largely disconnected from the personal financial situation of survey participants. Beyond

Figure 3.4: Effect of Role in Financial Decision Making on Interest Rate Expectations

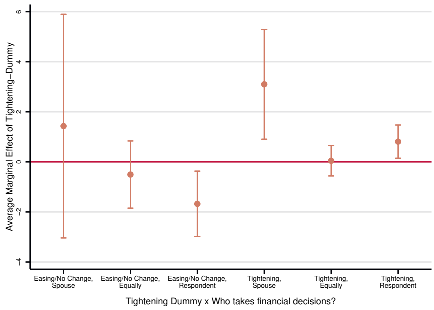

(a) Marginal Effect of Tightening Dummy by Financial Decision Making Category

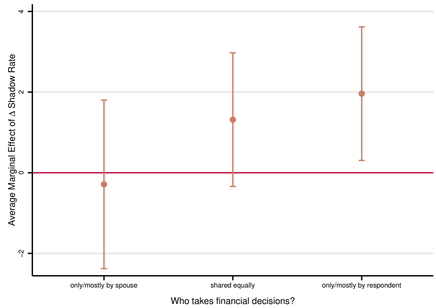

(b) Marginal Effect of Change in Shadow Rate by Financial Decision Making Category

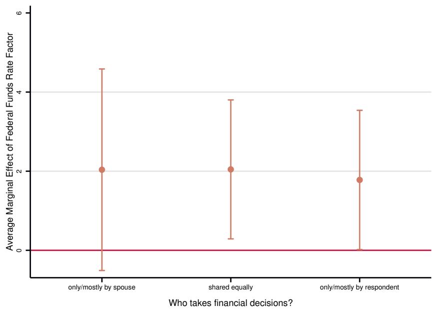

(c) Marginal Effect of FFR Factor by Financial Decision Making Category

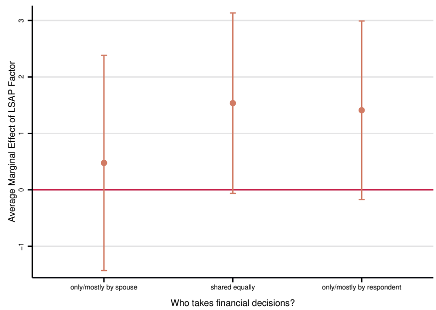

(d) Marginal Effect of LSAP Factor by Financial Decision Making Category

Note: Panel a) shows the marginal effect of both levels of the monetary policy tightening dummy variable by respondent characteristic. Panels b)-d) show the marginal effects of a one standard deviation increase in the monetary policy measure by respondent characteristic. All marginal effects are accompanied by a 95% confidence interval. The underlying regressions control for individual and FOMC meeting fixed effects. Standard errors are clustered at the respondent level.

Figure 3.5: Effect of Numerical Literacy on Interest Rate Expectations

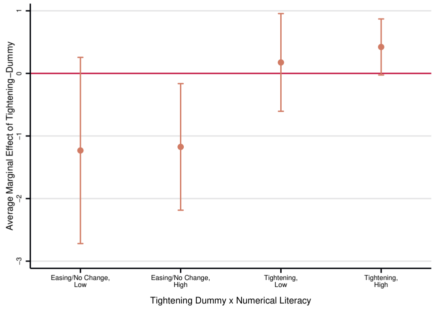

(a) Marginal Effect of Tightening Dummy by Numerical Literacy Category

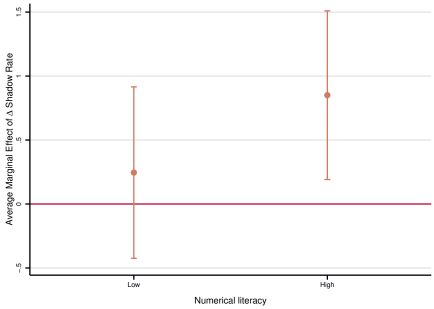

(b) Marginal Effect of Change in Shadow Rate by Numerical Literacy Category

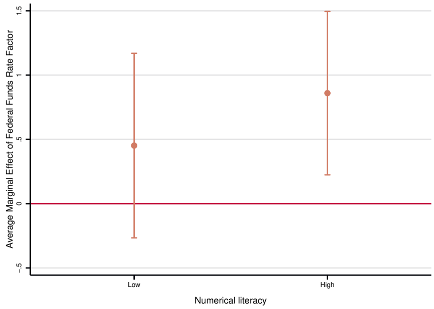

(c) Marginal Effect of FFR Factor by Numerical Literacy Category

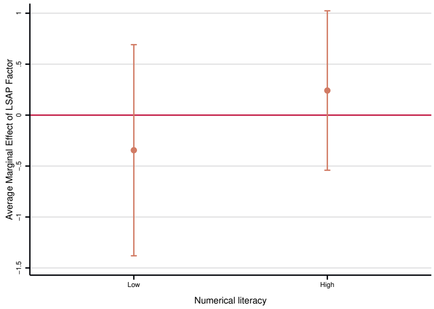

(d) Marginal Effect of LSAP Factor by Numerical Literacy Category

Note: Panel a) shows the marginal effect of both levels of the monetary policy tightening dummy variable by respondent characteristic. Panels b)-d) show the marginal effects of a one standard deviation increase in the monetary policy measure by respondent characteristic. All marginal effects are accompanied by a 95% confidence interval. The underlying regressions control for individual and FOMC meeting fixed effects. Standard errors are clustered at the respondent level.

Figure 3.6: Effect of Numerical Literacy on Inflation Expectations (12 months ahead)

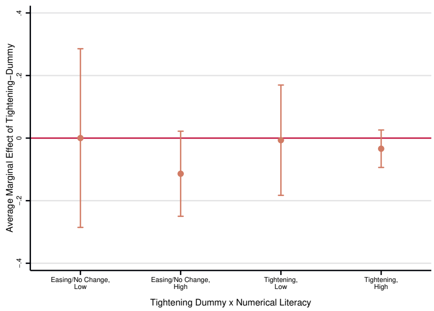

(a) Marginal Effect of Tightening Dummy by Numerical Literacy Category

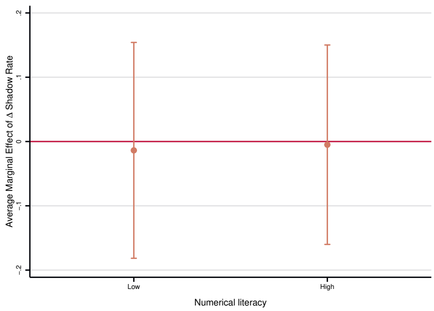

(b) Marginal Effect of Change in Shadow Rate by Numerical Literacy Category

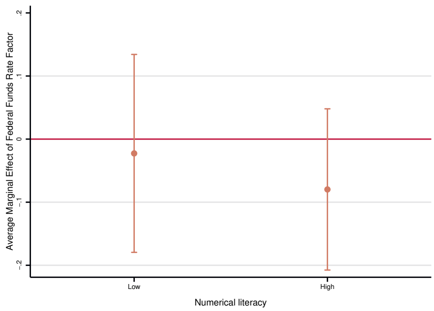

(c) Marginal Effect of FFR Factor by Numerical Literacy Category

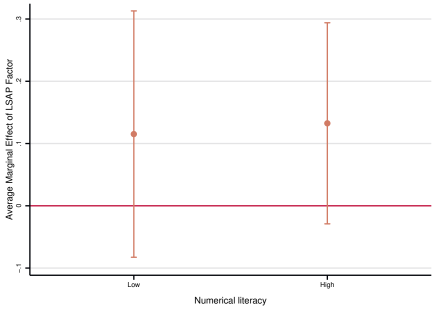

(d) Marginal Effect of LSAP Factor by Numerical Literacy Category

Note: Panel a) shows the marginal effect of both levels of the monetary policy tightening dummy variable by respondent characteristic. Panels b)-d) show the marginal effects of a one standard deviation increase in the monetary policy measure by respondent characteristic. All marginal effects are accompanied by a 95% confidence interval. The underlying regressions control for individual and FOMC meeting fixed effects. Standard errors are clustered at the respondent level.

the baseline results we explore response heterogeneity based on measures of financial or numerical literacy. Knowledge in these matters could make deciphering of FOMC decisions easier. We find some evidence of the importance of economic knowledge for announcement reactions. More numerically and financially literate respondents react more strongly to FOMC announcements in their interest rate expectations. Overall, however, they do not react on a wider range of expectations.

The experimental literature has identified rather strong effects of monetary policy news on household expectations while the evidence we provide using observational data does not reproduce these results. This could mean that the signals that real-world monetary policy news send are not strong enough to trigger the effects we observe in experimental settings. An alternative interpretation is that the Federal Reserve enjoys a high degree of credibility. The fact that expectations about unemployment and inflation, the two most important target variables of the Federal Reserve, do not react to changes in monetary policy, could be a byproduct of well anchored expectations and a resulting flat Phillips Curve. Under a steep Phillips Curve short-term inflation expectations could react to monetary policy even if they are anchored in the longer term. However, a flat Phillips Curve could mute the short-term effects on anchored expectations, which themselves may have contributed to the flattening of the Phillips Curve.

Our results are consistent with this, as not even the most salient FOMC announcements during our sample period exerted any effects on respondents' inflation expectations while having moderate effects on unemployment expectations. However, they are not easily reconciled with the strong effects from experimental evidence unless real world conditions, such as central bank credibility or the flatness of Phillips Curves, are part of the information set provided in experimental settings.

The common conclusion from both interpretations is that providing macroeconomic news in experimental settings may send fundamentally different signals than in the real world. Given that we cannot confirm or reject any of the two interpretations, the literature would benefit from more evidence about the effects of *real-world* macroeconomic news on expectations in the population, also in areas other than monetary policy. This would allow comparisons of different communication strategies and transmission mechanisms.

Appendix A Panel Conditioning

Panel conditioning in longitudinal surveys has received increased attention in recent years. It describes the problem that survey respondents that are repeatedly asked to participate in a survey might change their behavior because of their participation in that survey. An example could be a first time respondent in the Survey of Consumer Expectations who never informed herself about inflation. This respondent answers the SCE module on inflation expectations for the first time and realizes that she misses relevant information to form an expectation. Knowing that she will be asked the same question again in a month's time, she might start paying more attention to news about inflation in order to form a more accurate expectation the next time around.

Binder (2019) has shown that inflation expectations in the Survey of Consumer Expectations develop in a manner consistent with the mechanism described above. Respondents in their early rounds of participation consistently revise their inflation expectations downwards, irrespective of actual inflation dynamics. Even financially literate respondents make these revisions. The reason for the revisions seems to be increased attention to the topic by the respondents due to survey participation. Even though there is no indication in the survey description that average participation tenure should vary over time, we find that there is a significant difference in the distribution of survey tenure between those surveyed before and those surveyed after an FOMC meeting when using the entire sample. The null hypothesis of equal distributions is rejected in a Kolmogorov-Smirnov test with a p-value of 0.001. The difference vanishes when limiting the sample to respondents that have completed at least 7 survey rounds. We therefore analyze the impact of monetary policy announcements on deviations from a respondent's average expectation over 7 to 12 months.

By limiting our sample to respondents with at least 7 months of participation, we ensure that increasing attention to monetary policy news with longer survey participation does not induce a trend in the expectations that is unrelated to economic circumstances. The seven month threshold roughly coincides with the number of survey modules it takes for respondents to stop revising their expectations due to increasing tenure. Throughout our analysis we only compare responses of participants that fulfill this condition. Our results are, however, robust to modifying the threshold of survey participation. Whether or not we restrict the sample based on survey tenure, there are no structural differences in personal characteristics between respondents that answer the survey before or after FOMC meetings. Whitney-Mann-Wilcoxon tests with the null-hypothesis of equality of distributions of factor variables measuring education, financial literacy and income do not find a significant difference between the distributions of these variables before and after FOMC meetings.

Appendix B Numerical Literacy Questions

First-time respondents of the Survey of Consumer Expectations are asked the following seven questions eliciting their numerical and financial literacy:

- In a sale, a shop is selling all items at half price. Before the sale, a sofa costs $300. How much will it cost in the sale?
- Let's say you have $200 in a savings account. The account earns ten per cent interest per year. Interest accrues at each anniversary of the account. If you never withdraw money or interest payments, how much will you have in the account at the end of two years?
- In the BIG BUCKS LOTTERY, the chances of winning a $10.00 prize are 1%. What is your best guess about how many people would win a $10.00 prize if 1,000 people each buy a single ticket from BIG BUCKS?
- If the chance of getting a disease is 10 percent, how many people out of 1,000 would be expected to get the disease?
- The chance of getting a viral infection is 0.0005. Out of 10,000 people, about how many of them are expected to get infected?
- Imagine that the interest rate on your savings account was 1% per year and inflation was 2% per year. After one year, how much would you be able to buy with the money in this account? *Possible answers:* More than today, Exactly the same, Less than today
- Please tell me whether this statement is true or false: Buying a single company's stock usually provides a safer return than a stock mutual fund.

Appendix C Baseline Results 2014-2019

As we show in our baseline results section, three FOMC meetings in 2013 had strong effects on respondents' expectations. On the following pages, we present the regression results excluding all FOMC meetings of the year 2013.

Table 3.11: Baseline Regression Results 2014-2019: Interest Rate, Unemployment, Stock Market

	(1) Interest 12m	(2) Unempl. 12m	(3) Stocks 12m	(4) Interest 12m	(5) Unempl. 12m	(6) Stocks 12m	(7) Interest 12m	(8) Unempl. 12m	(9) Stocks 12m	(10) Interest 12m	(11) Unempl. 12m	(12) Stocks 12m
Post-FOMC	-0.911*	0.197	0.128	0.0484	-0.0960	-0.121	0.0714	-0.108	-0.132	0.0847	-0.114	-0.106
	(0.489)	(0.437)	(0.443)	(0.193)	(0.184)	(0.176)	(0.193)	(0.184)	(0.176)	(0.193)	(0.184)	(0.176)
Post-FOMC × Tightening	1.165**	-0.349	-0.299									
	(0.528)	(0.477)	(0.482)									
Post-FOMC × Δ Shadow Rate				0.361*	-0.00136	-0.0521						
				(0.201)	(0.191)	(0.181)						
Post-FOMC × Federal Funds Rate Factor							0.602***	-0.159	0.210			
							(0.188)	(0.175)	(0.174)			
Post-FOMC × Forward Guidance Factor							0.273	0.195	0.282			
							(0.211)	(0.193)	(0.192)			
Post-FOMC × LSAP Factor							-0.554	-0.505	-0.292			
							(0.378)	(0.342)	(0.340)			
Post-FOMC × Monetary Policy Shock										0.388*	0.0591	0.344*
										(0.222)	(0.204)	(0.202)
Post-FOMC × Information Shock										0.519**	-0.181	0.495**
										(0.244)	(0.228)	(0.222)
Constant	33.24***	35.10***	40.37***	33.24***	35.09***	40.37***	33.23***	35.10***	40.37***	33.24***	35.09***	40.37***
	(0.0964)	(0.0922)	(0.0880)	(0.0964)	(0.0922)	(0.0880)	(0.0965)	(0.0922)	(0.0880)	(0.0965)	(0.0922)	(0.0880)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FOMC Meeting FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	29897	29898	29900	29897	29898	29900	29897	29898	29900	29897	29898	29900
Respondents	6571	6572	6571	6571	6572	6571	6571	6572	6571	6571	6572	6571
R^2	0.741	0.681	0.718	0.741	0.681	0.718	0.741	0.682	0.718	0.741	0.681	0.718

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.12: Baseline Regression Results 2014-2019: Inflation

	(1) Inflation 12m	(2) Inflation 36m	(3) Inflation 12m	(4) Inflation 36m	(5) Inflation 12m	(6) Inflation 36m	(7) Inflation 12m	(8) Inflation 36m
Post-FOMC	-0.0925	-0.0854	-0.0411	-0.0627**	-0.0393	-0.0604*	-0.0437	-0.0648**
	(0.0707)	(0.0770)	(0.0303)	(0.0313)	(0.0303)	(0.0313)	(0.0304)	(0.0313)
Post-FOMC × Tightening	0.0610	0.0270						
	(0.0782)	(0.0841)						
Post-FOMC × Δ Shadow Rate			-0.00174	-0.000269				
			(0.0320)	(0.0348)				
Post-FOMC × Federal Funds Rate Factor					0.00312	-0.0105		
					(0.0305)	(0.0308)		
Post-FOMC × Forward Guidance Factor					-0.0353	-0.0483		
					(0.0320)	(0.0345)		
Post-FOMC × LSAP Factor					0.00491	0.000151		
					(0.0563)	(0.0623)		
Post-FOMC × Monetary Policy Shock							-0.0294	-0.0573
							(0.0340)	(0.0363)
Post-FOMC × Information Shock							-0.0550	-0.0715*
							(0.0370)	(0.0393)
Constant	4.185***	4.224***	4.185***	4.224***	4.185***	4.224***	4.185***	4.224***
	(0.0152)	(0.0157)	(0.0152)	(0.0157)	(0.0152)	(0.0157)	(0.0152)	(0.0157)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FOMC Meeting FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	28178	27982	28178	27982	28178	27982	28178	27982
Respondents	6349	6314	6349	6314	6349	6314	6349	6314
R^2	0.751	0.740	0.751	0.740	0.751	0.740	0.751	0.740

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.13: Baseline Regression Results 2014-2019: House Price Growth

	(1) House Prices 12m	(2) House Prices 36m	(3) House Prices 12m	(4) House Prices 36m	(5) House Prices 12m	(6) House Prices 36m	(7) House Prices 12m	(8) House Prices 36m
Post-FOMC	-0.0526	0.0400	-0.105***	-0.0931**	-0.104***	-0.0930**	-0.100***	-0.0913**
	(0.0848)	(0.0915)	(0.0352)	(0.0368)	(0.0353)	(0.0367)	(0.0355)	(0.0369)
Post-FOMC × Tightening	-0.0609	-0.157						
	(0.0929)	(0.0991)						
Post-FOMC × Δ Shadow Rate			0.0202	0.0188				
			(0.0366)	(0.0390)				
Post-FOMC × Federal Funds Rate Factor					-0.0714**	-0.0693**		
					(0.0356)	(0.0353)		
Post-FOMC × Forward Guidance Factor					-0.0491	-0.0304		
					(0.0380)	(0.0380)		
Post-FOMC × LSAP Factor					0.189***	0.180**		
					(0.0681)	(0.0735)		
Post-FOMC × Monetary Policy Shock							-0.0634	-0.0422
							(0.0404)	(0.0402)
Post-FOMC × Information Shock							-0.00507	-0.0249
							(0.0439)	(0.0453)
Constant	5.059***	4.888***	5.059***	4.887***	5.060***	4.888***	5.059***	4.887***
	(0.0177)	(0.0184)	(0.0176)	(0.0184)	(0.0177)	(0.0184)	(0.0176)	(0.0184)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FOMC Meeting FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	28296	27956	28296	27956	28296	27956	28296	27956
Respondents	6374	6323	6374	6323	6374	6323	6374	6323
R^2	0.708	0.696	0.708	0.696	0.709	0.697	0.708	0.696

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.14: Baseline Regression Results 2014-2019: Personal Financial Expectations

	(1) Lose Job 12m	(2) Spending 12m	(3) Income 12m	(4) Lose Job 12m	(5) Spending 12m	(6) Income 12m	(7) Lose Job 12m	(8) Spending 12m	(9) Income 12m	(10) Lose Job 12m	(11) Spending 12m	(12) Income 12m
Post-FOMC	-0.597	0.0231	-0.148	-0.115	0.000701	-0.0488	-0.0998	-0.00135	-0.0475	-0.0908	-0.00568	-0.0433
	(0.466)	(0.132)	(0.124)	(0.190)	(0.0552)	(0.0501)	(0.189)	(0.0554)	(0.0503)	(0.189)	(0.0556)	(0.0505)
Post-FOMC × Tightening	0.571	-0.0278	0.120									
	(0.509)	(0.146)	(0.137)									
Post-FOMC × Δ Shadow Rate				-0.0215	-0.0180	0.0313						
				(0.187)	(0.0575)	(0.0536)						
Post-FOMC × Federal Funds Rate Factor							0.136	0.0136	-0.0246			
							(0.200)	(0.0538)	(0.0481)			
Post-FOMC × Forward Guidance Factor							-0.145	0.0343	0.0198			
							(0.191)	(0.0589)	(0.0534)			
Post-FOMC × LSAP Factor							-0.361	-0.0524	-0.229**			
							(0.371)	(0.109)	(0.0969)			
Post-FOMC × Monetary Policy Shock										-0.178	0.0433	-0.00287
										(0.209)	(0.0616)	(0.0562)
Post-FOMC × Information Shock										0.148	-0.0341	0.0469
										(0.232)	(0.0645)	(0.0604)
Constant	13.96***	3.569***	3.713***	13.96***	3.569***	3.713***	13.96***	3.569***	3.713***	13.96***	3.569***	3.714***
	(0.0939)	(0.0278)	(0.0252)	(0.0939)	(0.0278)	(0.0252)	(0.0938)	(0.0278)	(0.0252)	(0.0939)	(0.0278)	(0.0252)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FOMC Meeting FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	16893	27772	27470	16893	27772	27470	16893	27772	27470	16893	27772	27470
Respondents	3903	6329	6261	3903	6329	6261	3903	6329	6261	3903	6329	6261
R^2	0.740	0.613	0.666	0.740	0.613	0.666	0.740	0.613	0.666	0.740	0.613	0.666

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4

Heterogeneity of Phillips Curve Slopes in Monetary Unions: Simulations [1]

[1] This chapter is based on Schuffels, Johannes, Clemens Kool, Lenard Lieb, & Tom van Veen. 2022. "Is the slope of the Euro Area Phillips Curve steeper than it seems? Heterogeneity and identification."

4.1 Introduction

The aim of this chapter is to direct attention to an endogeneity problem that arises when estimating Phillips Curve slopes using regional data and aggregations thereof: systematic slope heterogeneity. Theoretical models show that in monetary unions, the central bank's effort to stabilize a region-specific demand shock is a function of that region's Phillips Curve slope (among other variables) (Brissimis & Skotida 2008; Lee 2009). The reason is that the Phillips Curve moderates the effect of output gap variations on inflation. As a result, the output gap in regions with flatter Phillips Curves is more volatile than in regions with steep Phillips Curves. When estimating the union-wide slope of the Phillips Curve in reduced form regressions, the variation of the output gap that is due to the region-specific deviation from the union-wide slope enters the error term and biases the output gap coefficient towards zero unless one controls for slope heterogeneity. This result has implications for the ongoing debate about changing Phillips Curve slopes over time: according to our predictions obtained from theory, omitting controls for slope heterogeneity in reduced form estimations could spuriously suggest flattening slopes, while in reality regional slopes have only diverged.

Recently, several studies have shown the value of analyses using regional data in a monetary union as a *solution* to an endogeneity problem when estimating the slope of the union-wide Phillips curve: the endogenous response of monetary policy to shocks to inflation or output (McLeay & Tenreyro 2019; Bharadwaj & Dvorkin 2020). Since regional demand shocks cannot be fully stabilized by the union's central bank, they can serve as an exogenous shock to the output gap that affects inflation only through the Phillips Curve relationship. Both of the above cited analyses assume that the monetary union consists of structurally homogeneous economies, also with respect to the slope of the Phillips Curve. Under this strong assumption the conclusions drawn in these papers for identification of the slope are correct. However, the bias due to the underlying heterogeneity across regions has not received sufficient attention. Many recent analyses that use regional data for Phillips Curve estimations assume homogeneous slopes (McLeay & Tenreyro 2019; Eser et al. 2020; Hazell et al. 2020). Only Kapetanios et al. (2020) explicitly tackle cross-regional heterogeneity and take stock of how different (pooled) estimators perform in estimating Phillips Curve slopes using regional data. They do allude to the endogeneity problem outlined above. Their main objective however, is the modelling of common correlated effects, a separate problem from the one analyzed in this chapter. Additionally, their approach contains no quantification and differentiation between different sources of endogeneity, but rather a comparison of different estimators on the same underlying (real world) data.

There is ample empirical evidence that suggests heterogeneous slopes across countries and regions. Amberger & Fendel (2017) estimate country-specific reduced form Phillips Curve slopes in the Euro Area and find substantial heterogeneity. Imbs et al. (2011) give a potential explanation for this heterogeneity: substantial heterogeneity in the duration of nominal rigidities across industries leads to biased estimates of the Phillips Curve slope at the aggregate level. Cross-country heterogeneity in sectoral or firm structure could therefore drive slope heterogeneities across countries. Additionally, explanations for a flattening of the slope could apply to countries to different degrees. Section 4.2 describes how this chapter relates to the broader debate on a flattening of the slope. The aim of this chapter is to illustrate and isolate the effects that heterogeneous, regional Phillips Curve slopes have on estimating common monetary union slopes using aggregate and regional data.

The problem is illustrated using a simple New Keynesian monetary union model. It isolates the impact of slope heterogeneity on estimates of the union-wide slope and presents modifications to the estimation strategy that allow identification of the slope. The simulations show that both aggregate (i.e. union-level) as well as panel regressions suffer from this bias. From an econometric perspective this problem is not new. With the mean group estimator, Pesaran et al. (1999) presented a solution to biases due to slope heterogeneity more than 20 years ago. More recently, Breitung & Salish (2020) provide an estimation strategy tailor-made for *systematic* slope heterogeneity.

The rest of the chapter is structured as follows. Section 4.2 places the topic of this chapter in the wider context of the ongoing debates about the flattening of Phillips Curve slopes. Section 4.3 presents the monetary union model and illustrates the endogeneity problem by contrasting regression results on simulated data for a union with homogeneous slopes with a union with heterogeneous slopes. Section 4.4 concludes.

4.2 The Debate on the Flattening of the Phillips Curve Slope

For many years already, an observed flattening of the Phillips Curve slope has raised the attention of researchers and policy makers, leading to a wide range of mechanisms explaining the observed phenomenon. One leading explanation is that monetary policy has become much more successful at stabilizing the economy over the last decades Boivin & Giannoni (2006). This success relies on the endogenous response of monetary policy to movements in inflation or slack, thereby introducing a negative relationship between the two, while the Phillips Curve relationship implies a positive co-movement. Therefore, as McLeay & Tenreyro (2019) point out, it becomes impossible to observe the Phillips Curve relationship when analyzing aggregate data, similarly to the futile attempt of trying to recover demand curves from equilibrium price realisations. This explanation has also been echoed by policy makers, see e.g. Bullard (2018). Others suggest that there are important non-linearities in price and wage setting that standard analyses do not take into account (Lindé & Trabandt 2019) or that household inflation expectations can explain the missing deflation after the Great Recession (Coibion & Gorodnichenko 2015b). More recently, Lombardi et al. (2020) suggest that the falling unionisation in advanced economies has weakened the link between tight labor markets and wages and prices.

Many of these explanations are obtained from structural models or estimated in well identified settings. However, reduced form estimations of the Phillips Curve remain quite prominent in the literature, see for example Ball & Mazumder (2019) or Eser et al. (2020). In its most basic form this means running a regression of a measure of inflation on its own lags, a measure of slack and a measure of expected inflation. The recent innovation to focus on panels of regions instead of aggregate union-wide time series has already partially solved the problem that McLeay & Tenreyro (2019) describe - but, importantly, under the assumption of slope *homogeneity* across regions. The issue of slope heterogeneity remains largely untackled even though potential remedies exist as we show in this chapter. With this contribution we want to direct attention to the potential impact slope heterogeneity can have in theory and to what extent this matters in reduced form analyses on data coming from the Euro Area. As the empirical analysis in chapter 5 shows, slope heterogeneity can explain a portion of possible changes in the slope of the Euro Area Phillips Curve. However, our analysis does not intend to disprove any of the aforementioned mechanisms driving changes in the structure of the economy. Indeed, any change in slope heterogeneity could be driven by regions being subject to the mechanisms described above to different degrees. Therefore, we see our analysis as complementary to the many structural explanations put forward in the literature.

4.3 Simulation

As previewed in the introduction, we start the analysis by discussing the bias due to slope heterogeneity in a controlled setting. We simulate data according to a simple New-Keynesian model of a monetary union that is composed of two regions that can only differ in their idiosyncratic slope of the Phillips Curve. Varying the parameter that determines slope heterogeneity allows us to illustrate its theoretical impact on reduced form estimates of the union-wide slope.

4.3.1 Monetary Union Model

The monetary union model is characterized by a set of equations. Inflation in country i at time t is described by the following equation:

$$\pi_{it} = \beta E_t \pi_{it+1} + \kappa_i x_{it} + u_{it} \tag{4.1}$$

where κ_i is the Phillips Curve slope of country i and defined as $\kappa_i = \kappa + \eta_i$. η_i is the country-specific deviation from the union-wide Phillips Curve slope κ. In our model, η_i is distributed symmetrically around zero, yielding an expected value of κ_i across all regions equal to the union-wide slope κ. To isolate the effects of slope heterogeneity, differing values of κ_i are the only source of cross-region heterogeneity in this model. The supply shock is defined as an AR(1) process of the form $u_{it} = \rho_u u_{it-1} + \epsilon_{it}$ with the random variable $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$. The region's output gap is denoted by x_{it} and develops according to a standard IS equation:

$$x_{it} = E_t x_{it+1} - \sigma\left(\bar{i}_t - E_t \pi_{it+1}\right) + r_{it} \tag{4.2}$$

where σ measures the intertemporal elasticity of substitution and $\bar{i}_t$ is the union-wide nominal interest rate. The demand shock is defined as an AR(1) process of the form $r_{it} = \rho_r r_{it-1} + \mu_{it}$ with the random variable $\mu_{it} \sim N(0, \sigma_\mu^2)$.

The aggregate variables are defined as

$$\bar{\pi}_t = \Sigma_{i=1}^N w_i \pi_{it} \tag{4.3}$$

$$\bar{x}_t = \Sigma_{i=1}^N w_i x_{it}, \tag{4.4}$$

where w_i is country i's relative size in the monetary union, with $\Sigma_{i=1}^N w_i = 1$. The model is closed by an interest rate rule on the union-level of the following form

$$\bar{i}_t = \lambda_\pi \bar{\pi}_t + \lambda_x \bar{x}_t, \tag{4.5}$$

where λ_π is the central bank's weight on inflation stabilization and λ_x the corresponding weight on output stabilization. The target inflation rate is set to 0.

To illustrate and isolate the endogeneity problems arising due to slope heterogeneity we simplify the model by switching off biases stemming from the relative size of cost-push shocks and the persistence of both types of shocks. McLeay & Tenreyro (2019) show that the relative variance of cost-push (σ_u^2) versus demand shocks (σ_r^2) determines the bias. In case $\sigma_u^2 \gg \sigma_r^2$ the downward bias to the slope coefficient is large. Additionally, the persistence of cost-push and demand shocks, measured by ρ_u and ρ_r can blur identification of the slope. In order to achieve the most advantageous conditions for identification of the slope we set $\sigma_u^2 = 0$, $\sigma_r^2 > 0$ as well as $\rho_u = \rho_r = 0$. Imposing persistence of zero of the demand shocks simplifies the identification by making inflation and output gap expectations irrelevant. Since shocks don't have any effects beyond the period they occur in, the expected value for the subsequent period's inflation and output gap under rational expectations is always equal to its target, in this case zero for both variables. The demand shocks across the different regions are independent random variables. For simplicity, the monetary union is composed of two members of equal size ($w_1 = w_2 = 0.5$) and data is simulated for 1000 periods.[2] The model abstracts from many important features of monetary unions, such as trade between members or common correlated shocks. However, the purpose of this chapter is to isolate the impact of slope heterogeneity on the estimation of κ. It should be kept in mind that the bias that is demonstrated in what follows adds to the various biases that have been identified in the studies mentioned above.

4.3.2 Identification Under Homogeneous Slopes

Before introducing slope heterogeneity this section evaluates the identification problem under the assumptions described above as well as slope *homogeneity*: $\kappa_i = \kappa = 0.5$ for all i. The model is then simulated for the two regions. The purpose of this section is to provide a benchmark scenario in which identification of the slope is guaranteed in order to illustrate the contrast with results obtained under heterogeneous slopes.

Due to the absence of cost-push shocks the identification is straight-forward. Only demand shocks affect the economy and the central bank does not face a trade-off between controlling inflation or the output gap: a positive demand shock will affect output and inflation in the same direction. Co-movement between inflation and output is determined entirely by the Phillips Curve relationship. Figure 4.1a clearly shows the positive relationship between the output gap and inflation in the aggregate simulated data. A regression of the

[2] The remaining parameters are set to the following numerical values: $\sigma = 2, \lambda_\pi = 1.5, \lambda_x = 0.5, \sigma_{r_i}^2 = 0.7$.

Figure 4.1: Inflation and Output Gap in Simulations: Homogeneous Slopes

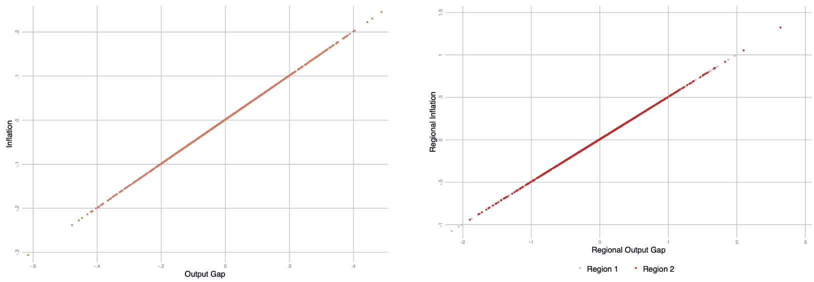

(a) Aggregate Data

(b) Regional Data

union-wide inflation rate on the output gap identifies the slope of the Phillips Curve, as column 1 in Table 4.1 shows.

Under these circumstances using regional data provides no added value. Both regions have exactly the same structure and are affected only by demand shocks. The regressions provide the same results, irrespective of the inclusion of period fixed effects as columns 2 and 3 in Table 4.1 show. Any of the three regressions identify the union-wide slope of the Phillips Curve.

4.3.3 Identification Under Heterogeneous Slopes

Next, we introduce slope heterogeneity. The union-wide (average) slope remains at $\kappa = 0.5$ but the country-specific deviations are no longer zero. We set $\eta_1 = 0.3$ and $\eta_2 = -0.3$. The aim of this exercise is to identify the union-wide slope κ under these conditions. Figures 4.2a and 4.2b show the aggregate and regional data for the simulated monetary union with slope heterogeneity. The picture looks very different from Figure 4.1. While the inflation-output gap realizations in regional data allow identification when estimating the slope separately by region, the aggregate relationship is blurry. Aggregate output gap and inflation are not helpful to identify the union's Phillips Curve slope. By contrast, the regional data in panel 4.2b clearly trace out the two regional Phillips Curve slopes. However, the relevant question for our analysis

Table 4.1: Regression Results on Simulated Data

	Homogeneous Slopes			Heterogeneous Slopes				
	(1) Aggregate OLS	(2) Pooled OLS	(3) Mean Group	(4) Aggregate OLS	(5) Pooled OLS	(6) Pooled OLS	(7) Mean Group	(8) Augmented GLS
$\bar{x}$	0.500***			-0.007				
	(0.000)			(0.019)				
x		0.500***	0.500***		0.368***	0.424***	0.500*	0.500***
		(0.000)	(0.000)		(0.010)	(0.000)	(0.300)	(0.000)
z								-1.652***
								(0.000)
Constant	0.000	0.000	0.000	0.000	0.001	0.139	0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.004)	(0.004)	(0.025)	(0.000)	(0.000)
Period Effects	No	No	CCE	No	No	PFE	CCE	No
Observations	1000	2000	2000	1000	2000	2000	2000	2000
R^2	1.000	1.000		0.000	0.651	0.968		

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Note: (1) and (4) regress aggregate inflation ($\bar{\pi}$) on the aggregate output gap ($\bar{x}$). (2) and (5) are pooled OLS regressions of regional inflation (π) on regional output gap (x). (3) and (7) are Mean Group regressions of regional inflation on regional output gap and common correlated effects (CCE), i.e. monthly average values of all regressors and the independent variable across all countries. (6) is an OLS regression of regional inflation on regional output gap and period fixed effects. (8) is a FGLS regression of regional inflation on regional output gap and an augmenting regressor (z) proposed by Breitung & Salish (2020).

is whether the union-wide slope will be identified in reduced form regressions of regional realizations of inflation on output gaps. The following subsections will depart from the purely graphical analysis of the data by deriving the bias in reduced form estimations from the monetary union model.

Figure 4.2: Inflation and Output Gap in Simulations: Heterogeneous Slopes

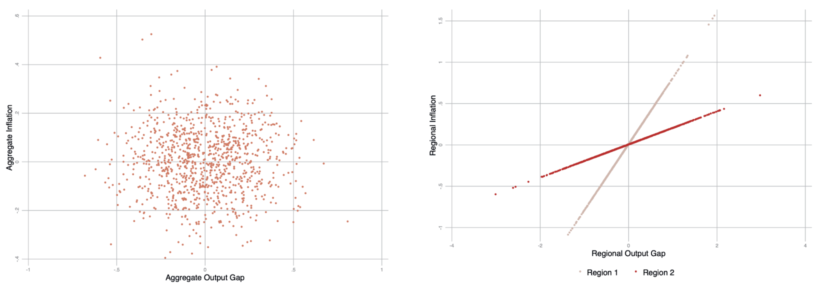

(a) Aggregate Data **(b)** Regional Data

Aggregate Data

As Figure 4.2a suggests, the slope of the aggregate relationship is very flat and imprecisely estimated compared to the results under homogeneous slopes. The only difference in the data between columns 1 and 4 in Table 4.1 is slope heterogeneity. The only parameter that differs between the two simulations is η_i. Slope heterogeneity alone led to a substantial flattening of the estimated slope. The bias can also be clearly linked analytically to slope heterogeneity. In column 4 the following regression was estimated

$$\bar{\pi}_t = \alpha + \kappa \bar{x}_t + v_t, \tag{4.6}$$

while the underlying data generating process can be expressed as the weighted average of each region's Phillips Curve:

$$\bar{\pi}_t = \Sigma_{i=1}^{N} w_i \left((\kappa + \eta_i) x_{it} + u_{it} \right). \tag{4.7}$$

Since κ is homogeneous across regions, equation 4.7 can be rewritten as

$$\bar{\pi}_t = \kappa \bar{x}_t + \bar{u}_t + \Sigma_{i=1}^{N} w_i \eta_i x_{it}. \tag{4.8}$$

Comparing equations 4.6 and 4.8, we therefore get the error term v_t

$$v_t = \Sigma_{i=1}^{N} w_i \eta_i x_{it}. \tag{4.9}$$

As explained above, the cost-push shocks in this simulation exercise are switched off. Therefore $\bar{u}_t$ is a constant equal to zero and can thus be dropped from the error term. The bias to $\hat{\kappa}$ depends on the correlation between $\bar{x}_t$ and $\Sigma_{i=1}^{N} w_i \eta_i x_{it}$. The estimator for $\hat{\kappa}$ can be written as

$$\hat{\kappa} = \frac{Cov(\bar{\pi}_t, \bar{x}_t)}{Var(\bar{x}_t)} = \kappa + \delta_{agg}, \tag{4.10}$$

where δ_{agg} is the coefficient in an OLS regression of the omitted variable $\Sigma_{i=1}^{N} w_i \eta_i x_{it}$ on $\bar{x}$. As the simulations show, the correlation between these two variables is strongly *negative*. The estimated coefficients for $\hat{\kappa}$ in the aggregate Phillips Curve estimation from 200 simulations with 500 periods each have a mean value of -0.023 with a standard deviation of 0.024. The theoretical value derived from the model in 4.3.1 is 0.022 (see Appendix A for the derivation of the bias for the aggregate regression).

Slope heterogeneity means that changes in the output gap have different effects on inflation in the two regions. If a positive demand shock hits region 1, the effect of that demand shock on regional and therefore aggregate inflation will be stronger than that of a demand shock of the same magnitude in region 2. For that reason and despite the equal weight both regions have in the interest rate rule (Equation 4.5), shocks in region 1 are stabilized to a larger degree than those in region 2. Previous studies, such as Brissimis & Skotida (2008) and Lee (2009), have shown this asymmetric stabilization in more complex models of monetary unions with cross-region heterogeneity.

How does this translate into a negative correlation with aggregate output? Due to the relatively strong reaction of monetary policy to demand shocks in region 1, region 2 experiences more volatile output gaps and inflation deviations from aggregate. Whenever the central bank stabilizes a positive demand shock in region 1 by substantially increasing the interest rate, output in region 2 falls. Due to its equal weight in the monetary union, the effect of the output reaction in region 2 on aggregate output will be relatively large. Since the effect of that output change on inflation is relatively mild compared to region 1, the central bank will accept these comparably larger output gap fluctuations in region 2. Precisely due to the central bank's relative disregard for demand shocks in region 2, there is a clear positive correlation between the output gap in region 2 and the aggregate output gap.

Figure 4.3a plots regional output gaps against the aggregate output gap. The reason for the sign of the correlation between $\bar{x}_t$ and $\eta_2 x_{2t}$ being negative is the negative sign of η_2. Figure 4.3b plots the aggregate output gap against the omitted variable v_t. The correlation between $\bar{x}_t$ and $\Sigma_{i=1}^{N} w_i \eta_i x_{it}$ is strongly negative.[3]

Regional Data

As McLeay & Tenreyro (2019) have shown, under persistent shocks and the existence of cost-push shocks regional data in a monetary union can alleviate the biases to the estimated Phillips Curve slope. It is especially powerful in controlling for the endogenous policy response since region-specific demand shocks cannot be fully stabilized by the central bank and therefore affect inflation through the Phillips Curve. Is this also true for the slope heterogeneity bias demonstrated above? As pointed out earlier, the raw data in Figure 4.2b suggest that cross-regional heterogeneity is clearly detectable. However, the aim of the analysis is to find out whether in reduced form regressions we can identify the union-wide slope κ. In column 5 in Table 4.1 we report results from the following pooled regression of regional inflation on the regional output gap

$$\pi_{it} = \alpha + \kappa x_{it} + e_{it} \tag{4.11}$$

where e_{it} is the error term. The regression yields a coefficient of the output gap of 0.368, much closer to the union-wide slope of 0.5 than when using aggregate data but nonetheless with a strong downward bias. To understand the source of this bias we derive the error term:

[3] The size of each region's economy, measured by w_i, plays a similar role as the region's Phillips Curve slope. A larger relative size of the economy also leads to less variation in the output gap as any shock to output in a big region will, everything else being equal, have larger effects on aggregate inflation than demand shocks in small regions. However, in order to illustrate the issue of slope heterogeneity, this mechanism is switched off in the model by assuming equally sized regions.

$$e_{it} = \eta_i x_{it} \tag{4.12}$$

Clearly, the regressor in equation 4.11 and the error term are correlated as they both contain x_{it}. Analogous to the aggregate case, the overall correlation is negative due to the region with $\eta_i < 0$ which experiences larger variations in the output gap, resulting in a downward bias of $\hat{\kappa}$. However, some of the variation in the regional output gaps x_{it} is being stabilized by the central bank, blurring the Phillips Curve relationship (McLeay & Tenreyro 2019). Period-fixed effects, or equivalently, regressions in terms of deviations from the union-wide aggregate, can control for this union-wide reaction of the interest rate to regional output gap fluctuations. Defining the regional deviation of a regional variable from its union-wide aggregate as $\tilde{y}_{it} = y_{it} - \bar{y}_t$, we estimate the following regression in column 6:

$$\tilde{\pi}_{it} = \alpha + \kappa \tilde{x}_{it} + h_{it} \tag{4.13}$$

The output gap coefficient now rises to 0.424 with a remaining bias equal to -0.076. Controlling for the endogenous response of monetary policy to regional output gap fluctuations has thus decreased the bias to $\hat{\kappa}$. The error term becomes

$$h_{it} = \eta_i x_{it} - \kappa \bar{x}_t \tag{4.14}$$

while the resulting bias - now exclusively due to slope heterogeneity - can be retrieved by estimating the following regression:

$$\eta_i x_{it} - \kappa \bar{x}_t = \delta_{reg} \tilde{x}_{it} + \gamma_{it}. \tag{4.15}$$

In 200 separate simulations of 500 periods each, δ_{reg} has a mean value of -0.0775 with a standard deviation of 0.0039. Accordingly, $\hat{\kappa}$ has a mean value of 0.4225. The intuition goes as follows and is closely related to the endogeneity in the aggregate data. The omitted variable, $\eta_i x_{it} - \kappa \bar{x}_t$, which can also be expressed as $\eta_i x_{it} - \bar{\pi}_{it}$, measures the region's idiosyncratic deviation from the aggregate inflation rate that is due to the region's deviation from the aggregate Phillips Curve slope. This deviation due to slope heterogeneity is correlated to the total region-specific deviation from aggregate output and here the reasoning resembles that in the aggregate case: The steeper the regional Phillips Curve, i.e. the larger η_i, the smaller the regional deviation from aggregate output will be because the central bank will not allow that region's demand shocks to play out their full effect on inflation in order to minimize aggregate inflation's deviations from target. Figure 4.4a plots the omitted variable, $\eta_i x_{it} - \kappa \bar{x}_t$, against the main regressor, $\bar{x}_t$. While the correlation is only negative in region 2 and positive in region 1, in a pooled setting the negative correlation dominates due to the wider dispersion of values in region 2. To further illustrate the fact that a steeper Phillips Curve slope compresses the deviations from the target output gap and therefore also the omitted variable, the kernel densities of the omitted variable for both regions are plotted in Figure 4.4b.

All previous results come from simulated data for a given parametrisation of the model regarding the heterogeneity of slopes presented in section 4.3.1. Next, we show how the bias develops for different degrees of heterogeneity. In Figure 4.5 we plot the theoretical values of δ_{agg} and δ_{reg} for different values of η_i, while leaving the model unchanged otherwise. η_i is always symmetrically distributed around $\kappa = 0.5$. As shown in section 4.3.2, if $\eta_i = 0$ for all i, the bias to both aggregate and regional coefficients is zero. The bigger the regional deviations from the aggregate slope, the bigger the biases. For small deviations, regional data can still help come reasonably close to identifying the slope, while the bias becomes substantial even in regional data with levels of heterogeneity assumed in the baseline parametrisation of the model. In aggregate data, already small levels of heterogeneity lead to substantial underestimations of the slope of the aggregate Phillips Curve. The graph shows clearly that growing heterogeneities among members of a currency union could - according to this simple model - offer an explanation for an observed flattening of the Phillips Curve when estimating a common coefficient using aggregate or panel data, while in reality the curve has only flattened for some members and steepened for others.[4]

Controlling for Slope Heterogeneity

There are several ways to identify the coefficient on the output gap despite slope heterogeneity. In this last part of the simulation exercise, we apply different techniques that deal specifically with slope heterogeneity and contrast these results with those obtained from aggregate or panel OLS regressions. In particular two solutions are examined. The first is a mean group estimator with common correlated effects (Pesaran et al. 1999) and has recently been applied to regional Phillips Curve estimations by Kapetanios et al. (2020). The authors estimate Phillips Curve slopes from U.S. state-level data using mean group estimators augmented by common correlated effects, i.e. average values across all panel groups of all variables in the regression for each time period. This approach takes into account slope heterogeneity by averaging over the state-specific slopes. Breitung & Salish (2020) confirm that mean group estimators are unbiased under systematic slope heterogeneity but point out that they are inefficient. They propose a GLS regression that is augmented by the regressor z_{it}:

$$z_{it} = x_{it} \left(\frac{1}{T} \Sigma_{t=1}^{T} x_{it}^2 - \frac{1}{NT} \Sigma_{i=1}^{N} \Sigma_{t=1}^{T} x_{it}^2 \right) \tag{4.16}$$

This approach specifically tackles the systematic nature of slope heterogeneity that is likely to be present when estimating Phillips Curves. A key pattern in the above simulation exercise is that output gaps in regions with

[4] Note that the pattern of growing biases to coefficients coming from reduced form regressions as a consequence of slope heterogeneity does not depend on the symmetric nature of heterogeneity assumed in the simulations or Figure 4.5. Any differences in regional slopes lead to differential implicit weights of member regions in the central bank's optimal monetary policy, and thus a bias to the coefficient in reduced form estimations.

flat Phillips Curves are much more volatile than those with steep curves (see Figure 4.4b). The augmenting regressor z_{it} captures this pattern in order to control for slope heterogeneity. The significance of this coefficient therefore also serves as a test for the existence of slope heterogeneity. If there is no systematic slope heterogeneity, the difference between second moments of the main regressor and their average across all panel groups shouldn't differ across panel groups and the coefficient on z should not be significantly different from zero.

In columns 3, 7 and 8 of Table 4.1, the two methods are applied to the simulated data under slope homogeneity and heterogeneity. Column 3 shows that the mean group estimator performs equally well under slope homogeneity as OLS estimators. It also recovers the area-wide slope under slope heterogeneity (column 8). The augmented GLS regression results in the same point estimate but is estimated much more precisely, which is in line with the result on the relative efficiency of the two estimators provided by Breitung & Salish (2020).

Figure 4.3: Aggregate, Regional and Weighted Regional Output Gaps

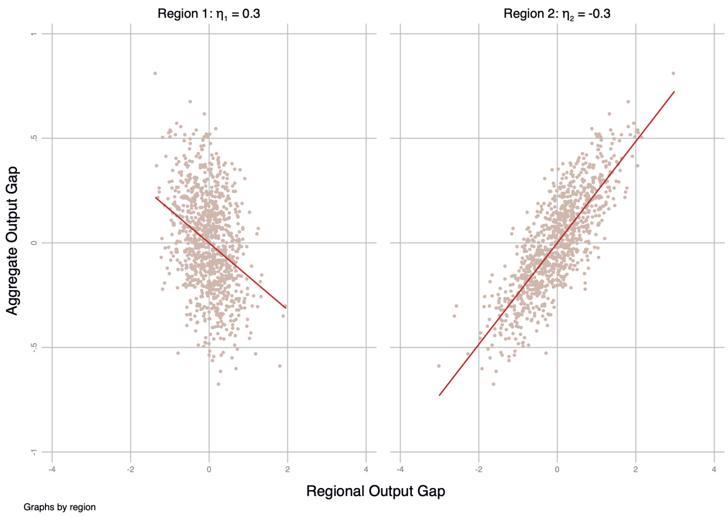

(a) Aggregate and Regional Output Gap by Region and Linear Fit

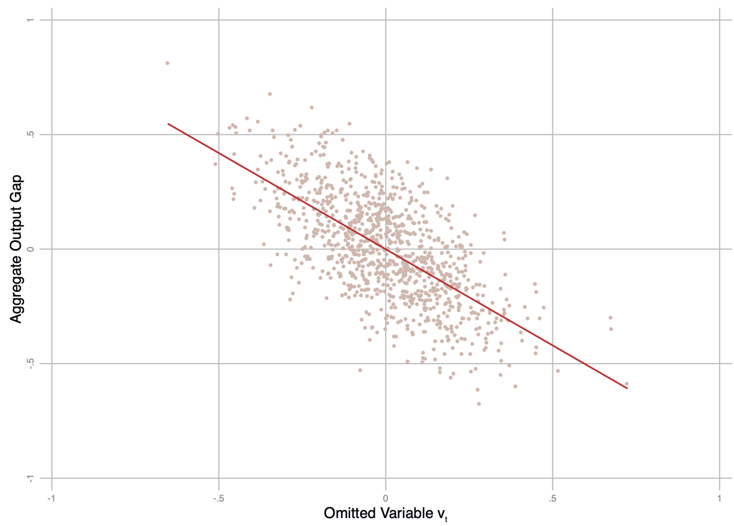

(b) Aggregate Output Gap and Omitted Variable

Figure 4.4: Omitted Variable $(\eta_i x_{it} - \kappa \bar{x}_t)$ in Regional Data

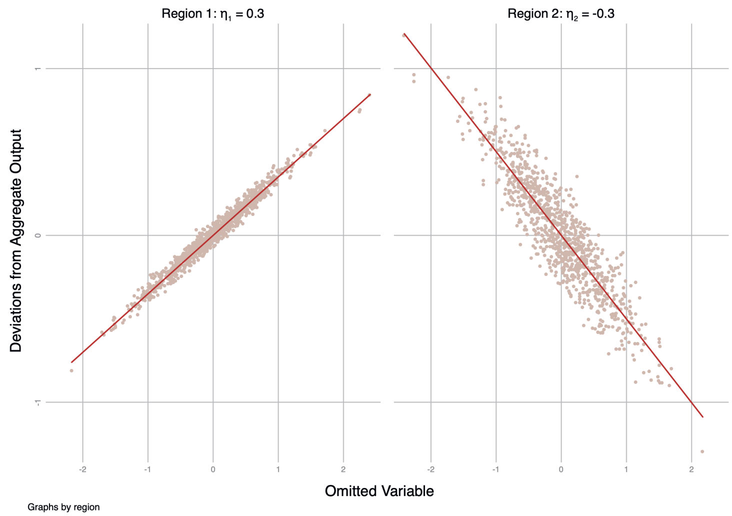

(a) Omitted Variable and Main Regressor by Region

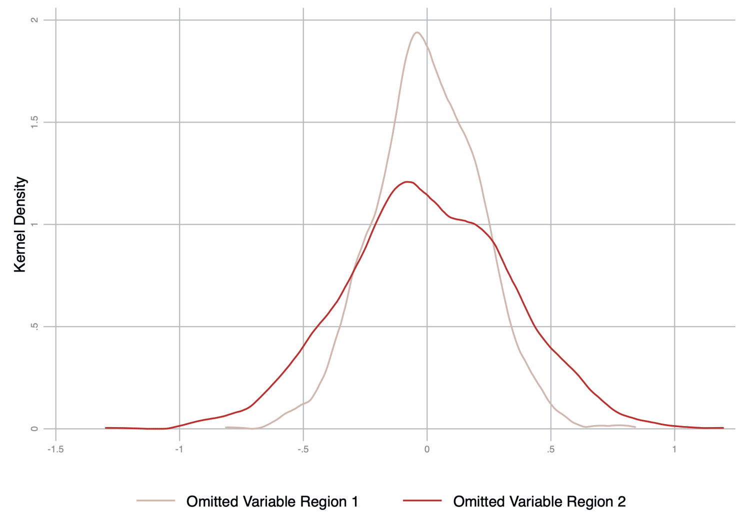

(b) Kernel Density of Omitted Variable by Region

Figure 4.5: Evolution of Bias with Varying Degrees of Slope Heterogeneity

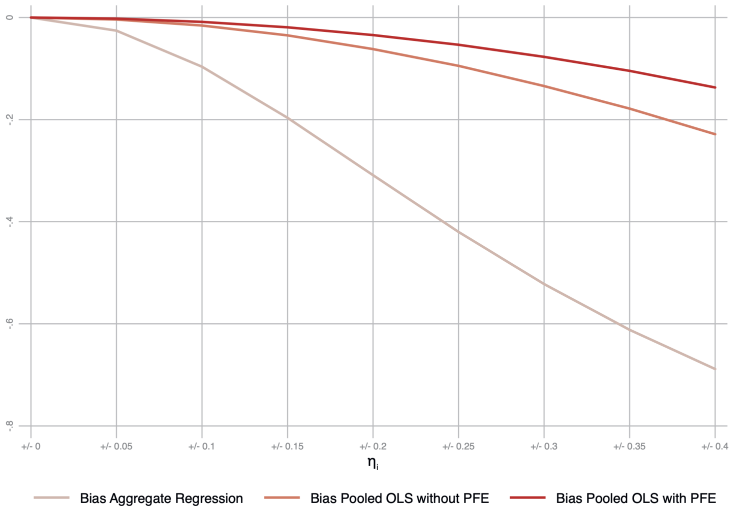

4.4 Conclusion

In this chapter we illustrate the impact of heterogeneous slopes on reduced form estimations of the Phillips Curve slope in a monetary union. If regional slopes differ, any attempt to estimate a union-wide slope either on aggregate or pooled regional data will suffer from substantial omitted variable bias. The reason is that the slope of the Phillips Curve changes the implicit weight a union member gets in the central bank's monetary policy rule. In regions with steep Phillips Curves, demand shocks will be stabilized to a larger degree than the mere size of the economy would justify as these demand shocks would impact aggregate union inflation to a larger degree than demand shocks in regions with flat Phillips Curves. The more pronounced the heterogeneity, the larger the bias to coefficients from aggregated or pooled estimations.

Our results can have important implications for reduced form analyses of the Phillips Curve relationship. We show theoretically that accounting for slope heterogeneity in a monetary union can raise the estimated coefficient on unemployment substantially. There are many recent examples of analyses that omit such controls and may therefore conflate trends of changing heterogeneity among sub-union entities with changes in the slope of the Phillips Curve, see for example Eser et al. (2020) for the Euro Area, Hazell et al. (2020) for an analysis of US state-level data and McLeay & Tenreyro (2019) for an analysis of US city-level data. The question whether reduced form analyses on Euro Area data may be biased due to slope heterogeneity is tackled in the following chapter.

Appendix A Derivation Omitted Variable Bias

Since expectations are zero in every period, the model can be easily simplified to the following two expressions for the output gap in both regions in terms of the demand shocks:

$$x_1 = \frac{1}{s}\left(r_1(\sigma^{-1} + s_1) - r_2 s_1\right) \tag{4.17}$$

$$x_2 = \frac{1}{s}\left(r_2(\sigma^{-1} + s_2) - r_1 s_2\right) \tag{4.18}$$

where $s = \lambda_x + \lambda_\pi \kappa + \sigma^{-1}$, $s_1 = \lambda_x w_2 + \lambda_\pi \kappa_2 w_2$ and $s_2 = \lambda_x w_1 + \lambda_\pi \kappa_1 w_1$. The coefficient δ_{agg} in a regression of $\Sigma_{i=1}^{N} w_i \eta_i x_{it}$ on $\bar{x}$ is therefore:

$$\begin{aligned}\delta_{agg} = &\Big\{\sigma_{r_1}^2\Big[\left(w_1\eta_1\left(s_1+\sigma^{-1}\right) - w_2\eta_2 s_2\right)\left(w_1\left(s_1+\sigma^{-1}\right) - w_2 s_2\right)\Big] \\ &+\sigma_{r_2}^2\Big[\left(w_2\eta_2\left(s_2+\sigma^{-1}\right) - w_1\eta_1 s_1\right)\left(w_2\left(s_2+\sigma^{-1}\right) - w_1 s_1\right)\Big]\Big\} \\ &\Big\{\sigma_{r_1}^2\Big[w_1\left(s_1+\sigma^{-1}\right) - w_2 s_2\Big]^2 \\ &+\sigma_{r_2}^2\Big[w_2\left(s_2+\sigma^{-1}\right) - w_1 s_1\Big]^2\Big\}^{-1}\end{aligned} \tag{4.19}$$

5

Heterogeneity of Phillips Curve Slopes in Monetary Unions: Evidence from the Euro Area [1]

[1] This chapter is based on Schuffels, Johannes, Clemens Kool, Lenard Lieb, & Tom van Veen. 2022. "Is the slope of the Euro Area Phillips Curve steeper than it seems? Heterogeneity and identification."

5.1 Introduction

In the previous chapter we illustrated a potential endogeneity problem that arises when measuring the slope of a monetary union's Phillips Curve in reduced form regressions in the presence of slope heterogeneity among the union's members. The analysis in the previous chapter is based on simulated data generated by a simple monetary union model that yields stylized results on potential biases to estimators of the output gap coefficient in reduced form Phillips Curve estimations.

The purpose of this section is to evaluate to what degree the patterns illustrated before matter in real world applications. We use the Euro Area as a case study. To that end, we first provide evidence on the existing degree of heterogeneity in country-specific inflation-unemployment trade-offs. The methodology to do so produces estimates of the impact of exogenous changes in country-level unemployment on inflation (Barnichon & Mesters 2021). Equipped with a gauge of the existing degree of heterogeneity, we move on to estimate different reduced form specifications on a panel of 10 Euro Area economies.

By estimating reduced form regressions following the method proposed by Breitung & Salish (2020), we test whether the observed degree of heterogeneity is sufficiently large to bias unemployment coefficients when heterogeneity is uncontrolled for.[2] This exercise shows that controlling for slope heterogeneity in Euro Area panel regressions of core inflation on unemployment (and other controls) can steepen the estimated slope. However, the augmenting regressor controlling for slope heterogeneity stays insignificant. By contrast, a subsample analysis reveals that the estimator detects slope heterogeneity in Euro Area since 2009, masking a stronger steepening of the slope if heterogeneity is uncontrolled for.

The rest of the chapter is organized as follows. Section 5.2 provides evidence on country-specific estimates of the inflation-unemployment trade-off in

[2] As output gap estimates are only available at annual frequency in the Euro Area, we use the unemployment rate as a proxy of the existing slack in the economy in all analyses of this section, relying on the relatively stable relationship between the two variables first documented by Okun (1962). This means that the theoretical sign of the Phillips Curve relationship reverses as unemployment and inflation should be negatively correlated.

the Euro Area. Equipped with this measure of the existing heterogeneity, section 5.3 presents evidence on the degree to which the observed heterogeneity biases can be obtained from reduced form regressions. Variation in the size of the bias over time is discussed in section 5.4. Section 5.5 concludes.

5.2 Heterogeneity in the Inflation-Unemployment Trade-off

To assess whether country-specific heterogeneity in Phillips Curve slopes exists in the Euro Area, reduced form evidence will not help us: country-specific regressions of inflation on unemployment may be misspecified as we cannot control for period fixed effects and panel regressions only recover a (potentially biased) union-wide estimate of the slope. Instead, we estimate the country-specific size of the inflation-unemployment trade-off using a methodology proposed by Barnichon & Mesters (2021). To avoid biases due to the endogenous reaction of monetary policy to demand shocks and due to the effects of cost-push shocks, Barnichon & Mesters (2021) exploit monetary policy surprises as exogenous changes in slack that affect inflation - and therefore allow the researcher to observe the relation between the two variables.

The key difference to reduced-form approaches at the country-level is that we identify co-movement in inflation and unemployment due to exogenous variations in unemployment. Applying this methodology at the country-level in the Euro Area can therefore expose the existing heterogeneity in country-level estimates of the inflation-unemployment trade-off. It should be noted that these estimates are not necessarily identical to estimates of the slope of the Phillips Curve. Due to the lack of country-level data on quantitative inflation expectations in the Euro Area, we can only quantify the total change in inflation due to a given monetary policy-induced change in unemployment.

This total trade-off can include second-round effects that run through inflation expectations. There are two cases in which the Phillips Multiplier would be equal to the slope: i.) expectations are fully anchored at the target or ii.)

expectations are entirely backward-looking.[3] This section therefore serves as a gauge of the dispersion of country-specific inflation-unemployment trade-offs around the Euro Area aggregate according to a common methodology that allows identification.

Specifically, we estimate a sequence ($h = 12, 13, .., 36$) of the following regression:

$$\Sigma_{j=0}^{h}\pi_{t+j} = \psi_h \Sigma_{j=0}^{h} u_{t+j} + X_t'\gamma_h + \epsilon_{t+h} \tag{5.1}$$

where $\Sigma_{j=0}^{h}\pi_{t+j}$ is the cumulative inflation rate from date t to date $t+h$, ψ_h the Phillips Multiplier at horizon h, $\Sigma_{j=0}^{h}u_{t+j}$ is the cumulative unemployment rate and X_t a vector of control variables, namely lagged unemployment and inflation. The cumulative unemployment rate $\Sigma_{j=0}^{h}u_{t+j}$ is instrumented by high frequency identified monetary policy surprises θ_t following the methodology of Jarociński & Karadi (2020). Specifically, we use the change in the 3-month ahead overnight indexed swaps (OIS) in a thirty minute window around the publication of the press release and the press conference after a monetary policy decision of the ECB Governing Council on days in which the OIS and the Euro Area stock market index EURO STOXX 50 moved in opposite directions. Jarociński & Karadi (2020) show that under these circumstances the actual monetary policy shock outweighs the effect of the publication of the central bank's outlook on future economic conditions.

We run these local projections for the ten founding members of the Euro Area and Greece as well as the aggregate Euro Area (changing composition). Inflation is measured with the core inflation rate (HCPI excluding energy and food). The sample runs from January 1999 to March 2019.[4] We estimate the Phillips Multiplier for horizons between 1 and 3 years to capture the full effects of a monetary policy shock after the transmission lag has passed. The results from this exercise will show to what degree the reaction of inflation to demand

[3] In the latter case, we would be controlling for inflation expectations through the inclusion of lagged inflation in equation 5.1.

[4] For Italy and Greece, the sample starts in January 2001

Figure 5.1: Phillips Multipliers: percent change in core inflation in response to monetary policy change that increases unemployment by 1pp

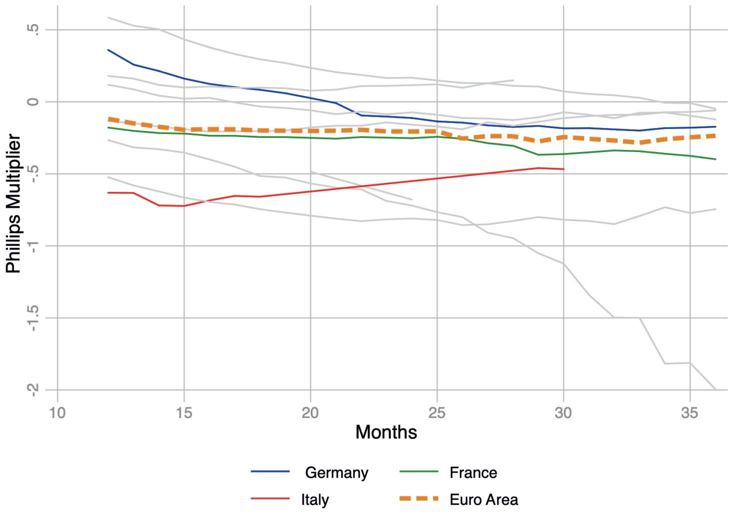

Note: The series of equations mentioned in equation 5.1 is run for each country individually. Control variables are 36 lags of core inflation and unemployment. Point estimates at individual horizons are excluded from the graph if 68% Anderson-Rubin confidence intervals are infinite or include gaps (Anderson & Rubin 1949).

shocks induced by monetary policy surprises differs i) across individual Euro Area economies and ii) from the union wide multiplier.

Figure 5.1 shows the Phillips Multiplier estimates between 12 and 36 months after the monetary policy shock.[5] The estimate of the multiplier for the Euro Area aggregate is just below -0.1 at 12 months after the impact of the shock, and drops to about -0.2 after 15 months and further to -0.25 after around 25 months. These estimates are roughly in line with the results in Eser et al. (2020).

[5] Due to the lag in the transmission of monetary policy, the Phillips Multiplier is initially indeterminate, we therefore only report results at horizons between 12 and 36 months.

To assess heterogeneity, we focus on the country-specific estimates of the Phillips Multiplier. For most countries, the estimated multiplier is quite stable between 12 and 36 months after the monetary policy shock. Only Austria's estimate drops throughout the estimation horizon. Table 5.1 shows the mean and median multiplier estimate between 12 and 36 months after the shock for each country as well as 68% confidence intervals corresponding to the median point estimate. The vast majority of point estimates is below zero with the exception of Belgium and Spain. It should be pointed out that the estimate for Belgium is falling steadily over the 3 year horizon and turns negative in the last months. For Spain, we cannot estimate the multiplier beyond a horizon of 28 months. Germany, France, the Netherlands and Portugal are closely clustered around the Euro Area wide multiplier. Austria and Ireland are the only clear outliers below the union wide estimate. Overall, these estimates do not suggest that the reaction of inflation to exogenous, monetary policy induced shocks to unemployment differs widely across the early members of the Euro Area. Austria and Ireland together account for less than 5% of the Euro Area's GDP. Additionally, those estimates that are further from the Euro Area aggregate multiplier tend to be estimated less precisely. Individual graphs of the estimated Phillips Multipliers by country including confidence intervals can be found in appendix A.

One prediction from the previous simulation exercise is that countries with steeper slopes should experience less volatility in their output gaps as the central bank will not tolerate changes in the output gap due to their outsized effect on union-wide inflation. As we use unemployment as a proxy for the output gap, Figure 5.2 plots the Phillips Multipliers estimated in this section against the standard deviation of the country's unemployment rate over the same time frame. There is a slight (yet insignificant) positive correlation between the two variables.

Figure 5.2: Mean Phillips Multiplier and Standard Deviation of Unemployment

Note: Phillips Multipliers are estimated according to the methodology proposed by Barnichon & Mesters (2021). The graph shows the mean estimate for each country between 12 and 36 months after the monetary policy shock. The underlying data for both variables runs from 1998 to 2019.

Table 5.1: Mean and Median Phillips Multiplier between 12 and 36 months after monetary policy shock and 68% Anderson-Rubin confidence intervals corresponding to the median Phillips Multiplier estimate

	Point Estimates		68% CI	
	Mean	Median	Upper	Lower
Austria	-0.876	-0.720	-0.248	-1.321
Belgium	0.202	0.165	0.248	0.087
Germany	-0.041	-0.113	0.021	-0.238
Spain	0.114	0.109	0.731	-0.062
France	-0.281	-0.253	-0.136	-0.365
Ireland	-0.762	-0.793	-0.545	-2.482
Italy	-0.626	-0.653	-0.138	-10.472
Netherlands	-0.054	-0.074	0.134	-0.371
Portugal	-0.147	-0.159	0.013	-0.347
Euro Area	-0.218	-0.205	-0.091	-0.323

Note: We do not report any results for Greece as Anderson-Rubin confidence intervals are either infinite or with gaps at all horizons. Finland is excluded from the table as multipliers and confidence intervals can only be estimated for two periods. The point estimate closest to zero is -0.485 and the corresponding confidence interval ranges from 0.265 to -32.638.

5.3 Reduced Form Evidence

The previous section equips us with unbiased estimates of the inflation-unemployment trade-off in the Euro Area founding member states and a first gauge of the degree of heterogeneity present among those countries. We found that while there are three countries whose (imprecisely estimated) Phillips Multiplier differs strongly from the union-wide estimate, the remaining country-specific estimates are clustered closely around the union-wide estimate. The aim of this section is to analyze whether this existing heterogeneity could bias reduced form panel regressions of inflation on a measure of economic slack (and other controls) at the sub-union level, a type of analysis that is being applied frequently (see, e.g. Eser et al. (2020), Hazell et al. (2020) or McLeay & Tenreyro (2019)). Contrary to the previous section, the aim of the reduced form analysis is to obtain estimates of the *union-wide* slope of the Phillips Curve.[6] We contrast results of regressions that omit slope heterogeneity controls with regressions that include those controls using the approach introduced by Breitung & Salish (2020). We showed in section 4.3.3 that this methodology is fully capable of detecting and controlling for slope heterogeneity. Additionally, it is more efficient than a mean group estimator and therefore our method of choice. As for the Phillips Multiplier estimates, we analyze the trade-off between core inflation and unemployment. Additionally, we control for (qualitative) household inflation expectations using the European Commission's business and consumer survey.[7] Previous research shows the informational value of household inflation expectations (Coibion & Gorodnichenko 2015b) and their usefulness in Phillips Curve estimations (McLeay & Tenreyro 2019). We also include for six lags of core inflation and month-fixed effects to control for seasonality. The sample runs from 1999 to the end of 2019.

Table 5.2 shows the results of these reduced form regressions. The dependent variable in all regressions is core inflation. Columns 1 to 3 are OLS regressions pooling all available data from the Euro Area founding members excluding Ireland and including Greece without controlling for slope heterogeneity. As others have shown before, the inclusion of period and region fixed effects steepens the estimated slope substantially McLeay & Tenreyro (2019). In column 3 we estimate a union-wide slope of -0.016 that is significantly negative.

Next, we move on to columns 4 to 6 in which we apply a FGLS estimator with panel group specific AR(1) autocorrelation structure that is suitable for

[6] Of course, the methodology applied in the previous section is also a reduced form analysis. The key difference is that it allows identification of the multiplier at the country-level and therefore a measure of the heterogeneity.

[7] The business and consumer survey reports country specific balance statistics on expected inflation constructed from the household survey answers on price expectations. It subtracts the share of respondents that expect falling prices from those expecting rising prices. It can therefore not be interpreted as a point estimate of the future inflation rate. Due to incomplete expectations data for Ireland, we exclude the country from the panel analysis.

panels in which $N < T$ (Parks 1967). This estimator eventually allows the estimation of the model controlling for slope heterogeneity proposed by Breitung & Salish (2020) and was introduced in section 4.3.3 of the previous chapter. The evolution of the slope estimate before and after the inclusion of country fixed effects (columns 4 and 5) shows a very similar pattern as in columns 2 and 3. Column 6 additionally controls for slope heterogeneity by including the augmenting regressor for unemployment defined in equation 4.16.[8] The augmenting regressor scales each country's unemployment rate by the difference between the second moments of that country's unemployment rate and the aggregate unemployment rate over the whole sample period. It is a direct measure of the mechanism by which countries with flat Phillips Curves will experience larger volatility of unemployment than those with steep slopes. We described this mechanism and the augmenting regressor in detail in sections 4.3.3 and 4.3.3 of the previous chapter.

While the coefficient of the augmenting regressor in column 6 is not significantly different from zero, the point estimate of the unemployment coefficient steepens from -0.015 (without slope heterogeneity control) to -0.019 (with slope heterogeneity control) giving the steepest slope estimate of all regressions presented in the table. Both point estimates are significantly different from zero. The insignificant coefficient on the augmenting regressor can be a consequence of the relatively homogeneous country-specific multipliers found in section 5.2.

Of course, any heterogeneity in size or elasticity of intertemporal substitution can trigger the same mechanism as heterogeneity in the Phillips Curve slope. Indeed, in the model presented in section 4.3.1 of the previous chapter, the total implicit weight that a country has in the central bank's objective function additionally depends on all other structural parameters of that coun-

8 When constructing the augmenting regressor, we replace the term $\frac{1}{NT}\Sigma_{i=1}^{N}\Sigma_{t=1}^{T}x_{it}^{2}$ with the average squared Euro Area unemployment rate instead of the average of all country specific rates to take into account the different weights with which countries enter the Euro Area's aggregate. However, the results remain qualitatively unchanged when exchanging the two variables. The variable is also scaled by a factor of 0.001 for better readability of the results. Additionally, we do not present regression results when allowing for heterogeneity in the coefficient of expected inflation across regions as we do not detect any heterogeneity when allowing for it.

Table 5.2: Reduced form Phillips Curve estimations with core inflation as dependent variable

	OLS			FGLS		
	(1)	(2)	(3)	(4)	(5)	(6)
Unemployment	-0.004*	-0.003	-0.016***	-0.004**	-0.015***	-0.019***
	(0.002)	(0.002)	(0.004)	(0.002)	(0.003)	(0.004)
Expected Inflation	0.213***	0.145**	0.218***	0.112**	0.209***	0.228***
	(0.039)	(0.048)	(0.030)	(0.047)	(0.058)	(0.060)
Augmenting Regressor Unemployment						0.003
						(0.002)
Constant	0.016	0.122	0.250**	0.132	0.195**	0.209**
	(0.032)	(0.096)	(0.097)	(0.082)	(0.085)	(0.085)
# Lags Inflation	6	6	6	6	6	6
Period Fixed Effects	No	Yes	Yes	Yes	Yes	Yes
Country Fixed Effects	No	No	Yes	No	Yes	Yes
Observations	2436	2436	2436	2436	2436	2436
R^2	0.895	0.913	0.910			
RMSE	0.317	0.304	0.302	0.292	0.290	0.290

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Note: (1) is a pooled OLS regression of country-level core inflation on unemployment and expected inflation. (2) additionally controls for period fixed effects and (3) for country fixed effects. (4) is a FGLS regression with panel group specific AR1 autocorrelation structure and period-fixed effects. (5) additionally controls for country fixed effects. (6) augments (5) with the augmenting regressor defined in equation 4.16 following Breitung & Salish (2020). All regressions include seasonal dummies.

try, such as its size or its elasticity of intertemporal substitution. To clearly illustrate the mechanism of slope heterogeneity, these other structural parameters were set to equal values for all regions. Using real world data, these other structural parameters may differ. However, a further look at Figure 5.2 shows that it is not necessarily large member states that have lower unemployment volatility.[9] Regarding elasticity of intertemporal substitution the determination is harder to make, mainly due to a lack of comparable cross-country estimates.[10]

9 Controlling for the relative size of member states also does not change any of the results presented in this section.

10 In a meta study of estimates on elasticity of intertemporal substitution across countries Havranek et al. (2015) report very heterogeneous results for Euro Area countries, but point out that for a number of countries very few estimates are available making the average for those countries highly dependent on individual modeling choices.

5.4 Slope Dynamics and Heterogeneity

The simulation results presented in the previous chapter clearly suggest that when slope heterogeneity is not controlled for, a change in the Phillips Curve slope in a monetary union observed in reduced form regressions could be entirely due to an increase in heterogeneity among members. This would mean that instead of a union-wide flattening, the curve has steepened for some members and flattened for others. The results from section 5.3 suggest that in the reduced form estimation on Euro Area data, unemployment coefficients steepen mildly when controlling for heterogeneity. However, the augmenting regressor meant to control for slope heterogeneity is insignificant, pointing to relatively homogeneous slopes across member economies. The aim of this section is to analyze to what degree slope heterogeneity has contributed to a potentially changing slope over time.

To detect potential changes in slope heterogeneity, unemployment and the augmenting regressor are both interacted with a dummy variable that takes the value 1 for the sub-period from 2009-2019.[11] The interaction terms indicate whether the coefficients on unemployment and the augmenting regressor differ between the two sub-periods. Table 5.3 presents the results. They suggest that the slope of the Euro Area Phillips Curve has steepened since 2009 - and that slope heterogeneity has masked the steepening to some degree if uncontrolled for.

Before 2009, we observe a coefficient on unemployment of about -0.01 in both specifications. Therefore, during the period 1999 to 2008, controlling for slope heterogeneity does not change the coefficient on unemployment. The augmenting regressor is insignificant. The results are different in the period between 2009 and 2019. Column 1 shows that without control for slope heterogeneity, the slope estimate steepens by a factor of roughly 1.6 compared to the period between 1999 and 2008.[12] In this sub-period however, introducing the augmenting regressor to control for slope heterogeneity leads to significant changes in the Phillips Curve slope estimates. First, the augmenting regressor

[11] As in the previous section, we do not allow for heterogeneity in the coefficient of expected inflation as all results are invariant to the inclusion of interaction terms.

[12] A t-test of equality of coefficients on unemployment in column 1 can only be rejected with a p-value of 0.2.

is significant between 2009 and 2019. Secondly and consequently, the coefficient on unemployment changes when controlling for slope heterogeneity. It steepens to -0.022 compared to -0.016 without control. As the coefficient on unemployment during the first period is essentially unchanged between columns 1 and 2, we can conclude that slope heterogeneity was more pronounced in the second sub-period and a failure to control for it leads to an underestimation of the observed but insignificant steepening in column 1. When controlling for slope heterogeneity, we estimate a steepening by a factor of roughly 2 between the two sub-periods, compared to 1.6 in column 1.[13]

Table 5.3: Reduced form Phillips Curve estimates by sub-period

	(1) FGLS	(2) Augmented FGLS
D1999-2008 × Unemployment	-0.010*	-0.011*
	(0.006)	(0.006)
D2009-2019 × Unemployment	-0.016***	-0.022***
	(0.003)	(0.005)
D1999-2008 × Augmenting Regressor Unemployment		0.006
		(0.004)
D2009-2019 × Augmenting Regressor Unemployment		0.005**
		(0.003)
Expected Inflation	0.194***	0.216***
	(0.060)	(0.060)
Constant	0.217**	0.255***
	(0.086)	(0.089)
# Lags Inflation	6	6
Period Fixed Effects	Yes	Yes
Country Fixed Effects	Yes	Yes
Observations	2436	2436
RMSE	0.290	0.290

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The results suggest that if slope heterogeneity is unaccounted for, the observed steepening of the slope is underestimated due to increasing slope heterogeneity within the Euro Area. The coefficient estimated with slope het-

[13] A t-test of equality of coefficients on unemployment in column 2 between the two sub-periods is rejected with a p-value of 0.1.

erogeneity control in column 2 is more than 30% steeper than the one in column 1. To further illustrate this result, Figure 5.3 shows yearly estimates of the unemployment coefficient without (top row) and with (middle row) control for slope heterogeneity. The bottom row shows the coefficient of the augmenting control variable by year. When controls for heterogeneity are omitted, the estimated coefficient on unemployment is at zero for most of the early 2000s before falling slightly below around 2008. The middle row shows that once heterogeneity is controlled for, the fall in the slope after 2008 is somewhat more pronounced. However, it should also be noted that confidence intervals around the point estimate under heterogeneity control widen.

Figure 5.3: Phillips Curve slope with and without heterogeneity control

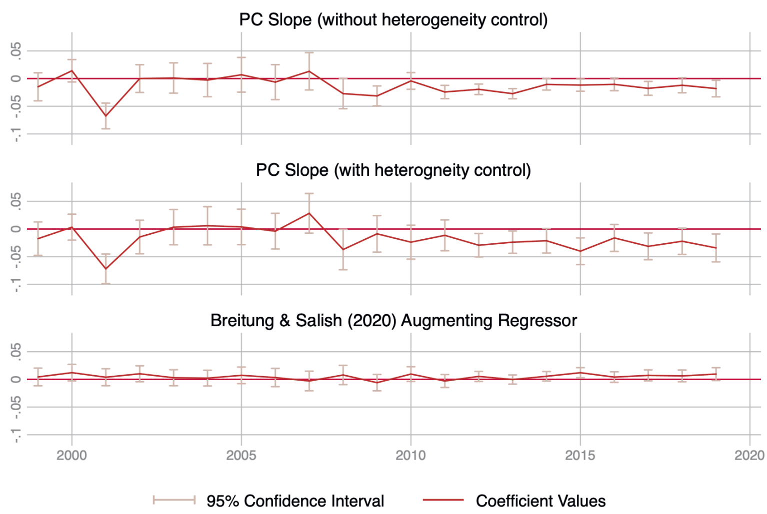

Note: The top panel shows marginal effect of unemployment on the core inflation rate by year in a FGLS regression with panel group specific AR1 autocorrelation structure without controlling for slope heterogeneity. The middle panel shows the marginal effect of unemployment when including the slope heterogeneity regressor proposed by Breitung & Salish (2020). The bottom row shows the marginal effect of the augmenting regressor by year. All regressions control for expected inflation and include seasonal dummies and period fixed effects.

5.5 Conclusion

In this chapter we test whether the patterns in simulated data hold in Euro Area country-level data. First, we provide evidence on causally identified estimates of the inflation-unemployment trade-off for a group of 10 Euro Area member states. We go on to show that applying controls for slope heterogeneity proposed by Breitung & Salish (2020) steepen the estimated Phillips Curve slopes to some degree when analyzing the Euro period as a whole. A sub-period analysis reveals that the slope of the Euro Area Phillips Curve may have steepened by more than what reduced form estimates omitting heterogeneity controls suggest in the period since 2009.

The aim of this chapter and the preceding chapter is to draw attention to the bias that slope heterogeneity among monetary union members can have when estimating the union-wide Phillips Curve relationship in reduced form regressions. The results are meant to encourage researchers presenting evidence on Phillips Curve slopes coming from reduced form panel analyses to control for slope heterogeneity in order to avoid conflating trends of increasing heterogeneity with a flattening of the union-wide slope.

The chapter also raises some questions for further research. Most importantly, it would be relevant for monetary policy to understand why the augmenting regressor turns significant in the period after 2009. The significance implies that the country-level Phillips Curve slope is correlated with the member countries' individual deviations from aggregate unemployment. Using the theoretical model used in chapter 4 or more complex versions of it (see, e.g. Brissimis & Skotida (2008) and Lee (2009)) we can clearly pinpoint the structural parameters that drive this phenomenon: the slope of the Phillips Curve, the size of the economy or elasticity of intertemporal substitution. All of these country-specific parameters determine to what degree the central bank allows a country's output gap to deviate from target. Further research on these heterogeneity patterns would therefore also be useful to better understand the policy reaction function of the central bank.

Appendix A Phillips Multipliers

Figure 5.4: Phillips Multiplier for the Euro Area

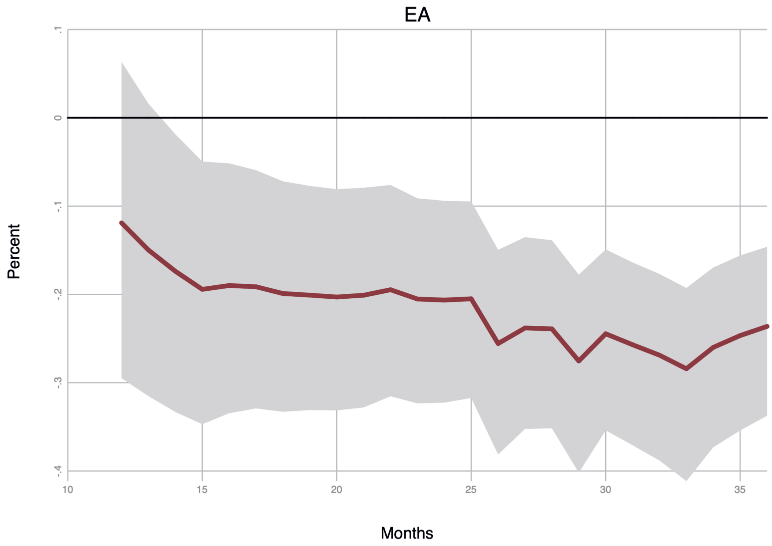

Note: The graph shows the Euro Area Phillips Multiplier over horizons between 12 and 36 months after the monetary policy shock for the Euro Area following Barnichon & Mesters (2021). Shaded area indicates 68% confidence intervals. Regressions of cumulative core inflation on cumulative unemployment include 36 lags of core inflation and unemployment. As the multiplier is indeterminate at short horizons due to the transmission lag of monetary policy, we only report horizons between 12 and 36 months.

Figure 5.5: Phillips Multipliers by Country

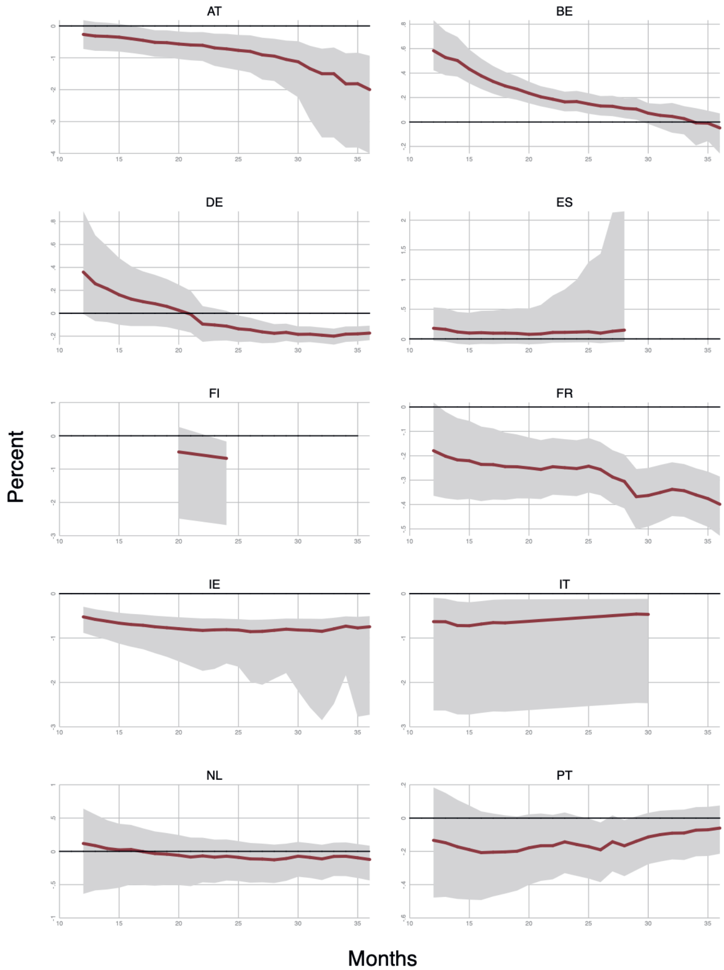

Note: The graphs show Phillips Multipliers over horizons between 12 and 36 months after the monetary policy shock by country following Barnichon & Mesters (2021). Shaded areas indicate 68% Anderson-Rubin confidence intervals. Regressions of cumulative core inflation on cumulative unemployment include 36 lags of core inflation and unemployment. As the multiplier is indeterminate at short horizons due to the transmission lag of monetary policy, we only report horizons between 12 and 36 months.

6
Conclusion

The research contained in this dissertation is concerned with several questions that emerged in the context of the macroeconomic situation of Western economies in the last decade, in particular the apparent disconnect between monetary policy and unemployment on the one hand and inflation on the other. As traditional tools of monetary policy proved to be ineffective to stabilize the economies of the United States, the Euro Area and other Western countries, the management of economic expectations of the general public, in particular its inflation expectations, became a centerpiece of monetary policy making. Chapters 2 and 3 contribute to the literature on the potential effects of such efforts on consumption behavior of households as well as the capability of central banks to influence economic expectations in the first place. Chapters 4 and 5 discuss econometric issues that could lead to spuriously flattening Phillips Curve slopes when estimating reduced form regressions of the relationship.

Research interest in the reaction of consumption to expected inflation has increased in recent years due to efforts by central banks to kick-start demand by steering inflation expectations. Chapter 2 contributes to this literature by analyzing whether various components of households' balance sheets determine how consumption reacts to expected inflation. Two channels in particular are conceivable: an increase in inflation expectations can raise consumption through direct increases in expected real wealth, e.g. for households with nominal financial liabilities. By affecting the real interest rate, expected inflation can interact with wealth if only those households can adapt their consumption to current real interest rates that are not budget constrained or sufficiently liquid to shift funds between consumption and savings. The chapter investigates these channels empirically using household-level information on balance sheets, durable consumption, and inflation expectations from the Dutch Central Bank's Household Survey. The analysis shows that investments in risky assets as well as a household's net worth moderates the relation between expected inflation and durable spending decisions. The net worth effect is most pronounced for households with fixed interest rate mortgages. Clearly, the chapter does not present causal evidence of the effect of inflation expectations on consumption choices. While there are studies documenting such causal effects, e.g. D'Acunto et al. (2016) and Coibion et al. (2019a), the balance sheet channel does not receive detailed attention. When trying to affect consumption, monetary policy making could benefit from a better understanding of heterogeneity in the reaction of households.

Chapter 3 studies the impact of the Fed's monetary policy announcements on households' expectations by comparing responses to the Survey of Consumer Expectations before and after Federal Open Market Committee (FOMC) meetings, over the period 2013-2019. The results show that Fed decisions affect expectations of interest rates on savings accounts, particularly for respondents with high financial and numerical literacy. The impact of monetary policy announcements on inflation expectations is muted, even in response to some of the most relevant meetings of the FOMC that took place during that period. Expectations of personal financial conditions are barely affected. The results stand in contrast to experimental studies that find strong effects of monetary policy and other macroeconomic news on expectations of households receiving a specific treatment, suggesting that the news naturally reaching the general population may provide weaker signals. Understanding the differences between results obtained in experimental settings and using observational data in this literature would greatly benefit the policy debate on this issue. If central banks are able to better emulate experimental settings, the impact of their communication efforts could be substantially increased. Additionally, the literature would benefit from evaluations of policy communication outside central banking. This would allow comparisons of a broader variety of communication strategies and policy environments.

Chapter 4 illustrates the impact of heterogeneous slopes on the estimation of the Phillips Curve slope in a monetary union. If regional slopes differ, any attempt to estimate a union wide slope either on aggregate or pooled regional data will suffer from substantial omitted variable bias. The reason is that the slope of the Phillips Curve changes the implicit weight a union member gets in the central bank's monetary policy rule. In regions with steep Phillips Curves, demand shocks will be stabilized to a larger degree than the mere size of the economy would justify as these demand shocks would impact aggregate union inflation to a larger degree than demand shocks in regions with flat Phillips Curves. The more pronounced the heterogeneity, the larger the bias to coefficients from aggregated or pooled estimations.

Chapter 5 tests whether the patterns in simulated data hold in Euro Area country-level data. First, it provides evidence on causally identified estimates of the inflation-unemployment trade-off for a group of 10 Euro Area member states. It goes on to show that applying controls for slope heterogeneity proposed by Breitung & Salish (2020) lead to an insignificant steepening of the estimated Phillips Curve slopes when analyzing the Euro period as a whole. A sub-period analysis reveals that the slope of the Euro Area Phillips Curve has steepened by more than what reduced form estimates omitting heterogeneity controls suggest in the period since 2009 and the augmenting regressor controlling for slope heterogeneity turns significant.

In the context of the last two chapters, further research should focus on more general cases than the simplified monetary union model used to obtain predictions in chapter 4, in particular regarding the distribution of region-specific deviations from the union-wide slope parameter. Chapter 4 assumes that the two regions' deviations are fixed. A natural extension could be to treat the slope coefficient itself as a random variable. From a monetary policy perspective, the chapters raise further questions, too. Most importantly, it is relevant for monetary policy to understand why the augmenting regressor turns significant in the period after 2009. The significance implies that the country-level Phillips Curve slope is correlated with the member countries' individual deviations from aggregate unemployment. Using the theoretical model used in chapter 4 or more complex versions of it (see, e.g. Brissimis & Skotida (2008) and Lee (2009)) we can clearly pinpoint the structural parameters that drive this phenomenon: slope of the Phillips Curve, the size of the economy or elasticity of intertemporal substitution. All of these country-specific parameters determine to what degree the central bank allows a country's output gap to deviate from target. Further research on these heterogeneity patterns would therefore also be useful to better understand the policy reaction function of the central bank.

Bibliography

Amberger, Johanna, and Ralf Fendel. 2017. "The Slope of the Euro Area Phillips Curve: Always and Everywhere the Same?" *Applied Economics and Finance* 4 (3): 77–88.

Anari, Ali, and James Kolari. 2002. "House Prices and Inflation." *Real Estate Economics* 30 (1): 67–84. ISSN: 1540-6229, accessed January 28, 2019.

Anderson, Theodore W., and Herman Rubin. 1949. "Estimation of the parameters of a single equation in a complete system of stochastic equations." Publisher: Institute of Mathematical Statistics, *Annals of Mathematical statistics* 20 (1): 46–63.

Andrade, Philippe, and Filippo Ferroni. 2021. "Delphic and Odyssean monetary policy shocks: Evidence from the euro area." *Journal of Monetary Economics* 117:816–832.

Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart. 2019. "Subjective Models of the Macroeconomy: Evidence from Experts and a Representative Sample." *Available at SSRN 3355356.*

Armantier, Olivier, Giorgio Topa, Wilbert Van der Klaauw, and Basit Zafar. 2017. "An overview of the survey of consumer expectations." *Economic Policy Review,* no. 23: 51–72.

Auclert, Adrien. 2019. "Monetary policy and the redistribution channel." *American Economic Review* 109 (6): 2333–67.

Bachmann, Rüdiger, Tim O. Berg, and Eric R. Sims. 2015. "Inflation expectations and readiness to spend: Cross-sectional evidence." *American Economic Journal: Economic Policy* 7 (1): 1–35.

Ball, Laurence M., and Sandeep Mazumder. 2019. "A Phillips curve for the euro area." *NBER Working Paper* (Nr. 26450).

Bank, European Central. 2021. "Monetary policy communication and the strategy review." European Central Bank. July 5. Accessed September 18, 2021. `https://www.ecb.europa.eu/home/search/review/html/monpol-communication.en.html`.

Barnichon, Regis, and Geert Mesters. 2021. "The Phillips Multiplier." Publisher: Elsevier, *Journal of Monetary Economics* 117:689–705.

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen. 2014a. "High-dimensional methods and inference on structural and treatment effects." *Journal of Economic Perspectives* 28 (2): 29–50.

———. 2014b. "Inference on treatment effects after selection among high-dimensional controls." *The Review of Economic Studies* 81 (2): 608–650.

Bernanke, Ben S. 1993. "Credit in the Macroeconomy." *Quarterly Review-Federal Reserve Bank of New York* 18:50–50.

———. 2013. "Transcript of Chairman Bernanke's Press Conference, June 19, 2013."

Bernanke, Ben S, and Kenneth N Kuttner. 2005. "What explains the stock market's reaction to Federal Reserve policy?" *The Journal of finance* 60 (3): 1221–1257.

Bharadwaj, Asha, and Maximiliano Dvorkin. 2020. "The Case of the Reappearing Phillips Curve: A Discussion of Recent Findings." *Federal Reserve Bank of St. Louis Review* 102 (3): 313–337.

Binder, Carola. 2017. "Fed speak on main street: Central bank communication and household expectations." Publisher: Elsevier, *Journal of Macroeconomics* 52:238–251.

———. 2019. "Panel Conditioning in the Survey of Consumer Expectations." *Available at SSRN 3330376.*

Boivin, Jean, and Marc P. Giannoni. 2006. "Has monetary policy become more effective?" Publisher: MIT Press, *The Review of Economics and Statistics* 88 (3): 445–462.

Bottone, Marco, and Alfonso Rosolia. 2019. "Monetary policy, firms' inflation expectations and prices: causal evidence from firm-level data." *Bank of Italy Temi di Discussione (Working Paper) No* 1218.

Breitung, Jörg, and Nazarii Salish. 2020. "Estimation of heterogeneous panels with systematic slope variations." Publisher: Elsevier, *Journal of Econometrics.*

Brissimis, Sophocles N., and Ifigeneia Skotida. 2008. "Optimal monetary policy in the euro area in the presence of heterogeneity." *Journal of International Money and Finance* 27 (2): 209–226.

Bruine de Bruin, Wändi, Wilbert Van der Klaauw, Julie S. Downs, Baruch Fischhoff, Giorgio Topa, and Olivier Armantier. 2010. "Expectations of inflation: The role of demographic variables, expectation formation, and financial literacy." *Journal of Consumer Affairs* 44 (2): 381–402.

Bullard, James. 2018. "The Case of the Disappearing Phillips Curve." 2018 ECB Forum on Central Banking.

Burbidge, John B., Lonnie Magee, and A. Leslie Robb. 1988. "Alternative transformations to handle extreme values of the dependent variable." *Journal of the American Statistical Association* 83 (401): 123–127.

Burke, Mary A., and Ali Ozdagli. 2013. "Household inflation expectations and consumer spending: evidence from panel data." *Working papers, Federal Reserve Bank of Boston.*

Case, Karl E., John M. Quigley, and Robert J. Shiller. 2005. "Comparing wealth effects: the stock market versus the housing market." *Advances in macroeconomics* 5 (1).

Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia. 2017. "Inflation expectations, learning, and supermarket prices: Evidence from survey experiments." *American Economic Journal: Macroeconomics* 9 (3): 1–35.

Claus, Edda, and Viet Hoang Nguyen. 2020. "Monetary policy shocks from the consumer perspective." Publisher: Elsevier, *Journal of Monetary Economics* 114:159–173.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico. 2020. "Monetary policy when households have debt: new evidence on the transmission mechanism." *The Review of Economic Studies* 87 (1): 102–129.

Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, and Maarten van Rooij. 2019a. *How does consumption respond to news about inflation? Field evidence from a randomized control trial.* National Bureau of Economic Research.

Coibion, Olivier, and Yuriy Gorodnichenko. 2012. "What can survey forecasts tell us about information rigidities?" *Journal of Political Economy* 120 (1): 116–159.

———. 2015a. "Information rigidity and the expectations formation process: A simple framework and new facts." *American Economic Review* 105 (8): 2644–78.

———. 2015b. "Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation." *American Economic Journal: Macroeconomics* 7 (1): 197–232.

Coibion, Olivier, Yuriy Gorodnichenko, and Rupal Kamdar. 2017. "The Formation of Expectations, Inflation and the Phillips Curve." *NBER Working Paper* (No. 23304).

Coibion, Olivier, Yuriy Gorodnichenko, Saten Kumar, and Mathieu Pedemonte. 2020. "Inflation expectations as a policy tool?" Publisher: Elsevier, *Journal of International Economics:* 103297.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber. 2019b. "Monetary policy communications and their effects on household inflation expectations." *NBER Working Paper* (No. 25482).

Crump, Richard K., Stefano Eusepi, Andrea Tambalotti, and Giorgio Topa. 2015. "Subjective Intertemporal Substitution." *Federal Reserve Bank of New York Staff Reports* Number 734.

Cumming, Fergus, and Paul Hubert. 2019. "The role of households' borrowing constraints in the transmission of monetary policy."

D'Acunto, Francesco, Daniel Hoang, and Michael Weber. 2016. *Unconventional Fiscal Policy, Inflation Expectations, and Consumption Expenditure.* SSRN Scholarly Paper ID 2760238. Rochester, NY: Social Science Research Network, March 3.

D'Amico, Stefania, and Mira Farka. 2011. "The Fed and the stock market: An identification based on intraday futures data." *Journal of Business & Economic Statistics* 29 (1): 126–137.

De Fiore, Fiorella, Marco Jacopo Lombardi, and Johannes Schuffels. 2021. "Are households indifferent to monetary policy announcements?" *BIS Working Papers 956.*

Di Maggio, Marco, Amir Kermani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao. 2017. "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging." *American Economic Review* 107, no. 11 (November): 3550–3588. ISSN: 0002-8282.

Dincer, N. Nergiz, and Barry Eichengreen. 2013. "Central bank transparency and independence: updates and new measures." Publisher: WP.

DNB. 2020. "Residential Mortgages - Statistics." Accessed February 8, 2020. `https://statistiek.dnb.nl/en/dashboards/residential-mortgages/index.aspx`.

Dräger, Lena, Michael Lamla, and Damjan Pfajfar. 2016. "Are survey expectations theory-consistent? The role of central bank communication and news." *European Economic Review* 85 (C): 84–111. ISSN: 0014-2921.

Duca, Ioana A., Geoff Kenny, and Andreas Reuter. 2018. "Inflation expectations, consumption and the lower bound: micro evidence from a large euro area survey." *ECB Working Paper Series* No 2196.

Eser, Fabian, Peter Karadi, Philip R. Lane, Laura Moretti, and Chiara Osbat. 2020. "The Phillips Curve at the ECB." Publisher: Wiley Online Library, *The Manchester School* 88:50–85.

Fama, Eugene F., and G. William Schwert. 1977. "Asset returns and inflation." *Journal of financial economics* 5 (2): 115–146.

Federal Reserve. 2014a. "FOMC Statement December 2014." Board of Governors of the Federal Reserve System. December 17. Accessed April 3, 2021. https://www.federalreserve.gov/newsevents/pressreleases/monetary20141029a.htm.

———. 2014b. "FOMC statement October 2014." October 29. Accessed April 3, 2021. https://www.federalreserve.gov/newsevents/pressreleases/monetary20141029a.htm.

———. 2020. "Fed Listens - Perspectives from the public." https://www.federalreserve.gov/publications/files/fedlistens-report-20200612.pdf.

Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack. 2011. "The financial market effects of the Federal Reserve's large-scale asset purchases." *International Journal of Central Banking* 7 (1): 45–52.

Gürkaynak, Refet S., Brian P. Sack, and Eric T. Swanson. 2004. "Do actions speak louder than words? The response of asset prices to monetary policy actions and statements." Publisher: FEDS Working Paper.

Havranek, Tomas, Roman Horvath, Zuzana Irsova, and Marek Rusnak. 2015. "Cross-country heterogeneity in intertemporal substitution." Publisher: Elsevier, *Journal of International Economics* 96 (1): 100–118.

Hazell, Jonathon, Juan Herreño, Emi Nakamura, and Jón Steinsson. 2020. "The Slope of the Phillips Curve: Evidence from US States." *NBER Working Paper* (Nr. 28005).

Hoesli, Martin, Colin Lizieri, and Bryan MacGregor. 2008. "The Inflation Hedging Characteristics of US and UK Investments: A Multi-Factor Error Correction Approach." *The Journal of Real Estate Finance and Economics* 36, no. 2 (February 1): 183–206. ISSN: 1573-045X, accessed January 28, 2019.

Ichiue, Hibiki, and Shusaku Nishiguchi. 2015. "Inflation expectations and consumer spending at the zero bound: Micro evidence." *Economic Inquiry* 53 (2): 1086–1107.

Imbs, Jean, Eric Jondeau, and Florian Pelgrin. 2011. "Sectoral Phillips curves and the aggregate Phillips curve." Publisher: Elsevier, *Journal of Monetary Economics* 58 (4): 328–344.

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

Kapetanios, George, Simon Price, Menelaos Tasiou, and Alexia Ventouri. 2020. "State-level wage Phillips curves." Publisher: Elsevier, *Econometrics and Statistics.*

Kim, Sangbae, and Francis In. 2005. "The relationship between stock returns and inflation: new evidence from wavelet analysis." *Journal of Empirical Finance* 12 (3): 435–444.

Lamla, Michael J., and Dmitri V. Vinogradov. 2019. "Central bank announcements: Big news for little people?" Publisher: Elsevier, *Journal of Monetary Economics* 108:21–38.

Lee, Jim. 2009. "Evaluating monetary policy of the euro area with cross-country heterogeneity: Evidence from a New Keynesian model." Publisher: Elsevier, *Economic Systems* 33 (4): 325–343.

Lewis, Daniel J., Christos Makridis, and Karel Mertens. 2019. "Do Monetary Policy Announcements Shift Household Expectations?" *FRB of New York Staff Report,* no. 897.

Lieb, Lenard, and Johannes Schuffels. 2022. "Inflation Expectations and Consumer Spending: The Role of Household Balance Sheets." *Empirical Economics.*

Lindé, Jesper, and Mathias Trabandt. 2019. "Resolving the missing deflation puzzle." Publisher: CEPR Discussion Paper No. DP13690.

Liu, Lucy Qian, Liang Wang, and Randall Wright. 2011. "On the "hot potato" effect of inflation: Intensive versus extensive margins." *Macroeconomic Dynamics* 15 (S2): 191–216.

Lombardi, Marco J., Marianna Riggi, and Eliana Viviano. 2020. "Bargaining power and the Phillips curve: a micro-macro analysis." *Bank of Italy Temi di Discussione (Working Paper) No* 1302.

Lombardi, Marco J., and Feng Zhu. 2018. "A Shadow Policy Rate to Calibrate US Monetary Policy at the Zero Lower Bound." Publisher: International Journal of Central Banking, *International Journal of Central Banking* 14 (5): 305–346.

Lusardi, Annamaria, and Olivia S. Mitchell. 2014. "The economic importance of financial literacy: Theory and evidence." *Journal of economic literature* 52 (1): 5–44.

McLeay, Michael, and Silvana Tenreyro. 2019. "Optimal inflation and the identification of the Phillips Curve." *NBER Working Paper* No. 25892.

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household balance sheets, consumption, and the economic slump." *The Quarterly Journal of Economics* 128 (4): 1687–1726.

Okun, Arthur M. 1962. "Potential GNP: Its Measurement and Significance." *Proceedings of the Business and Economic Statistics Section of the American Statistical Association.*

Park, JeeYeon. 2013a. "Dow, S&P 500 post all-time highs as Fed maintains stimulus." CNBC. September 18. Accessed February 27, 2021. `https://www.cnbc.com/2013/09/18/us-marketwrap.html`.

———. 2013b. "Fed to taper bond buying by $10 billion a month." CNBC. December 18. Accessed February 27, 2021. `https://www.cnbc.com/2013/12/18/fed-begins-taper-program.html`.

Parks, Richard W. 1967. "Efficient estimation of a system of regression equations when disturbances are both serially and contemporaneously correlated." Publisher: Taylor & Francis Group, *Journal of the American Statistical Association* 62 (318): 500–509.

Pence, Karen M. 2006. "The role of wealth transformations: An application to estimating the effect of tax incentives on saving." *The BE Journal of Economic Analysis & Policy* 5 (1).

Pesaran, M. Hashem, Yongcheol Shin, and Ron P. Smith. 1999. "Pooled mean group estimation of dynamic heterogeneous panels." Publisher: Taylor & Francis, *Journal of the American Statistical Association* 94 (446): 621–634.

Phillips, Alban W. 1958. "The relation between unemployment and the rate of change of money wage rates in the United Kingdom, 1861-1957." Publisher: JSTOR, *economica* 25 (100): 283–299.

Prial, Dunstan. 2013. "Bernanke Offers Possible Timetable for Tapering." Fox Business. June 19. Accessed February 27, 2021. `https://web.archive.org/web/20130622064449/http://www.foxbusiness.com/economy/2013/06/19/fed-decision-on-tap/`.

Roth, Christopher, and Johannes Wohlfart. 2020. "How do expectations about the macroeconomy affect personal expectations and behavior?" Publisher: MIT Press, *Review of Economics and Statistics* 102 (4): 731–748.

Schotman, Peter C., and Mark Schweitzer. 2000. "Horizon sensitivity of the inflation hedge of stocks." *Journal of empirical Finance* 7 (3): 301–315.

Schuffels, Johannes, Clemens Kool, Lenard Lieb, and Tom van Veen. 2022. "Is the slope of the Euro Area Phillips Curve steeper than it seems? Heterogeneity and identification."

Swanson, Eric T. 2017. "Measuring the effects of Federal Reserve forward guidance and asset purchases on financial markets." *NBER Working Paper.*

Vellekoop, Nathanael, and Mirko Wiederholt. 2017. "Inflation Expectations and Choices of Households." *Goethe University Frankfurt mimeo.*

Wiederholt, Mirko. 2014. "Empirical properties of inflation expectations and the zero lower bound." *Goethe University Frankfurt.*

Wu, Jing Cynthia, and Fan Dora Xia. 2016. "Measuring the macroeconomic impact of monetary policy at the zero lower bound." Publisher: Wiley Online Library, *Journal of Money, Credit and Banking* 48 (2): 253–291.

Yellen, Janet. 2013. "Communication in monetary policy." In *Society of American Business Editors and Writers 50th Anniversary Conference, Washington DC.*

Impact Paragraph

In the last decade, several questions concerning the effectiveness of monetary policy emerged in the context of the macroeconomic situation of Western economies. In particular, a disconnect between monetary policy and unemployment on the one hand and inflation on the other became apparent. This dissertation contributes to two separate strands of the broader macroeconomic literature that both have policy relevance because of this disconnect.

As traditional tools of monetary policy proved to be ineffective to stabilize the economies of the United States, the Euro Area and other Western countries, the management of economic expectations of the general public, in particular its inflation expectations, became an important issue for monetary policy makers. Additionally, Chapters 2 and 3 contribute to the literature on the potential effects of such efforts on consumption behavior of households as well as the capability of central banks to influence economic expectations in the first place. Chapters 4 and 5 discuss econometric issues that could lead to spuriously flattening Phillips Curve slopes when estimating reduced form regressions of the relationship.

Chapter 2 analyzes whether various components of households' balance sheets determine how consumption reacts to expected inflation. Two channels in particular are conceivable: an increase in inflation expectations can raise consumption through direct increases in expected real wealth, e.g. for households with nominal financial liabilities. By affecting the real interest rate, expected inflation can interact with wealth if only those households can adapt their consumption to current real interest rates that are not budget constrained or sufficiently liquid to shift funds between consumption and savings. The chapter investigates these channels empirically using household-level information on balance sheets, durable consumption, and inflation expectations from the Dutch Central Bank's Household Survey. The analysis shows that investments in risky assets as well as a household's net worth moderates the relation between expected inflation and durable spending decisions. The net worth effect is most pronounced for households with fixed interest rate mortgages. This chapter draws its policy relevance from repeated calls for central banks to use the steering of household inflation expectations to increase demand in a period during which short-term interest rates cannot be further lowered. The research shows that there seem to be households among which this effort could lead to increased spending on certain types of durable goods. However, the question whether this would be desirable given the dire financial situation of many of these households remains. In the contrary, the chapter raises the point that those households that seem most responsive to their inflation expectations when making consumption decisions, would rather benefit from more realistic, i.e. lower, inflation expectations. As such, the chapter contains lessons for policy makers, particularly regarding the importance of clear and direct communication about goals and strategy of monetary policy.

Chapter 3 studies the impact of the Fed's monetary policy announcements on households' expectations by comparing responses to the Survey of Consumer Expectations before and after Federal Open Market Committee (FOMC) meetings, over the period 2013-2019. The results show that Fed decisions affect expectations of interest rates on savings accounts, particularly for respondents with high financial and numerical literacy. The impact of monetary policy announcements on inflation expectations is muted, even in response to some of the most relevant meetings of the FOMC that took place during that period. Expectations of personal financial conditions are barely affected. The results stand in contrast to experimental studies that find strong effects of monetary policy and other macroeconomic news on expectations of households receiving a specific treatment, suggesting that the news naturally reaching the general population may provide weaker signals. This chapter is based on a paper co-authored with two economists from the Bank for International Settlements, an institution that brings together monetary policy makers from around the world. Therefore, the paper has received some attention after its publication as a Working Paper, for example in a speech given by the Head of the institution's Monetary and Economic Department in September 2021.[1] The paper

[1] "Back to the future: intellectual challenges for monetary policy", Speech by Claudio Borio, Head of the BIS Monetary and Economic Department at the David Finch Lecture, University of Melbourne, 2 September 2021, https://www.bis.org/speeches/sp210902a.htm

holds important lessons for monetary policy makers in that it shows that despite respondents' understanding that monetary policy affects interest rate, their personal economic decisions are virtually unaffected. This means that certain transmission mechanisms of monetary policy may work less smoothly than expected.

Chapter 4 illustrates the impact of heterogeneous slopes on the estimation of the Phillips Curve slope in a monetary union. If regional slopes differ, any attempt to estimate a union wide slope either on aggregate or pooled regional data will suffer from substantial omitted variable bias. The reason is that the slope of the Phillips Curve changes the implicit weight a union member gets in the central bank's monetary policy rule. In regions with steep Phillips Curves, demand shocks will be stabilized to a larger degree than the mere size of the economy would justify as these demand shocks would impact aggregate union inflation to a larger degree than demand shocks in regions with flat Phillips Curves. The more pronounced the heterogeneity, the larger the bias to coefficients from aggregated or pooled estimations.

Chapter 5 tests whether the patterns in simulated data hold in Euro Area country-level data. First, it provides evidence on causally identified estimates of the inflation-unemployment trade-off for a group of 10 Euro Area member states. It goes on to show that applying controls for slope heterogeneity proposed by Breitung & Salish (2020) lead to an insignificant steepening of the estimated Phillips Curve slopes when analyzing the Euro period as a whole. A sub-period analysis reveals that the slope of the Euro Area Phillips Curve has steepened by more than what reduced form estimates omitting heterogeneity controls suggest in the period since 2009 and the augmenting regressor controlling for slope heterogeneity turns significant.

Reduced form estimations of the Phillips Curve remain quite prominent in the literature. The last two chapters of this dissertation draw their policy relevance from this prominence. Given the relevance of the Phillips Curve relationship for monetary policy making, sound empirical evidence on its magnitude is crucial. The goal of chapters 4 and 5 is to encourage researchers to address the issue of slope heterogeneity when estimating slopes of aggregate

Phillips Curves. The chapters clearly demonstrate the necessity of controlling for slope heterogeneity to ensure sound reduced form evidence on the topic.

Summary

The research contained in this dissertation is concerned with several questions that emerged in the context of the macroeconomic situation of Western economies in the last decade, in particular the apparent disconnect between monetary policy and unemployment on the one hand and inflation on the other. As traditional tools of monetary policy proved to be ineffective to stabilize the economies of the United States, the Euro Area and other Western countries, the management of economic expectations of the general public, in particular its inflation expectations, became a centerpiece of monetary policy making. Chapters 2 and 3 contribute to the literature on the potential effects of such efforts on consumption behavior of households as well as the capability of central banks to influence economic expectations in the first place. Chapters 4 and 5 discuss econometric issues that could lead to spuriously flattening Phillips Curve slopes when estimating reduced form regressions of the relationship.

Research interest in the reaction of consumption to expected inflation has increased in recent years due to efforts by central banks to kick-start demand by steering inflation expectations. Chapter 2 contributes to this literature by analyzing whether various components of households' balance sheets determine how consumption reacts to expected inflation. Two channels in particular are conceivable: an increase in inflation expectations can raise consumption through direct increases in expected real wealth, e.g. for households with nominal financial liabilities. By affecting the real interest rate, expected inflation can interact with wealth if only those households can adapt their consumption to current real interest rates that are not budget constrained or sufficiently liquid to shift funds between consumption and savings. The chapter investigates these channels empirically using household-level information on balance sheets, durable consumption, and inflation expectations from the Dutch Central Bank's Household Survey. The analysis shows that investments in risky assets as well as a household's net worth moderates the relation between expected inflation and durable spending decisions. The net worth effect is most pronounced for households with fixed interest rate mortgages.

Chapter 3 studies the impact of the Fed's monetary policy announcements on households' expectations by comparing responses to the Survey of Consumer Expectations before and after Federal Open Market Committee (FOMC) meetings, over the period 2013-2019. The results show that Fed decisions affect expectations of interest rates on savings accounts, particularly for respondents with high financial and numerical literacy. The impact of monetary policy announcements on inflation expectations is muted, even in response to some of the most relevant meetings of the FOMC that took place during that period. Expectations of personal financial conditions are barely affected. The results stand in contrast to experimental studies that find strong effects of monetary policy and other macroeconomic news on expectations of households receiving a specific treatment, suggesting that the news naturally reaching the general population may provide weaker signals.

Chapter 4 illustrates the impact of heterogeneous slopes on the estimation of the union-wide Phillips Curve slope in a monetary union. If regional slopes differ, any attempt to estimate a union wide slope either on aggregate or pooled regional data will suffer from substantial omitted variable bias. The reason is that the slope of the Phillips Curve changes the implicit weight a union member gets in the central bank's monetary policy rule. In regions with steep Phillips Curves, demand shocks will be stabilized to a larger degree than the mere size of the economy would justify as these demand shocks would impact aggregate union inflation to a larger degree than demand shocks in regions with flat Phillips Curves. The more pronounced the heterogeneity, the larger the bias to coefficients from aggregated or pooled estimations.

Chapter 5 tests whether the patterns in simulated data hold in Euro Area country-level data. First, it provides evidence on causally identified estimates of the inflation-unemployment trade-off for a group of 10 Euro Area member states. It goes on to show that applying controls for slope heterogeneity proposed by Breitung & Salish (2020) lead to an insignificant steepening of the estimated Phillips Curve slopes when analyzing the Euro period as a whole. A sub-period analysis reveals that the slope of the Euro Area Phillips Curve has steepened by more than what reduced form estimates omitting heterogeneity controls suggest in the period since 2009 and the augmenting regressor controlling for slope heterogeneity turns significant.

Nederlandse Samenvatting

Het in dit proefschrift opgenomen onderzoek is gericht op een aantal vraagstukken dat naar voren komt in de context van de macro-economische situatie van de West-Europese economieën over de laatste tien jaar. Daarbij gaat het met name om de schijnbare loskoppeling van monetair beleid en werkloosheid enerzijds en inflatie anderzijds. De traditionele hulpmiddelen van het monetair beleid bleken er niet langer in te slagen de economieën van de Verenigde Staten, de eurozone en de overige westelijke landen te stabiliseren. Hierdoor kwam binnen het monetair beleid het beheer van de economische verwachtingen van het grote publiek, met name ten aanzien van inflatie, centraal te staan. Hoofdstuk 2 en 3 vormen een bijdrage aan de literatuur over de potentiële effecten van dergelijke inspanningen op het consumptiegedrag van huishoudens alsook over het vermogen van centrale banken om de economische verwachtingen überhaupt te beïnvloeden. Hoofdstuk 4 en 5 gaan in op een aantal econometrische kwesties die zouden kunnen leiden tot een oneigenlijk afvlakkende helling van de Phillips-curve bij het schatten van de 'reduced form'-regressies van de relatie.

Als gevolg van de inspanningen van centrale banken om de vraag weer op gang te helpen door de inflatieverwachtingen te sturen, is de interesse van onderzoekers in de reactie van de consumptie op de verwachte inflatie de laatste jaren toegenomen. Hoofdstuk 2 draagt bij aan dit onderzoek door te analyseren in hoeverre de verschillende componenten van de balansen van huishoudens bepalen hoe de consumptie reageert op de verwachte inflatie. Daarbij kan met name aan twee kanalen worden gedacht: een toename van de verwachte inflatie kan de consumptie doen toenemen door middel van een directe stijging van het verwachte reële vermogen, bijvoorbeeld voor huishoudens met nominale financiële verplichtingen. Doordat de verwachte inflatie de werkelijke rentes beïnvloedt, kan er ook een wisselwerking ontstaan met het vermogen wanneer alleen huishoudens die geen budgetbeperkingen kennen of die voldoende liquide zijn om hun middelen heen en weer te schuiven tussen consumptie en spaargeld, hun consumptie kunnen aanpassen aan de actuele reële rentes. In dit hoofdstuk worden deze kanalen empirisch onderzocht met behulp van gegevens van de balansen op huishoudensniveau en de inflatieverwachting van de enquête onder huishoudens van DNB. Uit dit onderzoek blijkt dat investeringen in risicovol vermogen alsook de nettowaarde van de huishoudens een dempend effect hebben op de relatie tussen de verwachte inflatie en beslissingen over duurzame uitgaven. Het effect van de nettowaarde is het meest uitgesproken bij huishoudens met een hypotheek met vaste rente.

Hoofdstuk 3 gaat in op het effect van de monetaire beleidsaankondigingen van de Fed op de verwachtingen onder huishoudens. Gedurende de periode 2013-2019 worden de reacties op consumentenenquêtes van voor en na bijeenkomsten van de Federal Open Market Committee (FOMC) met elkaar vergeleken. Uit de resultaten blijkt dat de besluiten van de Fed van invloed zijn op de verwachtingen ten aanzien van de rentes op spaarrekeningen, met name onder ondervraagden met grote financiële kennis en rekenkundig inzicht. Het effect van monetaire beleidsaankondigingen op de inflatieverwachtingen is gematigd, zelfs bij een aantal van de belangrijkste bijeenkomsten van de FOMC die tijdens die periode plaatsvonden. De verwachtingen ten aanzien van de persoonlijke financiële omstandigheden worden nauwelijks beïnvloed. Deze resultaten staan in schril contrast met die van experimentele studies die juist wijzen op een sterk effect van het monetair beleid en overig macro-economisch nieuws op de verwachtingen onder huishoudens die een speciale behandeling krijgen, hetgeen wel suggereert dat een zwakker signaal uitgaat van nieuws dat van nature het grote publiek bereikt.

Hoofdstuk 4 illustreert het effect van heterogene hellingen op de schatting van de uniebrede helling van de Phillips-curve binnen een monetaire unie. Als de regionale hellingen verschillen, heeft elke poging om de uniebrede helling te schatten op basis van getotaliseerde of gebundelde regionale gegevens te lijden van een aanzienlijke vertekening van de resultaten vanwege weggelaten variabelen. De reden daarvan is dat door de helling van de Phillips-curve de impliciete weging die een lid van de unie krijgt binnen een monetaire beleidsregel van de centrale bank, wordt gewijzigd. In regio's met een steile Phillips-curve worden schokken aan vraagzijde in hogere mate gestabiliseerd dan de pure omvang van de economie zou rechtvaardigen, aangezien dergelijke schokken aan vraagzijde heviger zouden doorwerken op de totale inflatie van de unie dan schokken aan vraagzijde in regio's met een vlakke Phillips-curve. Hoe nadrukkelijker de heterogeniteit, des te groter wordt de vertekening voor coëfficiënten van getotaliseerde of gebundelde schattingen.

In hoofdstuk 5 wordt getoetst of de patronen in gesimuleerde gegevens bij gegevens op het niveau van de eurozone overeind blijven. Allereerst wordt hierdoor bewijs geboden voor oorzakelijk geïdentificeerde schattingen van de uitruil tussen werkloosheid en inflatie voor een groep van 10 lidstaten van de eurozone. Vervolgens wordt aangetoond dat toepassing van controles voor de heterogeniteit van de hellingen zoals voorgesteld door Breitung Salish (2020), bij een analyse van de europeriode als geheel tot aanzienlijke steiler geschatte hellingen van de Phillips-curve leidt. Uit analyse van een subperiode blijkt dat de helling van de Phillips-curve van de eurozone steiler is geworden dan wat de schattingen van de gereduceerde vorm bij weglating van de controles van de heterogeniteit in de periode vanaf 2009 suggereren, en dat de versterkte regressor ter controle van de heterogeniteit van de hellingen van belang wordt.

About the Author

Johannes Schuffels was born on March 1, 1993 in Munich, Germany. He started studying economics at LMU Munich in 2011, completing his Bachelor's degree in 2015. Johannes continued his studies at Université Paris 1 - Panthéon Sorbonne. In 2017, he obtained his Master's degree in Financial Economics with *mention bien*. The same year, he started his doctoral studies at Maastricht University under the supervision of Prof. Clemens Kool, Prof. Tom van Veen and Dr. Lenard Lieb. This dissertation contains his research.

During his PhD, Johannes taught several undergraduate courses in macroeconomics and international economics at the School of Business and Economics as well as the Faculty of Arts and Social Sciences. He also represented the PhD students of the department for Macro, International and Labor Economics in the faculty's PhD Committee. In October 2019 he was co-organizer on a workshop on Globalisation and Structural Change at Maastricht University. Between March and July 2020, Johannes worked at the Directorate-General for Economic and Financial Affairs of the European Commission as a Blue Book Trainee.

In November 2020, Johannes joined the Organization for Economic Cooperation and Development (OECD) as Junior Economist in the Statistics and Data Directorate, where he contributed to the development of new methods for the compilation of regional price statistics on real estate. Since December 2021, Johannes works as Economic Analyst at the Directorate-General for Economic and Financial Affairs of the European Commission.