Expectations, Returns and the Macroeconomy

Franziska Hünnekes

Munich 2020

Expectations, Returns and the Macroeconomy

Inaugural-Dissertation
zur Erlangung des Grades
Doctor oeconomiae publicae (Dr. oec. publ.)
an der Ludwig-Maximilians-Universität München

2020

vorgelegt von
Franziska Hünnekes

Referent: Prof. Dr. Gerhard Illing
Korreferent: Prof. Dr. Gernot J. Müller
Promotionsabschlussberatung: 22.07.2020

Datum der mündlichen Prüfung: 06.07.2020
Berichterstatter: Gerhard Illing, Gernot J. Müller, Christoph Trebesch

Danksagung

Die letzten dreieinhalb Jahre waren wohl die lehrreichsten und spannendsten meines Lebens. Am Ende haben sie zu dieser Dissertation geführt. Möglich wurde das Ganze nur durch die Unterstützung zahlreicher Menschen, denen ich hier meinen Dank aussprechen möchte.

Zunächst geht ein großer Dank an meinen Betreuer Gerhard Illing. Mit Ihrer Bereitschaft meine Promotion an der LMU zu unterstützen, haben Sie mir den Weg gleich zu Anfang geebnet. Im Laufe der Zeit waren Sie dann einer meiner wichtigsten Berater, nicht nur für die wissenschaftlichen Fragen, sondern auch für die anderen Aspekte des Promotionsalltags. Ich danke Ihnen, dass Sie immer ein offenes Ohr für mich hatten und mich in allen Entscheidungen unterstützt haben.

Einen weiteren großen Anteil am Entstehen dieser Dissertation haben meine Koautoren. Ich danke Zeno Enders, Gernot Müller, Moritz Schularick und Christoph Trebesch für ihr Vertrauen in diese ihnen doch weitestgehend unbekannte Doktorandin. Die Arbeit mit euch hat mir große Freude bereitet, mich vieles gelehrt und mich bestens auf die Fertigstellung dieser Dissertation und alle weiteren beruflichen Herausforderungen vorbereitet. Für eure Unterstützung, auch über die Papiere hinaus, bin ich euch sehr dankbar.

Für die tägliche Arbeit waren natürlich auch meine Kollegen am Lehrstuhl entscheidend. Ich danke Sebastian Horn, Jonas Schlegel, Alexander Schramm, Alexander Schwemmer und Peter Zorn für spannende Diskussionen und regelmäßige Runden durch den Englischen Garten. Ein großer Dank geht auch an Agnès Bierprigl, für ihr offenes Ohr und die wunderbare Unterstützung bei allen administrativen und sonstigen Angelegenheiten.

Ein weiterer wichtiger Teil des Arbeitsalltags waren die anderen Doktoranden meiner Kohorte. Danke an euch für dreieinhalb Jahre konstante Mensa-Mittagessen am Donnerstag, Stefan für "Kaffee"-Pausen am ifo, die immer wieder zu neuen Einsichten bezüglich der ifo-Daten geführt haben, und Michael und Brendan für die technische Unterstützung (ohne euch wären meine Tabellen und Graphen höchstens halb so schön).

Für Kapitel 1 und 3 habe ich einen großen Teil meiner Promotionszeit am ifo, spe-

ziell am EBDC, verbracht. Danke an Heike Mittelmeier, Valentin Reich, Sebastian Wichert und Shuyao Yang für gefühlt unendliche Datenexporte, immer doch längere Öffnungszeiten und Flexibilität, wenn ich dann spontan noch ein weiteres letztes Mal kommen musste.

Ein großer Dank geht auch an die Egon-Sohmen-Stiftung für ihre Unterstützung, ohne sie hätte ich wohl keinen Platz an der LMU gefunden. Das Stipendium hat mir die wunderbare Möglichkeit eröffnet an der LMU zu promovieren und durch die flexible Gestaltung viel Freiraum für die Forschung geschaffen.

Mein größter Dank gilt meinen Freunden und meiner Familie. Meinen Freunden dafür, dass ihr mich so nehmt wie ich bin und mir immer zuhört, auch wenn wir mittlerweile doch unterschiedliche Sachen machen. Ich bin dankbar, dass ihr auch über zum Teil große Distanzen und lange Jahre immer zu mir gehalten habt. Und meine Familie ist seit ich denken kann meine größte Unterstützung. Dass ich das hier gewagt und geschafft habe, liegt in erster Linie an euch. Ihr habt nie in Frage gestellt, dass mein Weg mich weg von zuhause und zu solchen Dingen führen wird, und ihr habt nie Zweifel daran aufkommen lassen, dass das einzig richtige ist.

Und zu guter Letzt Markus: von Anfang hast du mich durch Höhen und Tiefen begleitet, und warst durch alles hindurch immer für mich da, ohne einen einzigen Zweifel. Ohne dich wäre diese Dissertation wohl nicht fertig geworden.

Franziska Hünnekes

München im März 2020

Contents

List of Figures

List of Tables

Preface

This thesis has two parts. Chapters 1 to 3 consider the role of expectations in the economy and their implications for monetary policy. Chapter 4 explores the profitability of Germany's capital exports. Both parts discuss topics highly relevant for economic policy. First, central banks are increasingly using the management of expectations as a new policy tool, but the exact effects are still little understood. Second, the size of Germany's current account surplus has been the topic of much debate, both inside and outside of Germany. In this context, the profitability of the capital side of the surplus is an important aspect, which has not yet received a lot of attention. The aim of this thesis is to provide insights into these issues using empirical analysis, and thus help inform the debates among economists and policy-makers.

In recent years, macroeconomists have (re-)discovered expectations, especially those of firms and households measured in surveys. One key reason are the inflation developments since the great financial crisis in 2008. Central bankers and macroeconomists alike were puzzled, first, by the lack of disinflation, and then by the lack of inflation during the recovery (see e.g., Coibion and Gorodnichenko 2015b). The lack of inflation, in turn, has forced central banks to come up with new tools as their policy rates hit the effective lower bound but inflation was still far below target. In addition to asset purchase programs, forward guidance measures were introduced to influence expectations of firms and households directly. Before delving into the details of these developments, I will first provide an overview of the role of expectations in macroeconomics in general.

Expectations are central to economics. Forward-looking behavior distinguishes humans from inanimate objects, and with that distinguishes economics from the natural sciences.[1] The importance of expectations for macroeconomics in particular was noted early on, for example by Keynes (1936). Expectations entered models formally in the 1950s and 1960s, usually in the form of adaptive expectations. This static approach to expectations was criticized first by Muth (1961) and then famously by Lucas (1972).

[1] As Bachmann (2019) puts it: "Particles, molecules, and stars do not look into the future or have their behavior influenced by the future" (p. 65).

These authors introduced the notion of model-consistent or rational expectations arguing that agents are sufficiently sophisticated to understand economic principles and to adapt to new developments.

Early versions of the rational expectations formulation focused on how agents use information to form expectations, but they still allowed for limits on the amount of information available to agents. However, over time the assumption of full information was added to the rational expectations framework, i.e. most authors assumed that agents are not only aware of the underlying structural relationships but are also able to find and process all relevant information immediately (Coibion and Gorodnichenko 2012). This assumption made research, which considers what actual expectations look like, seem unnecessary from a macroeconomic perspective. Thus, the introduction of rational expectations led to a sustained period, in which expectations were not considered much in macroeconomic research.

Only in the early 2000s, a renewed interest into deviations from this assumption developed, mainly on the theoretical side. This research was motivated by the discrepancy between the observed inflation dynamics and those predicted by models with nominal rigidities based on price stickiness. In particular, several papers introduced information rigidities while maintaining the rational expectations assumption.[2] With the introduction of these theoretical approaches, also some empirical literature reemerged, mainly driven by Carroll (2003). However, these studies focused on time-series data. The analysis of micro-level survey data remained very limited, in part due to skepticism towards survey data, especially in the US (Bachmann 2019).

Then the financial crisis of 2008 and its aftermath left many advanced economies in a new macroeconomic environment with sustained low inflation rates despite strong easing monetary policy. In light of these developments, expectations came back into focus. Central banks struggling with the effective lower bound of interest rates and an impaired transmission mechanism, introduced forward guidance as a tool to influence the expectations of key economic actors directly. Based on earlier research on liquidity traps (e.g., Eggertsson and Woodford 2003; Krugman 1998), central bankers and macroeconomists alike expected these policies to be quite effective, especially in the low interest rate environment. However, it quickly became clear that the effects were rather weak (Del Negro, Giannoni, and Patterson 2012), and that some of the announcements even lowered expectations, i.e. achieved the opposite of what was intended (Campbell, Evans, Fisher, and Justiniano 2012).

This observation initiated a careful reevaluation of how monetary policy interacts with expectations. Campbell et al. (2012) first introduced the idea that forward guid-

[2]The most relevant examples are sticky information by Mankiw and Reis (2002), and noisy information due to rational inattention (Sims 2003) and imperfect common knowledge (Woodford 2002).

ance – and central bank communication in general – can have very different effects depending on how it is perceived by economic agents. On the one hand, a clear announcement of a strong easing policy may lead firms and households to develop a more positive view of economic prospects as they expect monetary policy to support the recovery (Odyssean view). On the other hand, the same announcement may be perceived as revealing that the economy is doing worse than previously thought. In this case, firms and households will revise their expectations downwards (Delphic view). This view of central bank communication goes back to the idea that the central bank is better informed about the economy than other agents, turning its announcements into a source of information about the economic outlook (Romer and Romer 2000).

These discussions of forward guidance have led to a reconsideration of how monetary policy in general affects expectations. In particular, Melosi (2017) and Nakamura and Steinsson (2018) integrate the notion of Delphic effects into dynamic general equilibrium models. In both cases, any decision or announcement by the central bank carries information about the state of the economy. Following, Nakamura and Steinsson (2018), this effect is referred to as the information effect of monetary policy.

Importantly, models incorporating the information effect of monetary policy allow for deviations from the full information assumption. They thus also led to a renewed interest in the empirical analysis of expectations, especially of the effect of monetary policy on expectations. However, most studies so far consider only professional forecasters and financial markets (see e.g., Campbell et al. 2012; Nakamura and Steinsson 2018). While these actors are relevant, we know that their expectation formation process differs substantially from the one of firms and households, who are the ones setting prices and driving demand in the economy. Therefore, Chapters 1 and 2 consider the effects of monetary policy surprises on expectations of those two groups. Both chapters find evidence of the information effect.

An underlying assumption of these discussions about expectations and monetary policy is that expectations matter, i.e. that expectations reported in a survey are indeed relevant for economic decisions. There are different ways to test this assumption. For example, Coibion and Gorodnichenko (2015b) show that household inflation expectations from the Michigan Survey of Consumers improve the estimation of the Phillips curve and help solve the missing disinflation puzzle after the financial crisis. In addition, one can consider how expectations and actions interact directly in surveys. Coibion, Gorodnichenko, and Ropele (2020) and Coibion, Gorodnichenko, and Kumar (2018) show that inflation expectations do affect the decisions of firms by introducing information treatments to different surveys. In both studies, expectations are about aggregate consumer price inflation, and the sample period is limited to the recent years. Chapter 3 adds additional evidence by considering a long-run panel of firms and their

expectations about their own production. The chapter shows that these expectations also matter for the decisions of firms regarding prices and production.

In the following, I describe the first three chapters in more detail.[3] Chapter 1 considers how expectations of German firms react to monetary policy by the European Central Bank (ECB). We find evidence of non-linear effects. After a small tightening monetary policy surprise, i.e. an interest rate increase, firms expect their own prices and production to fall. This response is in line with standard theory, which assumes that higher interest rates dampen demand and lower inflation, and thus lower expectations. However, for larger surprises this is not the case, instead firms do not change expectations significantly, or even expect higher prices and higher production. The results are symmetric for easing surprises, i.e. interest rate decreases. The observed reaction to large surprises is in line with the information effect of monetary policy offsetting the standard effect. Specifically, we interpret this finding as firms only paying attention to the information component of the monetary policy announcement when the surprise is relatively large.

The results is this chapter are obtained using data from the ifo institute's survey of German manufacturing firms. The key advantage of our approach is that this survey provides information on the day each firm returns its questionnaire. Therefore, we can compare firms responding within two working days before the ECB announcement to firms responding within two working days after the announcement. This tight window reduces the likelihood that expectations change due to macroeconomic developments other than the monetary policy announcement. We measure monetary policy using high-frequency identified surprises in interest rates. Specifically, we use changes in overnight interest rates in a tight window around each ECB meeting to capture only the unexpected part of the monetary policy decision.

While our interpretation of the non-linear effects as evidence that the information effect matters more for larger surprises is compelling, we cannot directly provide evidence for this interpretation with the available data. Since the non-linear response has important implications for monetary policy, a better understanding of the information effect is necessary. In particular, it is interesting to see whether the results are limited to the special case of the ifo firm survey. Therefore, Chapter 2 considers whether households also respond in a similar fashion to monetary policy. In addition, this chapter provides more direct evidence for the presence of information effects by analyzing news reporting about monetary policy.

Specifically, Chapter 2 considers the effect of monetary policy on household infla-

[3]Chapters 1, 3 and 4 are based on joint work with coauthors. Therefore, in the text below I use plural pronouns when discussing these chapters. The coauthors are Zeno Enders and Gernot Müller for Chapters 1 and 3, and Moritz Schularick and Christoph Trebesch for Chapter 4.

tion expectations in the United States. Results differ from the ones in Chapter 1 in that the information effect dominates for US households. In particular, a tightening monetary policy surprise, i.e. an interest rate increase, leads to higher inflation expectations. In turn, an easing monetary policy surprise lowers expectations. Both results imply that households adjust inflation in reaction to the implied information about the current outlook for inflation, not in reaction to the effect of the changed monetary policy stance on future inflation. The effects are robust across many different specifications.

As in Chapter 1, I measure monetary policy using high-frequency changes in interest rates, this time around Federal Open Market Committee (FOMC) meetings. The data on inflation expectations is from the Michigan Survey of Consumers. The Michigan Survey has a limited panel dimension – a subset of households is interviewed twice. This allows me to control for unobserved household heterogeneity.

In order to verify that the positive response is indeed due to the information effect, I analyze newspaper reporting after FOMC meetings. The reason for this additional step is that we know from previous research that households do not closely follow monetary policy news and that they do not know a lot about central banks (Binder 2017b). Therefore, it is somewhat surprising that households react so significantly and so clearly in line with information effect to the monetary policy decision. One plausible driver of this effect is media reporting because we know household use this as an information source when forming expectations (D'Acunto, Malmendier, Ospina, and Weber 2019b). If newspapers report about the implications of the central bank announcement for the economic outlook in line with the information effect, households may notice this information and react accordingly.

I test this hypothesis by analyzing articles published in US newspapers within a week after FOMC meetings. In particular, I use text analysis to classify articles based on whether they refer to high or increasing inflation or whether they refer to low or decreasing inflation. I find that newspapers indeed report more about high or increasing inflation after a surprise tightening announcement. The reporting in turn significantly increases households' inflation expectations. These results confirm that information effects are driving the households' response. Thus, Chapter 2 shows that US households not only change their inflation expectations significantly after monetary policy announcements but also do so according to the information effect hypothesis. Given that this response is potentially contrary to the one intended, central banks need to consider the existence of these effects when designing communication strategies.

After analyzing the effect of monetary policy on expectations in the first two chapters, I turn to the effect of expectations on actions in Chapter 3. This chapter shows that, in response to their expectations for the following months, firms already change their price and production decisions in the current month. The effect is significant both

for firms whose expectations turn out to be correct ex-post and for firms whose expectations turn out to be incorrect. The latter indicates that noisy or undue expectations matter for firms as well. In addition, these undue expectations have a significant effect on aggregate production and prices.

The results in this chapter are based on the same data from the ifo survey of German manufacturing firms used in Chapter 1. In order to establish the causal effect of expectations on actions at the firm level, we rely on propensity score matching. This approach allows us to compare firms with the same observable characteristics but different expectations. Assuming we control for all joint drivers of expectations and outcomes, the difference in outcomes between firms with different expectations is only due to the latter. We find that optimistic firms (expecting a production increase) are more likely to report higher production and prices in the month the expectation was reported. Effects for pessimistic firms are symmetric for lower production and prices.

There are two reasons firms may react to their expectations in this manner. One is that they simply know something about their future, which is not yet reflected in the observable fundamental values, and adjust accordingly (news view). The other is, that they are just optimistic or pessimistic for some reason unrelated to future fundamentals and react based on this (noise view). We disentangle these two views by considering the ex-post forecast errors of firms. Specifically, we categorize firms as correct and incorrect based on their reported actual production in the next three months. We find that the majority of the effect observed before is due to the correct firms. However, also those firms, which turn out to be incorrect ex-post, are significantly more likely to change production and prices, albeit to a smaller extent. Only incorrectly pessimistic firms do not adjust prices.

In order to understand whether these micro-level effects have macroeconomic consequences, we construct aggregate measures of what we call "undue" optimism and pessimism, inspired by the micro-level analysis. Using local projections, we assess the effect of these two expectations measures on industrial production and the producer price index in Germany. We find that undue optimism indeed increases industrial production and the producer prices significantly. Undue pessimism on the other hand, has almost no significant effect.

In conclusion, Chapter 3 shows that expectations matter significantly for firm outcomes, and that this effect translates to the aggregate economy, at least for positive expectations. This highlights the importance of understanding how expectations are formed. It also implies that the effects of monetary policy on expectations found in the first two chapters are likely to lead to real effects.

The second part of this thesis studies the returns on foreign assets. The accumula-

tion of foreign assets is related to the current account. Mechanically, current account surpluses are accompanied by financial account surpluses, which in turn increase the net foreign asset position. In addition, increased trade activity is often accompanied by investment abroad, which also increases gross flows.

In the last decade, Germany has been running an exceptionally large current account surplus: net exports worth 8% of GDP are unprecedented both compared to Germany's history of sustained surpluses as well as internationally. This surplus has sparked a lot of debate, inside as well as outside of Germany. On the one hand, the surplus results from large exports, which are linked to a high degree of competitiveness (see e.g., Fuest 2017; German Council of Economic Experts 2017). On the other hand, low imports and high domestic savings are also important drivers of the surplus as well. In this view, the surplus is an indicator of low domestic demand and excess savings. In addition, it not only represents domestic imbalances but may also reinforce those in the euro area (see e.g., Bernanke 2015; Krugman 2017). While criticism of the surplus was initially rejected in Germany, the debate has now shifted towards a more skeptical view as well (for example, Board of Academic Advisors 2019; Weidmann 2019).

One aspect largely missing from the German debate, however, are the implications of the capital side of the surplus. Internationally, some authors have warned of the financial implications (Bernanke 2015), but within Germany, the effects of the significant build-up of both the net and the gross foreign asset position have been mainly ignored.[4] In Chapter 4, we add evidence on this issue by comparing the return on Germany's foreign assets to the returns of other countries on their foreign assets.

The international comparison of gross returns is a new approach. So far, the international finance literature has focused on comparing return differentials, i.e. the difference between what countries earn on their foreign assets and what they pay on their foreign liabilities. These comparisons were generally done in context of the so called exorbitant privilege of the United States – the ability to continuously earn substantially higher returns on its assets than it is paying on its liabilities (see for example, Curcuru, Dvorak, and Warnock 2013; Gourinchas and Rey 2007a; Lane and Milesi-Ferretti 2003; Meissner and Taylor 2008). Instead, we compare returns on gross foreign assets directly. The idea is that generally all advanced economies can invest in very comparable foreign assets, while the foreign liability composition may be due to specific country characteristics. Therefore, asset returns are more informative about the profitability of foreign investment as such.

Specifically, in Chapter 4, we compute returns on the gross foreign asset positions

[4]There are a few recent exceptions, in particular, Bundesbank (2018) and Fiedler, Görg, Hornok, Jannsen, Kooths, Marchal, and Potjagailo (2018). However, they focus on a much shorter time period and lack an international dimension.

of Germany and a set of other advanced economies using balance of payments data. The estimated returns reveal a significant difference between Germany and the other countries: since the 1980s, the German return was 2 percentage points lower on average. This difference is relatively constant over time, statistically significant and can mostly not be explained by commonly considered factors. In particular, we show that valuation changes due to exchange rate effects are similar across countries and thus do not contribute to the return differential. The composition regarding type of asset is more relevant – up to 25% of the observed difference can be explained by Germany investing more in low return categories. However, the majority of the difference between the return on foreign assets of Germany and the other countries is due to return differences within each asset category. Geographical composition and riskiness cannot explain these differences. We find the largest differential for foreign direct investment and equity. German returns are 3 and 4 percentage points lower, respectively. For the other two asset categories (debt and 'other investment'), we find only small and insignificant differences.

In addition, German foreign asset returns are also lower than returns earned on a portfolio of domestic German assets. The difference was 3 percentage points on average since the 1960s, and actually widened recently. Finally, we show that Germany's foreign investment does not provide additional benefits, such as consumption insurance and hedging against demographic risks.

In conclusion, Chapter 4 shows that German foreign asset returns have been low, both in international and domestic comparison. Most of the difference cannot be explained by commonly considered factors. Instead, we find particularly large residual differentials for foreign direct investment and equity. These two categories usually earn the highest return within the foreign asset position. The results may hint a systematic failure of German foreign investment, which should be explored further using more detailed data than available in the balance of payments.

This thesis contributes to two policy-relevant literatures. On the one hand, the expectations of firms and households are highly important for the transmission and effectiveness of monetary policy. On the other hand, the performance of Germany's investment abroad adds important arguments to the debates about the current account surplus in Germany.

Chapters 1 and 2 highlight the non-trivial effect of monetary policy on expectations. While monetary policy indeed significantly affects expectations, the effect may often be in the opposite direction of what was intended. This leads to important trade-offs that central banks have to consider. On the one hand, the management of expectations provides an additional powerful policy tool, which supports and expands the trans-

mission of monetary policy. On the other hand, it is ex-ante difficult to predict how private actors will perceive specific announcements and therefore how they will react to this. More research is needed to understand how monetary policy communication can be effective without leading to adverse effects via the information effect. Chapter 3 underscores the relevance of this question by confirming that firm expectations as reported in surveys do matter for firm decisions, and that the effects of expectations translate to the aggregate economy.

Chapter 4 on the other hand, provides evidence that Germany's capital exports are not performing well abroad. In light of these results, debates about the German surplus and the fiscal stance have to be reconsidered. Given the relatively good performance of domestic assets, more domestic investment, both private and public, may be beneficial. This will in turn potentially increase import demand and reduce the current as well as the financial account surplus. In addition, we need a better understanding of the reasons why German foreign investments are outperformed by those of other countries.

I hope this thesis will spark more debate on these issues and lead to more research to improve our understanding of the role of expectations in the economy as well as of the determinants of foreign asset returns.

Chapter 1

Monetary Policy Announcements and Expectations: Evidence from German Firms*†‡

*A version of this chapter has been published in the *Journal of Monetary Economics*, Vol. 108, p. 45–63, in December 2019, doi: https://doi.org/10.1016/j.jmoneco.2019.08.011.

†This chapter is based on joint work with Zeno Enders (University of Heidelberg) and Gernot Müller (University of Tübingen).

‡We thank our discussant Rüdiger Bachmann and the participants of the Annual Research Conference of the National Bank of Ukraine and the Narodowy Bank Polski 2019 "Central Bank Communications: From Mystery to Transparency" in Kiev for very valuable comments. We are also grateful to Kai Carstensen for sharing useful insights into the setup of the Ifo Business Survey, to Marek Jarociński and Peter Karadi for sharing their shock series, and to Refet Gürkaynak for providing us early access to the Euro Area Monetary Policy Database. Enders and Müller thank the DFG for financial support in the framework of the SPP 1859. We also gratefully acknowledge the technical support by the team of the LMU-ifo Economics & Business Data Center (EBDC) in Munich as well as Mark Brisske.

1.1 Introduction

Do monetary policy announcements affect firm expectations and, if so, how? In this study, we take up the question empirically as we analyze the effect of policy announcements by the European Central Bank (ECB) on expectations of German firms during the period 2004 to 2018. We rely on the Ifo Business Survey Industry (IBS) of firm expectations and three distinct measures of monetary policy innovations. First, we focus on the announcements of specific non-conventional policy measures in the period since the global financial crisis. Second, more broadly, we consider monetary policy surprises as captured by high-frequency interest rate changes around monetary policy events. Third, we employ more structural measures of central bank information and communication shocks. Throughout, we focus on how firm expectations of production and prices change in response to these measures.

The main result of our analysis is twofold: we find a) that monetary policy surprises do indeed affect firm expectations significantly and b) that they do so in a non-linear way. In the first part of our analysis, we show that the announcements of non-conventional policies by the ECB hardly affected firm expectations and, to the extent that they did, they *lowered* expectations of prices and production – even though these policies were arguably designed to be expansionary. Once we turn to monetary policy surprises more broadly defined, we find that they affect firm expectations significantly. Moreover, the effect is non-linear in the size of the surprise. A moderate surprise increase of the interest rate reduces firm expectations, while surprise reductions raise them – in line with conventional wisdom. The strength of the effect declines as the size of the surprise increases, both for positive and negative surprises. Very large surprises no longer affect firm expectations significantly. Lastly, we find that central bank information shocks also affect firm expectations. Yet, they cannot account for the non-linear effect of monetary surprises on expectations, and neither can direct central bank communication during monetary policy events.

Our findings are consistent with the notion that monetary policy announcements induce market participants to update their views about the economy quite generally and not only about monetary policy. Romer and Romer (2000) argue that this is to be expected if the information sets of the central bank and the public are not perfectly aligned. Indeed, they show that the forecasts by the US Federal Reserve, which are unknown to market participants in real time, are useful in predicting inflation, given professional forecasts of inflation. More recently, Melosi (2017) puts forward and estimates a model where monetary policy shocks operate via a "signaling channel" as they induce price setters to update their belief about the state of the economy. The analysis of Nakamura and Steinsson (2018) also lends support to an "information effect,"

according to which market participants update their beliefs about the natural rate in response to monetary surprises. Finally, Jarociński and Karadi (2020) use the stock-market reaction to monetary policy surprises in order to strip monetary surprises of their information content. In this way, they are able to measure central bank information shocks on which we rely in our analysis below.

Our results highlight a new margin along which the information effect may operate: the size of the monetary surprise. According to our estimates, firm expectations adjust to monetary policy surprises in conventional ways to the extent that the surprise is small. The effect becomes weaker as the size of the shock increases – possibly because in this case the information effect becomes stronger. This pattern may emerge because the attention that firms pay to the potential information conveyed by monetary policy announcements is endogenous to the size of the policy surprise: inattentiveness is strong for modest surprises, while firms become more attentive to the information content as the surprise gets bigger. We stress that while we find this explanation of our empirical results compelling, we cannot rule out other explanations based on our analysis.

It may seem that our results are in conflict with recent findings by Kumar, Afrouzi, Coibion, and Gorodnichenko (2015) and Coibion et al. (2018). Using survey data for New Zealand, they find that firms tend to be inattentive to monetary policy or, more specifically, that firm expectations about aggregate inflation are not well anchored and widely dispersed across firms. We make three observations to reconcile our findings with this evidence. First, our analysis concerns firm expectations about firms' own prices rather than about aggregate inflation. Second, we focus on the effect of interest rate changes, which are likely more relevant for firms' profits than aggregate inflation. It may hence be more worthwhile for firms to pay closer attention to them. Last, Coibion et al. (2018) find that inflation errors are less pervasive in the manufacturing sector and in the case when firms face many competitors. It is therefore noteworthy that our results are based on data from the highly competitive manufacturing sector in Germany.

Specifically, our analysis uses survey data for a large panel of German firms. Firms are surveyed on a monthly basis and asked about their expectations of future production and prices. Our sample runs from July 2004 to June 2018. Crucial for our identification strategy is the information about the specific day on which firms submit their responses online. This allows us to distinguish between firms who responded before a specific policy measure was announced, and those firms that have responded afterwards. In a nutshell, the difference in expectations across the two groups of firms provides a measure of the effect of the policy announcement. Moreover, our data set allows us to control for a large set of firm characteristics.

Given our sample of German firms, we focus on the policy announcements of the ECB. First, we consider announcements of non-conventional policies, starting with the one-year long-term refinancing operations (LTRO), announced in June 2009. The last announcement in our sample is the termination of the expanded asset purchasing program (APP) in June 2018. These announcements of non-conventional policies are often discussed in the context of "forward guidance" since they pertain mostly to future policies, even though this holds to some extent for conventional policy announcements as well (Gürkaynak, Sack, and Swanson 2005). Still, this policy dimension has become even more pertinent during the last decade as policy rates were constrained by the zero lower bound. Under these circumstances, central banks relied heavily on announcements in order to communicate unconventional policy measures and to manage the expectations of the public.[1]

According to our estimates, the effect of the ECB's non-conventional policy announcements on firm expectations is limited. Moreover, if there is an effect at all, it tends to be negative. That is, the announcements of policies, which are arguably meant to be expansionary, tend to reduce expectations of production and prices. Such a reduction may be surprising in light of conventional wisdom but can be rationalized through information effects. Similar effects have been observed for the US and have brought to the fore the potentially "Delphic" nature of forward guidance (Campbell et al. 2012).[2] In this case – rather than being perceived as a commitment to a future policy as "Odyssean" forward guidance would have it – market participants update their belief about the state of the economy. In light of the information implicitly conveyed by the policy announcements, the outlook appears less benign to market participants and they revise expectations accordingly.

The non-conventional policies that were announced during our sample period were quite exceptional, in terms of both their specifics and their scope. Hence, we turn to monetary policy surprises more broadly defined. In particular, we study the effect of interest rate innovations around monetary policy events. We obtain these innovations from the Euro Area Monetary Policy Event-Study Database (Altavilla, Brugnolini, Gürkaynak, Motto, and Ragusa 2019). A plain-vanilla regression of firm expectations on interest rate surprises yields a counterintuitive effect: it is weak but positive, even as we limit our analysis to firms that respond to the ifo survey in a narrow window around the monetary policy event. We conjecture that this result may be due to infor-

[1] At a theoretical level, the effectiveness of such announcements is still a subject of considerable controversy. The canonical New Keynesian model predicts that "forward guidance" generates implausibly large effects (Carlstrom, Fuerst, and Paustian 2015; Del Negro et al. 2012; McKay, Nakamura, and Steinsson 2016).

[2] Policy announcements related to non-conventional policy measures have lowered yields significantly (Altavilla and Giannone 2017; Krishnamurthy and Vissing-Jorgensen 2011), but failed to raise expectations of inflation and output growth (Del Negro et al. 2012).

mation effects that differ in strength depending on the size of the surprise. Indeed, once we slice the data accordingly, we find a highly non-linear relationship between monetary policy surprises and firm expectations. Moderate surprises move expectations significantly in the expected direction, large surprises – both positive and negative – do much less so.

To explore the role of information effects further, we turn, in a third step, to more detailed measures of central bank information and communication. Specifically, we rely on a measure of central bank information shocks by Jarociński and Karadi (2020). Positive central bank information shocks capture favorable news about the economy – revealed through monetary policy announcements. As we assess the effect of these shocks on firm expectations, we indeed find that they raise price expectations, and significantly so. The effect on output expectations is insignificant. However, we find that once we control for these shocks in our baseline regression, monetary policy surprises continue to have a non-linear effect on expectations. As we dissect our results further, we find that they are driven by the news conveyed in the press releases of the ECB, rather than in the communication during the press conference. This suggests that the (non-linear) information effect operates directly through the policy measure, rather than through the communication thereof.

In our view, our analysis makes two contributions. First, our results suggest that the potential information conveyed by a monetary policy announcement may be endogenous to the size of the policy change. Second, our analysis provides evidence that monetary policy announcements affect the expectations of individual firms. In this way, we address a certain shortcoming of existing work that has focused on the effect of monetary policy announcements on professional forecasters (Campbell et al. 2012; Del Negro et al. 2012). According to theory, monetary policy operates through the expectations of price setters. Hence, it is of particular interest to assess whether firms, rather than professional forecasters, respond to monetary policy announcements. Our results show that they do and non-linearly so.

More generally, our study relates to a number of recent contributions that provide new evidence on expectation formation. Coibion and Gorodnichenko (2012, 2015a) show that professional forecasters adjust forecasts only sluggishly to shocks. As for evidence on price setters, Boneva, Cloyne, Weale, and Wieladek (2020) use a panel of expectations of UK firms to estimate a version of the New Keynesian Phillips curve. Coibion et al. (2020) use a survey of Italian firms to estimate the causal effect of firm expectations on firm decisions.

There is also work on expectation formation based on the ifo survey. An early study by Nerlove (1983) finds evidence in support of an adaptive expectations model. More recently, Bachmann and Elstner (2015) show that at most one third of the firms

in the ifo survey systematically overpredict or underpredict their production growth one-quarter ahead. Massenot and Pettinicchi (2018), in turn, identify various factors, which account for forecasting errors of firms in the ifo sample. Buchheim and Link (2017) analyze to what extent the expectation formation of firms depends on aggregate rather than on firm-specific information. Enders, Hünnekes, and Müller (2019a) investigate the role of firm expectations on price setting and production. In a similar vein, a number of recent papers investigate the effects of uncertainty on production and price-setting decisions using data from the ifo survey. Bachmann, Elstner, and Sims (2013) show that surprise movements in uncertainty lead to significant reductions in production, while Bachmann, Born, Elstner, and Grimme (2019) find that the frequency of price adjustment increases in idiosyncratic business volatility.

The remainder of this chapter is structured as follows. The next section introduces our data set. In Section 1.3, we study the effect of the ECB's announcements of non-conventional policies on firm expectations. Section 1.4 presents results for how monetary policy surprises, more broadly measured, affect firm expectations. Section 1.5 analyzes the role of central bank information and communication in more detail. A final section offers conclusions.

1.2 Data

In what follows, we briefly describe our data set. First, we provide some details on the survey of German firms from which we obtain a measure of firm expectations. Second, we turn to the monetary policy announcements of the ECB, both the non-conventional policy measures announced since 2009 and a broader set of monetary policy surprises.

1.2.1 Firm expectations

For our analysis, we use the Ifo Business Survey Industry (IBS), maintained at the LMU-ifo Economics & Business Data Center (EBDC) in Munich (IBS-IND 2018a). The survey contains monthly observations for several thousand German firms from the manufacturing sector. In each month, firms report expectations regarding their production and their prices for the next 3 months. The answers are qualitative: production and prices may increase, not change, or decrease. In addition, the survey includes qualitative questions about realized production and price changes in the previous month. The survey has broad coverage in terms of German industry. It is also used to construct the ifo business climate index, a widely observed leading indicator for current and future economic activity in Germany (Becker and Wohlrabe 2008).

The unit of observation in the IBS is either a product or a plant, depending on

Figure 1.1: Ifo business survey, descriptive statistics

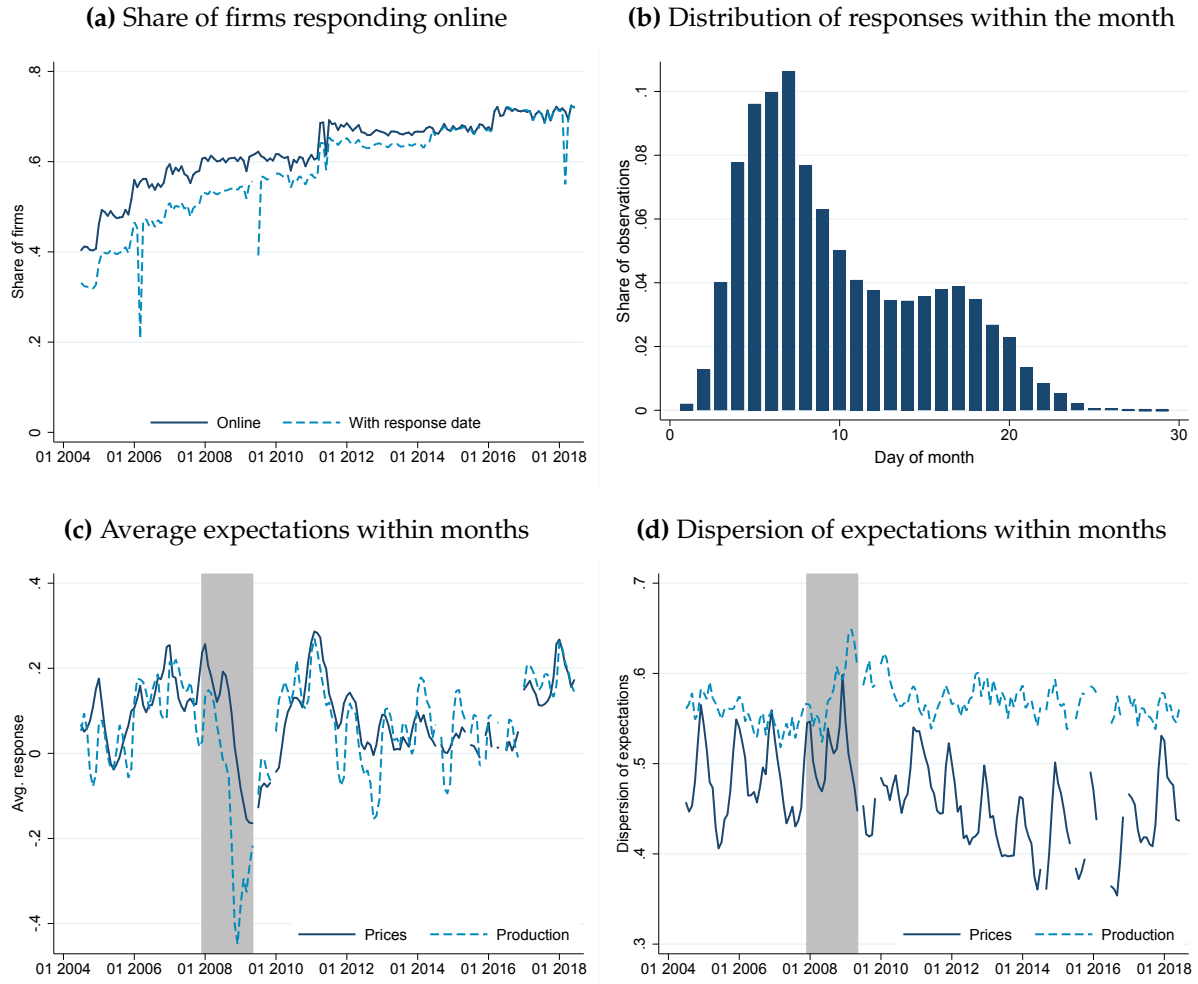

Notes: Descriptive statistics for our sub-sample of the ifo business survey (IBS). Panel (a) shows share of firms with a response date. Panel (b) reports response days within month. Panel (c) shows average expectations over time. Panel (d) displays dispersion measure suggested by Bachmann et al. (2013). Averages are arithmetic means, no weights used. Response dates are not available in the following months: 06-2009, 12-2009, 08-2014, 11-2015, 03-2016, 05-2016, 06-2016, and 12-2016. Shaded areas mark recession periods as defined by the German Council of Economic Experts.

the firm. As a result, some firms provide several responses per month. We conduct our analysis at the product/plant level and do not explicitly account for whether a product/plant is part of a multi-product firm. In our analysis below, we refer to the individual observation as a "firm" in order to simplify the exposition.

The IBS starts in 1980. However, only since 2004 can firms respond online to the survey. By now, the majority of firms use this option, as shown in panel (a) of Figure 1.1. Whenever firms answer online, the time and date of their response is recorded. The majority of firms respond in the first 10 days of the month. Panel (b) of Figure 1.1 displays the distribution of answers across the days of the month. We rely on the response date in our econometric strategy as explained below. For this reason, we limit

our analysis to those firms for which the response date is available. Our sample runs from July 2004 to June 2018. During this period, some 2300 firms fill out the survey in each month on average. Unfortunately, for a few months during our sample period no response time was recorded in the database. As a result, we have to drop eight months from our analysis below.[3]

In our analysis, we focus on expectations regarding the change in production and prices. The questions are as follows (our translation):

Q1 Expectations for the next 3 months: Taking changes of terms and conditions into account, our domestic sales prices (net) for product XY will probably increase [1], not change [0], or decrease [-1].

Q2 Expectations for the next 3 months: Our domestic production activity regarding good XY will probably increase [1], not change [0], or decrease [-1].

Note that these questions ask for qualitative answers and permit three realizations only: prices and production may increase, decrease, or remain unchanged. Consequently, whenever a firm reports to expect an increase in the previous period, the change in expectations cannot be larger than zero, and conversely for pessimistic firms. When estimating how expectations change in response to monetary policy announcements, we therefore control for the lags of expectations. Last, note that the IBS does not provide fixed-event forecasts (for instance, expectations regarding June 2009) but fixed-horizon forecast (for instance, expectations for the next 3 months going forward). The change in expectations is thus not a forecast revision in a strict sense, as the two forecasts do not pertain to the exact same period.

Panel (c) of Figure 1.1 displays the average responses to these questions for the months in our sample. The average expectation across firms fluctuates considerably over time, both for prices and production. The two time series tend to co-move, with some exceptions. For example, the pronounced downturn during the financial crisis was larger for production expectations than for price expectations, reflecting the "missing disinflation" (Hall 2011). More generally, we also observe that over time production expectations are more volatile than price expectations.

Panel (d) shows the dispersion of price and production expectations within each month.[4] The dispersion is generally larger for production expectations. During the crisis, the dispersion in both variables increased. After the crisis, however, it remained

[3]Specifically, we omit the following months: 06-2009, 12-2009, 08-2014, 11-2015, 03-2016, 05-2016, 06-2016, and 12-2016.

[4]Dispersion of expectations based on qualitative survey data is measured as $\sqrt{\text{frac}^+ + \text{frac}^- - \left(\text{frac}^+ - \text{frac}^-\right)^2}$, where frac^+ and frac^- are the fraction of positive and negative responses in each month, respectively. This measure is also used by Bachmann et al. (2013).

at a higher level for production expectations only. Finally, dispersion is more volatile over time for price expectations compared to production expectations.

In addition, we use answers to further survey questions as control variables. These include questions on past production, prices, and demand, as well as current orders and current capacity utilization. All questions, except for the one on capacity utilization, ask for qualitative answers with a similar answer structure as the price and production expectations questions. In each case, there are three possible answers: positive, neutral, and negative. Accordingly, we code the answers in a similar fashion as the price and production expectations. Table A.1 in Appendix A1 provides details on all questions.

Finally, our econometric strategy requires information about the day the firm responded to the survey. As discussed above, this date is available only for a subset of firms. We know that responses of the full sample of the IBS are useful indicators for the German economy. Therefore, we want to make sure that our subset of firms is not too different from the full sample. For this purpose, we compare the mean and standard deviation of all variables used. We find that firms responding online do not differ much from the full sample regarding their average responses. This is despite the fact that the sub-sample includes a larger share of firms from more recent years, since the share of firms with a response date increased over time. However, firms in our subset are slightly larger: the average number of employees in the full sample is 489 with a standard deviation of 3560, compared to 548 employees on average in our subset of firms, with a standard deviation of 3770. Table A.2 in Appendix A1 provides details on the descriptive statistics in both samples.

1.2.2 Monetary policy announcements

In our analysis, we use three distinct measures of monetary policy surprises, which are due to monetary policy announcements. First, we consider directly the announcements of unconventional policies by the ECB in the wake of the financial crisis. Second, we employ high-frequency changes in overnight-index-swap (OIS) interest rates around monetary policy events. Last, we rely on a decomposition of monetary policy surprises due to Jarociński and Karadi (2020). We briefly discuss these measures in what follows.

First, we identify 16 announcements of non-conventional policies by the ECB between May 2009 and June 2018, such as the announcement of the first long-term refinancing operations (LTROs), the different asset purchase programs, and the first forward guidance announcements. Our list of events largely follows Dedola, Georgiadis,

Gräb, and Mehl (2018).[5] Table A.3 in Appendix A1 provides an overview. These policy measures differ along a number of dimensions and we allow for different effects of each announcement in our analysis below. An aspect common to most measures is – with the notable exception of the OMT – that they brought about an expansion of the ECB's balance sheet. The ECB engaged, in other words, in "quantitative easing."

Additionally, we obtain a broader measure of monetary policy surprises from the Euro Area Monetary Policy Event-Study Database (EA-MPD) compiled by Altavilla et al. (2019). It records high-frequency changes of interest rates in a small window around monetary policy events. Given the small window size, these changes are likely to capture the surprise component of the monetary policy announcement relative to what market participants had expected prior to the event. The EA-MPD is a rich resource in that it provides data on changes of various interest rates and exchange rates for monetary policy events in the euro area. For each event, there are three event windows:

W1 Press-Release Window: change in the median quote during the interval from 13:25-13:35 before the press release to the median quote during the interval 14:00-14:15 after it.

W2 Press-Conference Window: change in the median quote during the interval from 14:15-14:25 before the press conference to the median quote for the interval from 15:40-15:50 after it.

W3 Monetary-Event Window: change in the median quote during the interval from 13:25-13:35 before the press release to the median quote for the interval 15:40-15:50 after it.

In our baseline analysis, we rely on the full Monetary-Event Window (W3) in order to capture the joint effect of the press release and the press conference. In Section 1.5, we consider the monetary policy surprise for each window in isolation. In the EA-MPD, interest rate surprises are measured by the change in the OIS rate for different maturities. In our analysis, we use 1-month OIS rates.

In our sample period, around 90% of the observed changes are smaller than 3 basis points in absolute terms, while 50% lie within -0.55 and 0.3 basis points, see Figure 1.2. For what follows, we stress that a surprise by 1 basis point is a relatively large surprise. The largest surprises are close to 20 basis points. We use dashed lines to indicate the surprises on the announcement dates of non-conventional policies. Roughly speaking, only half of them bring about a reduction of interest rates, the other half induces the

[5] Dedola et al. (2018) analyze the effect of quantitative easing measures and hence consider only policies that affect the ECB's balance sheet. Our focus is broader, such that we include additional announcements like the first forward guidance communication or the announcement of the Outright Monetary Transactions (OMT).

Figure 1.2: Monetary policy surprises, 2004 to 2018

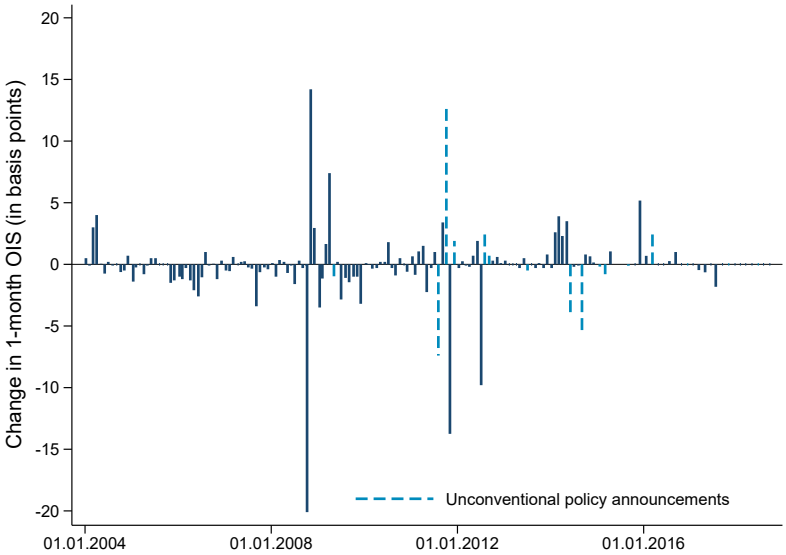

Notes: Changes in the 1-month OIS rate around ECB meetings, as provided by Altavilla et al. (2019). Days with unconventional announcements are indicated by dashed lines.

OIS rate to rise. In terms of size, some of the surprises associated with the announcements of non-conventional policies are particularly large, but we also note that some of the announcements brought about almost no change in OIS rates.

As discussed in the introduction, interest rate surprises are not necessarily pure monetary policy shocks. Monetary policy events may also induce a change in interest rates to the extent that central bank announcements relate news about the state of the economy to market participants. In our analysis below, we seek to account for this possibility in various ways. One strategy that we pursue is to rely on a decomposition of monetary policy surprises by Jarociński and Karadi (2020). They disentangle pure monetary policy shocks from unexpected information contained in the ECB's communication during the policy announcements based on high-frequency data and sign restrictions. First, they use high-frequency data to measure monetary policy surprises around monetary policy events, similar to those recorded in the EA-MPD. They also include 9 speeches of the ECB's president in their analysis. Based on this data set, the authors measure the reaction of stock prices (Euro Stoxx 50) and interest rates (3-month Eonia interest swaps) in a window from 10 minutes before to 20 minutes after the announcement. Next, they use sign restrictions to distinguish central bank information shocks from monetary policy shocks. They classify the shocks that move interest rates and stock prices in opposite directions as monetary policy shocks, in line with the received wisdom. Shocks that move both variables in the same direction are classified as central bank information shocks. The intuition is straightforward: because pure mon-

etary policy shocks should push interest rates and stock prices in opposite directions, any positive co-movement due to monetary policy surprises reflects new information about the economy released during the monetary policy event in question.

We now proceed in three steps. First, we investigate the effects of the specific non-conventional policy announcements of the ECB on firm expectations. We then consider how monetary policy surprises more broadly defined affect firm expectations. Last, we assess the role of monetary policy and central bank information shocks identified by Jarociński and Karadi (2020) as well as possible differences between the surprises measured during the press-release window (W1) and the press-conference window (W2).

1.3 Non-conventional monetary policy announcements

We now assess the response of firm expectations to the announcements of non-conventional monetary policies by the ECB. For this purpose, we compute the change in expectations in the month of the announcement relative to the previous month. To isolate the effect of an announcement, we consider a four-working-day window centered around the announcement and compare firms that reply within two working days after an announcement to firms that reply two days prior to the announcement.[6]

Our empirical strategy largely follows Del Negro et al. (2012). In contrast to their analysis, however, we focus directly on firms rather than on professional forecasters. Moreover, we are better able to capture the effect of an announcement because we have information about the timing of a firm's response. This allows us to focus our analysis on the expectations of those firms that respond within a narrow window around the announcement and, as a result, it is less likely that expectations change due to macroeconomic developments other than the monetary policy announcement. We assume throughout that the timing of a firm's response is unrelated to the announcement.

We pool the observations of firm responses between January 2009 and June 2018 and estimate the following model:

$$\Delta f(y)_{i,t} = \alpha + \sum_m \beta_m D_{i,m} + \delta_1 f(y)_{i,t-1} + \delta_2 Z_{i,t-1} + u_{i,t}.$$

Here, $f(y)_{i,t}$ are expectations of firm i regarding variable y in the next 3 months reported in month t (production or prices), and $\Delta f(y)_{i,t} = f(y)_{i,t} - f(y)_{i,t-1}$ is the change of the expectation between the current and the previous month. $Z_{i,t-1}$ is a vector of

[6]Since all ECB announcements occur on Thursday and almost no firm responds on the weekend, we consider all firms which answer between Friday and Monday following the announcement as being subjected to the announcement "shock".

control variables. It includes lagged realizations of expectations and several variables capturing economic activity at the firm level, such as the state of orders or capacity utilization.

We use the index m to refer to the announcements of the ECB. $D_{i,m}$ is a dummy variable that indicates whether expectations of firm i have potentially been affected by announcement m. We set $D_{i,m} = 1$ for those firms that respond within two working days after announcement m. Correspondingly, we set $D_{i,m} = 0$ for firms which respond two working days before the announcement. For example, the dummy variable for the introduction of 12-months LTROs, which was announced on May 7, 2009, is 1 for firms responding between May 8 and May 11, and 0 for firms responding on May 5 and May 6. We drop all firm observations on the day of the event as well as those outside the four-working-day window of an announcement in the month of the announcement, but include firm-month observations for those months without an announcement in order to enlarge the "control group." In this case, we set the dummy variable to zero. Our sample includes 16 announcements of non-conventional policies by the ECB. In what follows, we focus on 12 of these events, since 4 announcements either occur too early or too late in the month. In this case there are too few observations in the four-working-day window around the announcement, see Panel (b) of Figure 1.1.

Table 1.1 reports our estimates. Columns (1) to (3) display results for price expectations, columns (4) to (6) for production expectations. Consider first the results without time fixed effects, shown in columns (1), (2), (4), and (5). Here, we find that several announcements alter firm expectations significantly. When including the full set of controls (columns (2) and (5)), three announcements significantly affect price and production expectations. The announcement of 12-month LTROs in May 2009, one of the first measures with a large effect on the ECB's balance sheet (Dedola et al. 2018), *lowered* expectations regarding prices. Similarly, the announcement of 12/13-month LTROs in October 2011 *reduced* both price and production expectations. The announcement of the details of the OMT program also lowered production expectations, but there is no effect on price expectations. The 36-month LTROs announcement in December 2012 is an exception in that it raised price expectations.

In the specification with time fixed effects, shown in columns (3) and (6), we can no longer study the effect of four specific events since the announcement happened so early in each month that there are no observations in the first part of the event window. The remaining announcements are insignificant for price expectations, see column (3), but some announcements continue to have a significant negative effect on production expectations, see column (6).

These results are broadly consistent with earlier findings based on US data (Campbell et al. 2012; Del Negro et al. 2012). However, these studies assess the effect of an-

Table 1.1: Effect of unconventional monetary policy on firm expectations

	Dependent variable: change in the expectations for prices			production		
	(1) Baseline	(2) Further controls	(3) Time FE	(4) Baseline	(5) Further controls	(6) Time FE
12-month LTROs	-0.156***	-0.101***	-0.005	-0.140***	-0.066	-0.056
	(0.032)	(0.031)	(0.038)	(0.041)	(0.041)	(0.051)
6-month LTROs	-0.036	-0.034	-0.043	-0.046	-0.015	-0.025
	(0.027)	(0.026)	(0.031)	(0.036)	(0.035)	(0.041)
12/13-month LTROs	-0.029	-0.064**	-0.041	-0.136***	-0.153***	-0.080*
	(0.026)	(0.025)	(0.028)	(0.038)	(0.040)	(0.044)
36-month LTROs	0.070**	0.086**	0.056	-0.003	0.027	0.070
	(0.035)	(0.035)	(0.046)	(0.042)	(0.040)	(0.056)
OMT details	-0.054**	-0.038	-0.034	-0.192***	-0.135***	-0.123***
	(0.026)	(0.026)	(0.029)	(0.039)	(0.040)	(0.044)
Forward Guidance	-0.030**	-0.019		-0.005	0.001	
	(0.013)	(0.012)		(0.019)	(0.018)	
TLTROs	-0.070	-0.055	-0.023	-0.042	0.010	0.048
	(0.052)	(0.052)	(0.056)	(0.067)	(0.069)	(0.074)
ABSPP+CBPP3	-0.011	-0.006		-0.036*	0.008	
	(0.013)	(0.013)		(0.021)	(0.021)	
APP details	0.006	-0.003		0.028	0.030	
	(0.020)	(0.020)		(0.026)	(0.027)	
PSPP share limit	-0.027	-0.019		0.064**	0.101***	
	(0.017)	(0.017)		(0.031)	(0.033)	
APP end	0.034	0.028	-0.006	-0.013	-0.011	-0.055
	(0.028)	(0.033)	(0.048)	(0.043)	(0.045)	(0.067)
Expectation, t-1	X	X	X	X	X	X
Further Controls		X	X		X	X
Monthly time FE			X			X
Observations	236635	201212	201212	230028	197239	197239
Adjusted R^2	0.22	0.29	0.29	0.25	0.32	0.33

Notes: Results from regression of changes in expectations on dummy variables indicating monetary policy announcements. Only firms responding within a four-working-day window around the respective events are included. Firms from months without events are included as additional control observations. Firm expectations are obtained from IBS. Further controls included but not shown. For details, see Table A.8 in Appendix A2. Robust standard errors displayed in parentheses. * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

nouncements of non-conventional policies in the US based on the predictions of professional forecasters, rather than on the expectations of price setters. For the sake of comparability, we also turn briefly to professional forecasters' expectations in the euro area. Specifically, we focus on the revisions of forecasts for the HICP inflation and for real GDP growth in the Survey of Professional Forecasters (SPF) run by the ECB. The survey is conducted only at quarterly frequency during a number of specific days. The ECB publishes the time at which the participants receive the questionnaire as well as the deadline for handing it in. We use this information to create a dummy variable for each event. Specifically, only responses from the survey round directly following the announcement are allowed to be affected by the announcement.[7] Otherwise, our econometric specification follows Del Negro et al. (2012): to ensure that we do not capture other macroeconomic news revealed between two surveys, we control for a large set of macroeconomic surprises.[8] Overall, we find that professional forecasters in the euro area, just like firms in the German manufacturing sector, often reduced their forecasts in response to the announcements of non-conventional policies by the ECB. We report results in the appendix, see Table A.4.

In sum, there is little evidence that announcements of non-conventional monetary policies by the ECB raised firm expectations of production and prices, and similarly for professional forecasters. Remarkably, however, there are some instances where the announcements reduced expectations. To account for this observation, we note that some of the announcements of the ECB may have revealed bad news about non-monetary fundamentals. As discussed above, Campbell et al. (2012) obtain similar results for the US and rationalize these results based on the notion of "Delphic" forward guidance. Delphic forward guidance, as opposed to "Odyssean" forward guidance, does not involve a commitment about future policies but rather reveals information about the likely path of future policies given the policy maker's estimate of current and future non-monetary fundamentals.[9] Hence, the central idea that underlies Delphic forward guidance is the information effect of monetary policy surprises discussed in the introduction.

In concluding this section, we stress that the findings above have to be taken with a grain of salt due to two caveats. First, the announcements of the ECB may have

[7]For example, for the first 12-month LTROs announcement on May 07, 2009 the associated survey round is 2009Q3 which was conducted between July 15, 2009 and July 17, 2009. We provide a full list of the SPF rounds associated with each announcement in Appendix A1, see Table A.3.

[8]Macroeconomic surprises are measured as the difference between a macroeconomic release and the respective Reuters poll forecast before the release. Here our data source is Thomson Reuters Datastream.

[9]Just like us, Campbell et al. (2012) consider non-conventional policy announcements that have a flavor of forward guidance. The non-conventional policy measures which have been announced by the ECB do not qualify as forward guidance in a narrow sense (Eggertsson and Woodford 2003). Still, to the extent that quantitative easing operates through a signaling channel, it contains an element of forward guidance (Bauer and Rudebusch 2014).

been to some extent anticipated. As a result, the actual announcements of specific non-conventional measures may have fallen short of the expectations about their scope and strength entertained by market participants prior to the announcements. Indeed, in line with this conjecture, we observe positive interest rate surprises around some of the announcements, see Figure 1.2. In order to address this concern, we directly study the effect of interest rate surprises on firm expectations in our analysis below.

As a second caveat, note that we seek to capture the effect of a single monetary announcement on firm expectations within a four-working-day window. Our data set is unique in that it allows us to capture the effect of macroeconomic policy announcements on firm expectations in such a narrow window. It is nevertheless possible that expectations are also subject to other shocks within this window. As such, our estimates are potentially contaminated by noise. We run some placebo regressions to assess this hypothesis and find indeed some significant announcement effects for days where no announcements occurred.[10] In the next section, we will address this shortcoming, as we estimate the effect not of individual policy announcements but of a generic interest rate surprise of which there are many in our sample. This will allow us to estimate the effect of monetary policy announcements more precisely.

1.4 Monetary policy surprises

We now take a broader perspective on how monetary policy announcements affect firm expectations. We no longer focus on the effect of specific non-conventional policy measures, but assess how firm expectations respond to monetary policy surprises in general. For this purpose, we rely on the Euro Area Monetary Policy Event-Study Database (EA-MPD) provided by Altavilla et al. (2019). As explained in Subsection 1.2.2, it provides measures of monetary policy surprises that are constructed using high-frequency data. In our baseline specification, we measure monetary policy surprises as the change of the 1-month OIS rate in the window that spans both the press release and the press conference.

1.4.1 Linear effects

Our sample runs from July 2004 to June 2018. Within this period, there are 155 monetary events, that is, 155 meetings of the governing council of the ECB, followed by a press release and press conference. For 136 events, sufficiently many firms submit their survey response in the four-working-day window around the event. Other events are taking place during the first days or the last days of the month. In this case, we lack a

[10] Results are available on request.

sufficient number of firm responses within the four-working-day window around the event.[11] Furthermore, in some months no response dates are available for any firm, as discussed in Subsection 1.2.1. As before, in order to measure their effect on firm expectations, we contrast the responses of firms, which responded in the two working days after the event to those which answered in the two working days preceding it.

We again use $\Delta f(y)_{i,t}$ to denote the change in expectations regarding either prices or production reported by firm i at time t. Time is measured in months and there may be several monetary events within a month.[12] We use ε_m to denote the monetary surprise, measured by the change of the 1-month OIS rate at the event date in basis points. $D_{i,m}$ is a dummy variable, which equals 1 if firm i responded in the two working days after the monetary event m and 0 if it responded during the two working days prior to the event. We do not consider firms outside the windows. We pool observations across monetary events and estimate the following model:

$$\Delta f(y)_{i,t} = \alpha + \beta D_{i,m}\varepsilon_m + \delta_1 f(y)_{i,t-1} + \delta_2 Z_{i,t-1} + u_{i,t}. \tag{1.1}$$

Vector $Z_{i,t-1}$ includes additional control variables, such as the lagged dependent variable and the average state of business across all firms in the previous month to capture the state of the business cycle. Furthermore, we include additional variables to control for the firm's current situation. These are changes in prices or production and demand in the previous two months, the state of domestic and foreign orders in the previous month, as well as the capacity utilization in the previous month. All variables are coded on an ordinal scale with three outcomes, with the exception of capacity utilization, which is reported in percent. Table A.1 in Appendix A1 provides a detailed list of all questions.

Table 1.2 reports the results for price expectations in columns (1) and (2) and for production expectations in columns (3) and (4). The columns refer to alternative specifications in terms of control variables. The effect of monetary surprises on expectations is positive, but weak in the sense that the effect is only marginally significant. Nevertheless, taken at face value, this result implies that a surprise increase of interest rates raises price and production expectations, while a surprise reduction lowers expectations. This result conflicts with the received wisdom about the monetary transmission mechanism, but is in line with what we have established in the previous section regarding the effect of announcements of expansionary, non-conventional policies.

[11] See again panel (b) of Figure 1.1 for the average distribution of responses over the month.

[12] In our sample this happens only twice. In this case, we consider two windows per month.

Table 1.2: Effect of changes in the 1-month OIS rate

	Dependent var.: change in the expectations for prices		production	
	(1) Baseline	(2) Further controls	(3) Baseline	(4) Further controls
OIS, 1-month	0.0007	0.001*	0.002*	0.0001
	(0.0007)	(0.0007)	(0.0009)	(0.0009)
Expected prices, *t-1*	-0.454***	-0.578***		
	(0.004)	(0.006)		
Expected production, *t-1*			-0.495***	-0.622***
			(0.004)	(0.005)
Average state of business, *t-1*	0.134***	0.077***	0.133***	0.092***
	(0.008)	(0.009)	(0.011)	(0.012)
Further controls		X		X
Observations	65003	58779	62968	57379
Adjusted R^2	0.23	0.28	0.24	0.33
Observations before	31978	28761	30960	28058
Observations after	33025	30018	32008	29321

Notes: Results for regression of changes in expectations on monetary policy surprises in four-working-day windows around the respective events. Firm expectations are obtained from IBS. Monetary policy surprises are measured using high-frequency interest rate changes, taken from Altavilla et al. (2019). Further controls included but not shown. For details, see Table A.9 in Appendix A2. Robust standard errors displayed in parentheses. * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

1.4.2 Non-linear effects

The announcements of non-conventional policies had a relatively large effect on the 1-month OIS rate. On average, the changes of the OIS rate were twice as large as for conventional monetary policy events.[13] More generally, the size of the surprise varies considerably across monetary events. Against this background, it appears possible that the linear estimate of the response of expectations to monetary surprises masks some non-trivial heterogeneity. Specifically, large monetary policy surprises may induce a larger information effect that offsets the conventional effects of monetary surprises. Smaller monetary surprises may instead induce more conventional effects.

To assess this hypothesis formally, we rely on a modified version of equation (1.1). Specifically, we now sort monetary surprises prior to estimation according to their size. Next, we discretize the space of monetary surprises based on the distribution of interest rate changes over our sample.

[13] To see this please refer to Figure 1.2.

In particular, we estimate the following model with 7 bins, $b = 1, ..., 7$:

$$\Delta f(y)_{i,t} = \alpha + \sum_{b=1}^{7} \beta_b D_{i,m} \varepsilon_{m,b} + \delta_1 f(y)_{i,t-1} + \delta_2 Z_{i,t-1} + u_{i,t}. \tag{1.2}$$

Here, $\varepsilon_{m,b}$ assumes the value of the monetary surprise ε_m in case it falls into bin b. Otherwise, it is zero.

Figure 1.3 shows the results. In the top row, we display the effect of monetary policy surprises on expectations conditional on the size of the surprise. In Panel (a), we consider the effect on price expectations, in Panel (b) the effect on production expectations. In each panel, the horizontal axis indicates the bins (in basis points of interest rate changes) for which the estimates of β_b are displayed along the vertical axis. Each bin contains at least 5 different events, and several thousand firm observations. Recall that at high frequency, monetary policy surprises appear as somewhat moderate changes in the interest rate, with few changes exceeding $\pm$ 2 basis points (around 15%, see Subsection 1.2.2). In both panels, the horizontal line indicates the estimate that we obtain if we do not condition on the size of monetary surprises, reported in columns (2) and (4) of Table 1.2, with shaded area indicating 90% confidence bounds.

Our main result is that the effect of monetary policy surprises on firm expectations is now in line with the received wisdom, if the surprise is moderate. The estimated coefficients display a U-shaped pattern. A small surprise increase of the short-term interest rate lowers price expectations significantly. This may be because fewer firms expect price increases or because more firms expect price declines, or both. Below, we disentangle the effect using a binarized outcome variable. Note also that a small surprise reduction of short-term interest rates increases price expectations, according to our estimates. Importantly, however, we find that for *large surprises*, both negative and positive, the effect on expectations is no longer significantly negative and, at times, even positive. A very similar pattern emerges for production expectations, shown in Panel (b) of Figure 1.3. They respond significantly only to small monetary surprises.

In sum, we find that monetary policy affects firm expectations as conventional wisdom suggests. Interest rate increases (reductions) lower (raise) price and production expectations – but only for as long as the surprise is moderate. Larger surprises fail to affect firm expectations. This pattern is consistent with the notion that larger monetary policy surprises tend to trigger larger information effects. Consider a large surprise reduction of the monetary policy rate. While such a surprise should raise price and production expectations in the absence of information frictions, it may fail to do so if it induces firms to revise their assessment of the economy downward, because they realize that the central bank holds a more pessimistic view about the prospects of the economy than they did prior to the monetary announcement. Likewise, a large sur-

Figure 1.3: Effect of monetary policy surprises on firm expectations

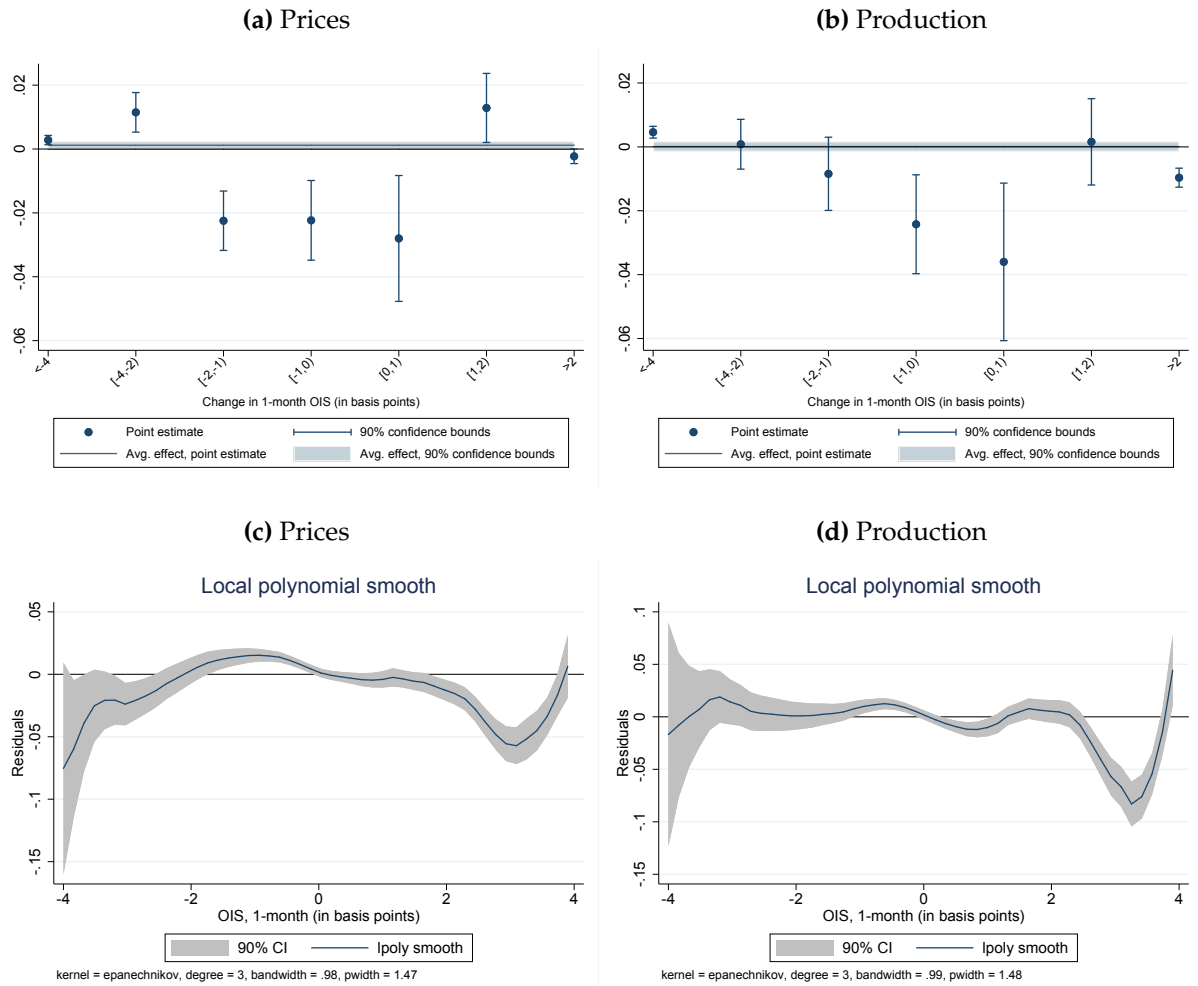

Notes: Effects estimated in four-working-day windows around monetary events. Effects are allowed to vary for different sizes of the surprise. Upper panels display estimates of 7 separate bins of surprise sizes. Horizontal line represents estimate based on linear specification. Bottom panels display kernel estimates, excluding largest/smallest OIS changes. Degree: 3. Firm expectations are obtained from IBS. Monetary policy surprises are measured using high-frequency interest rate changes from Altavilla et al. (2019).

prise increase of interest rates may convey a benign assessment of the economy by the central bank. As a result, it may not lower production and price expectations as a full information rational expectations model would predict.

Earlier work highlights the importance of information effects in accounting for monetary policy transmission (Jarociński and Karadi 2020; Melosi 2017; Nakamura and Steinsson 2018). Our results suggest a new margin along which the information effect may operate: the attention that firms pay to the potential information content of monetary policy announcements might be endogenous to the size of the policy change. Firms may pay little attention to the information content of a monetary announcement if the surprise is moderate. In a sense, firms operate in a business-as-usual regime –

they simply incorporate the direct effect of the interest rate change in their forecasts. As the surprise becomes bigger, however, firms might gradually pay more attention and reassess their views about non-monetary fundamentals in light of the monetary policy announcement.[14] In fact, our estimates suggest that the information effect can completely offset the conventional interest rate effect if the surprise is sufficiently large.

To investigate further the non-linear relationship between interest rate surprises and firm expectations, we also rely on a non-parametric approach. Specifically, in a first step, to control for factors other than monetary policy surprises, we estimate equation (1.1), except that we no longer include the interest rate surprise. We compute the residuals from this regression and estimate a kernel, which relates these residuals to the monetary policy surprises in a fully non-linear way. In doing so we employ an Epanechnikov kernel.

Figure 1.3 shows the results for price and production expectations in the bottom panels. Note that the kernel estimates in Panels (c) and (d) provide a direct measure of the *effect* of monetary policy surprises on expectations, rather than the *coefficient* estimates shown in Panels (a) and (b). That is, Panels (c) and (d) account for the fact that the monetary impulse increases as we move from the middle of the spectrum to the outer region in terms of the size of the monetary surprise. In the middle of the spectrum, the effect of monetary policy surprises is in line with conventional wisdom: positive interest rate surprises lower expectations and, conversely, rate reductions raise expectations. However, once we consider large surprises, the effect becomes weaker and even starts to reverse its sign for very large surprises. Again, this pattern is consistent with the notion that large monetary policy innovations trigger an information effect, which counteracts the direct interest effect.

In order to obtain a comprehensive measure of the non-linear relationship that characterizes the data, we estimate yet another model. For this purpose, we include a cubic term in equation (1.1). In this way, we can test for the significance of the non-linearity while economizing on the degrees of freedom. Specifically, we estimate the following model:

$$\Delta f(y)_{i,t} = \alpha + \beta D_{i,m}\varepsilon_m + \gamma D_{i,m}\varepsilon_m^3 + \delta_1 f(y)_{i,t-1} + \delta_2 Z_{i,t-1} + u_{i,t}. \tag{1.3}$$

Table 1.3 reports the estimates for this model, for price expectations in columns (1) to (3) and for output expectations in columns (4) to (6). For the specifications with the full set of control variables (columns (2) and (4)), we obtain a significant effect of the

[14] This notion is closely related to rational inattention as formalized by Sims (1998) and Maćkowiak and Wiederholt (2009, 2015). It is, however, distinct from these models as the attention is split between the "routine" monetary-policy part and the information content of interest rate changes. The former is more or less constantly monitored, the latter only during unusual episodes, see Nimark (2014) for a related setup.

Table 1.3: Effect of changes in the 1-month OIS rate, with cubic changes

	Dependent variable: change in the expectations for prices			production		
	(1) Baseline	(2) Further controls	(3) Ex. largest OIS	(4) Baseline	(5) Further controls	(6) Ex. largest OIS
OIS, 1-month	-0.0005	-0.0007	-0.0035**	-0.0018	-0.0038**	-0.0039*
	(0.001)	(0.001)	(0.0017)	(0.002)	(0.0015)	(0.002)
OIS, 1-month, cubic	0.074	0.114*	1.024***	0.210***	0.242***	1.005***
(coeff. & s.e.$\times 10^{-4}$)	(0.065)	(0.067)	(0.273)	(0.078)	(0.077)	(0.371)
Expected prices, *t-1*	-0.454***	-0.577***	-0.576***			
	(0.004)	(0.006)	(0.006)			
Expected prod., *t-1*				-0.495***	-0.622***	-0.622***
				(0.004)	(0.005)	(0.005)
Average state of business, *t-1*	0.135***	0.078***	0.083***	0.134***	0.094***	0.095***
	(0.009)	(0.009)	(0.009)	(0.011)	(0.012)	(0.013)
Further controls		X	X		X	X
Observations	65003	58779	56491	62968	57379	55155
Adjusted R^2	0.23	0.28	0.28	0.24	0.33	0.33
Observ. before	31978	28761	27395	30960	28058	26731
Observ. after	33025	30018	29096	32008	29321	28424
Excl. largest OIS changes			X			X

Notes: Results for regression of changes in expectations on monetary policy surprises and surprises cubed in four-working-day windows around the respective events. Firm expectations are obtained from IBS. Monetary policy surprises are measured using high-frequency interest rate changes from Altavilla et al. (2019). In columns (3) and (6) the largest four surprises are excluded (absolute value larger 10 basis points). Further controls included but not shown. For details, see Table A.10 in Appendix A2. Robust standard errors displayed in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

cubic term. For production we now also obtain a significant estimate of the coefficient β, which captures the linear effect.[15] The opposite sign of the linear and the cubic term imply that, as before, small positive interest rate surprises reduce expectations. For larger surprises, instead, the cubic term dominates and counteracts the conventional effect.

We also make sure that outliers do not drive our results: in columns (3) and (6), we report results for a sample excluding the four largest surprises (which exceed 10 basis points in absolute value). We find that the coefficients on the cubic terms are larger and in the case of prices also more significant.

Figure 1.4 visualizes the results reported in Table 1.3. We plot a straight line that

[15] If firms are weighted with the number of employees, all coefficients are highly significant.

Figure 1.4: Effect of changes in the 1-month OIS rate, with cubic changes

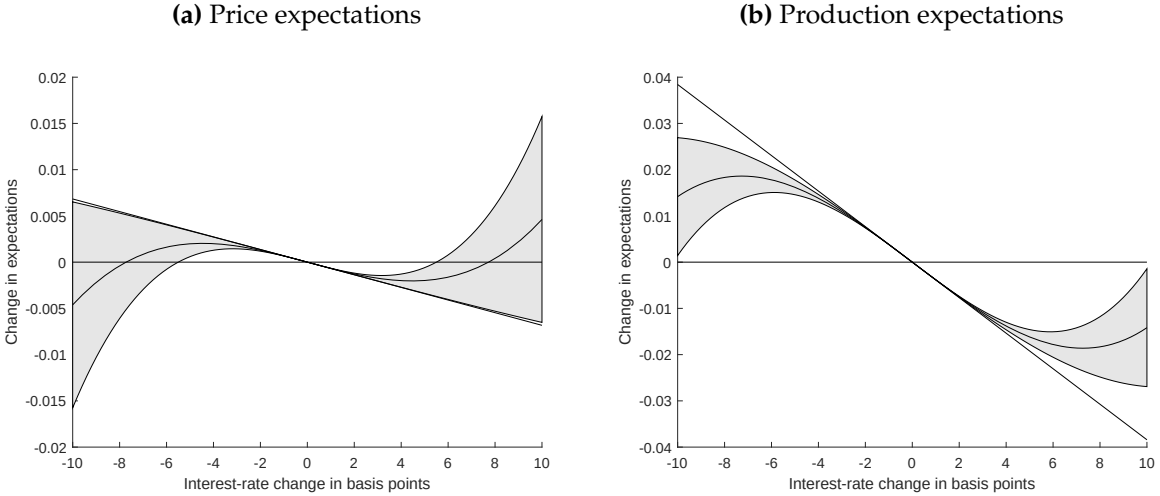

Notes: Graphical representation of results from the regression of changes in expectations on monetary policy surprises and surprises cubed in four-working-day windows around the respective events (Table 1.3, columns (2) and (5)). Straight line represents estimate of linear term. Shaded area indicates 90% confidence interval around the cubic component. Horizontal axis measures interest rate changes (bp), vertical axis measures change in expectations. Firm expectations are obtained from IBS. Monetary policy surprises are measured using high-frequency interest rate changes from Altavilla et al. (2019).

represents the linear effect, based on our estimate of β, see equation (1.3). In addition, we plot the total effect of monetary surprises, that is, the sum of the linear and the cubic terms, based on the estimate of β and γ, respectively. The shaded area indicates the 90% confidence intervals regarding our estimate γ. In this way, we can easily assess for which size a monetary surprise induces an effect that is significantly different from a linear response. We find this to be the case for monetary surprises larger than 4 basis points in absolute value for production expectations. For price expectations, the difference is only marginally significant.

In order to interpret the quantitative effect of monetary surprises on firm expectations, it is important to note that our dependent variable is measured on an ordinal scale with more than two possible answers. Hence, expectations may decline because fewer firms expect an increase or because more firms expect price declines, or both. To disentangle these effects, we proceed as follows. We create two new binary variables, which separately measure whether there was an upward or downward revision of expectations. In each case, the variable takes a value of 1 if there was a revision of expectations and 0 otherwise. We then estimate equation (1.3) again for both binary variables separately. The new models are linear probability models and the estimated coefficients correspond to the increase of the probability of an upward/downward revision in expectations following a monetary policy surprise by 1 basis point.

Table 1.4 shows the results for prices in columns (1) and (2), and for production in columns (3) and (4). It turns out that the change in expectations is driven more by

Table 1.4: Effect of changes in the 1-month OIS rate, binarized dependent variable

	Dependent var.: change in the expectations for			
	Prices		Production	
	(1) Upward rev.	(2) Downward rev.	(3) Upward rev.	(4) Downward rev.
OIS, 1-month	0.0001	0.0008	-0.001	0.002**
	(0.001)	(0.001)	(0.001)	(0.001)
OIS, 1-month, cubic	0.022	-0.101**	0.066	-0.158***
(coeff. & s.e.$\times 10^{-4}$)	(0.037)	(0.049)	(0.049)	(0.054)
Controls	X	X	X	X
Observations	58779	58779	57379	57379
Adjusted R^2	0.12	0.20	0.17	0.21
Observations before	28761	28761	28058	28058
Observations after	30018	30018	29321	29321

Notes: Results based on regression of changes in expectations on monetary policy surprises and surprises cubed in four-working-day windows around the respective events. Changes in expectations are binarized, i.e., two new variables separately indicate upward and downward revisions. Firm expectations are obtained from IBS. Monetary policy surprises are measured using high-frequency interest rate changes from Altavilla et al. (2019). Control variables included but not shown. For details, see Table A.11 in Appendix A2. Robust standard errors displayed in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

changes in downward revisions than upward revisions. Accounting for both the linear and cubic term, a monetary policy surprise of +1 basis point (bp) raises the probability by 0.08 percentage points (pp) that a given firm lowers its price expectations. For a surprise of +5bp the corresponding number is 0.27pp, while for a surprise of +10bp, it is -0.22pp. For production expectations the effects are larger: a surprise of +1bp yields a 0.22pp increase in the probability of a downward revision, a surprise of +5bp a 0.93pp increase, and a surprise of +10bp an increase of merely 0.66pp.

So far, we have been concerned with the response of firm expectations to monetary policy announcements. This response is of particular interest to the extent that firm expectations matter for firm actions. We take up the issue in related work as we investigate systematically the effect of firm expectations on firm actions as well as on aggregate outcomes (Enders et al. 2019a). The analysis uses the same data set as this chapter. We establish evidence that expectations do indeed matter strongly, both for firm actions and for aggregate outcomes.[16] However, in Enders et al. (2019a) we do not condition on monetary policy surprises. Therefore, we briefly take up the issue in the context of the present study. Specifically, we use the survey questions regarding

[16]In a related study, Balleer and Zorn (2019) investigate the effect of monetary policy on price-setting behavior of firms in the ifo panel. They also use the shocks identified by Jarociński and Karadi (2020). The authors furthermore analyze whether firm heterogeneity, for example regarding credit constraints, matters for the price-setting response.

realized changes in prices and production.[17] We replace the dependent variable in our baseline regression from equation (1.3) with the actual change of prices and production between the period of the monetary event and the following one. We find the same pattern as for expectations: positive (negative) interest rate changes lead to significantly lower (higher) prices and production, but in a non-linear way. The cubic term turns out to be positive and significant, see Table A.5 in Appendix A1.[18]

In a last experiment, we assess whether monetary policy announcements also affect firms' views on current business conditions. This may be expected if information effects play a major role. For this purpose, we focus on the current state of business as a dependent variable (Q9 in Table A.1). We find no significant effect of the interest rate surprises, neither linear nor non-linear. This suggests that the effect of monetary policy actions is delayed and that the information revealed by monetary policy announcements pertains to future realizations of non-monetary fundamentals rather than current ones. We do find, however, that firms' assessment of the current stock of inventories (Q12 in Table A.1) changes in response to monetary announcements in line with the response of expectations, although the effect of the cubic term is only marginally significant. We also report these results in Appendix A1, see Table A.6.

1.4.3 Robustness checks

Our results are robust to various alternative specifications. Table 1.5 reports estimates for a number of sensitivity analyses. It displays the results for prices in the upper panel, and those for production in the bottom panel. In order to streamline the exposition, we omit the estimated coefficients of the additional control variables in the table.

First check, we increase the window around the monetary policy event from two working days before and after the event to four working days before and after. Column (1) of Table 1.5 shows that this does not change our results much. Also, once we employ the "full window" and consider the responses of all firms that have been submitted after an event (prior to the next event), we still find a highly significant effect of the cubic term and all linear terms are highly significant as well. Column (2) of Table 1.5 shows the results.

Columns (3) to (5) report the estimates once we control for distinct features of the macroeconomic environment. First, we exclude the period where monetary policy in the euro area was close to the zero lower bound on interest rates from our sample. That is, we estimate our model on observations for the period from 2004 to 2011 only.

[17] These questions are listed in Table A.1 as Q3 and Q4 in the appendix. They are also qualitative in nature.

[18] Note that monetary events affect all firms, independently of whether they answer before or after the event, because all firms may adjust production and prices afterwards.

Table 1.5: Robustness checks

	(1) 8-working-day window	(2) Full window	(3) 2004-2011	(4) Crisis dummy	(5) Uncertainty	(6) Firm fixed effects	(7) Std. errors clustered by firm	(8) Only firms expecting no change in t-1
	Panel (a): Dependent variable: change in the expectations for prices							
OIS, 1-month	-0.0008	-0.0016***	-0.0009	-0.0008	0.000004	-0.0006	-0.0007	-0.0009
	(0.001)	(0.0005)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
OIS, 1-month, cubic	0.138**	0.080***	0.121	0.129	0.175*	0.100	0.114*	0.172***
(coeff. & s.e.$\times 10^{-4}$)	(0.067)	(0.021)	(0.085)	(0.085)	(0.090)	(0.069)	(0.067)	(0.057)
OIS, 1-month				-0.0006				
$\times$ crisis dummy				(0.002)				
OIS, 1-month					-0.002			
$\times$ uncertainty dummy					(0.002)			
Controls	X	X	X	X	X	X	X	X
Observations	72013	188211	24329	58779	58779	58779	58779	45258
Adjusted R^2	0.28	0.29	0.29	0.28	0.28	0.32	0.28	0.10
Observations before	41939		13905	28761	28761	28761	28761	22209
Observations after	30074		10424	30018	30018	30018	30018	23049
	Panel (b): Dependent variable: change in the expectations for production							
OIS, 1-month	-0.004**	-0.005***	-0.010***	-0.006***	-0.001	-0.004**	-0.004**	-0.002
	(0.002)	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
OIS, 1-month, cubic	0.224***	0.342***	0.448***	0.506***	0.453***	0.237***	0.242***	0.149*
(coeff. & s.e.$\times 10^{-4}$)	(0.075)	(0.025)	(0.094)	(0.104)	(0.102)	(0.077)	(0.077)	(0.082)
OIS, 1-month				-0.011***				
$\times$ crisis dummy				(0.003)				
OIS, 1-month					-0.008***			
$\times$ uncertainty dummy					(0.003)			
Controls	X	X	X	X	X	X	X	X
Observations	70239	184184	23827	57379	57379	57379	57379	37627
Adjusted R^2	0.33	0.33	0.34	0.33	0.33	0.37	0.33	0.12
Observations before	40864		13625	28058	28058	28058	28058	18247
Observations after	29375		10202	29321	29321	29321	29321	19380

Notes: Table shows several robustness checks for the non-linear effects of monetary policy surprises on firm expectations. Firm expectations are obtained from IBS. Monetary policy surprises are measured using high-frequency interest rate changes from Altavilla et al. (2019). In columns (1) and (2) window size varies. In columns (3) to (8) we use the four-working-day window. "Crisis dummy" is 1 from 01/2008 to 04/2009 and 0 otherwise, based on data from the German Council of Economic Experts. "Uncertainty dummy" is 1 if VSTOXX_{t-1} is larger than its sample mean plus one standard deviation, and 0 otherwise. The volatility series is the V2TX series from www.stoxx.com. Full set of control variables included but not shown. Robust standard errors in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

We find that the effects for production are larger and more significant in this case, while the effects for prices are roughly unchanged but turn insignificant. In column (4) we show results based on a specification which features a dummy variable to account for the financial crisis in 2008 and 2009. Again, we find that results for prices are weaker than in the baseline, while effects for production are strong, in terms of both size and significance. Finally, we also control for economic uncertainty as measured by a stock market volatility index (VSTOXX). Specifically, we capture periods of high uncertainty by a dummy variable, which equals one whenever the volatility index exceeds the sample mean plus one standard deviation. We interact the linear change in the OIS with this dummy variable. We find that our results regarding the cubic term are robust to accounting for economic uncertainty and conclude that the non-linear effect of our baseline specification is not merely capturing increased uncertainty or other features of the macroeconomic environment.

In the last three columns of Table 1.5, we provide results for additional robustness checks. First, we include firm fixed effects and report results in column (6). We find our estimates largely unchanged relative to the baseline, but the cubic term for prices becomes insignificant. Still, results for production remain highly significant. Column (7) shows results for a specification where we cluster standard errors at the firm level in order to account for potential correlation within firms over time. It turns out that this is not consequential for the precision of our estimates. Finally, we want to make sure that the qualitative nature of our dependent variables does not drive the results. One issue in this regard is that the responses of firms which expect an increase (or decrease) in the previous period are constrained in that they cannot report an even more optimistic (or pessimistic) outlook. Therefore, we estimate a specification in which we only include firms that expected no change in the previous month. The results in column (8) are fairly close to the baseline.

We further explore the relevance of working with qualitative variables as we turn to an additional variable. The ifo survey also features a question on the expected state of business of the responding firm. This question is answered twice using different measures. Once, firms simply choose from the qualitative answers as for the other questions (improve/stay the same/worsen). The second time, firms choose a scalar value between 0 and 100 by moving a slider, but without seeing the full scale. Instead, they only see marks at 0, 50, and 100, as well as that these values indicate "worsen", "stay the same", and "improve", respectively. We estimate a version of our model with this variable and find that results for expectations regarding the expected state of business expectations are very similar to production expectations. We provide details in Table A.6 in Appendix A1. We conclude that the qualitative nature of the data does not drive our results for the baseline case.

1.5 Central bank information and communication

The analysis in the previous section shows that firm expectations respond to monetary policy surprises. Moreover, the response depends on the size of the surprise. Building on earlier work that has established the information effect of monetary policy surprises, we put forward the following hypothesis to account for our finding: the extent of attention that firms pay to the potential information content of monetary policy announcements increases in the size of the policy change. As a result, the direct effect of a monetary surprise may be partially offset by the information effect, or even completely, if the surprise is very large. In what follows, we shed further light on the relationship between monetary policy surprises and firm expectations. In particular, we assess whether available measures of central bank information and communication can account for the observed effects. In a first step, we rely on the series of central bank information shocks by Jarociński and Karadi (2020). In a second step, we use the information in the Euro Area Monetary Policy Event-Study Database (EA-MPD) about the timing of information releases by the central bank to disentangle effects of monetary policy decisions as such and central bank communication.

1.5.1 Central bank information shocks

Under our maintained hypothesis, central bank information is crucial to account for our findings. We thus investigate whether a more direct measure of central bank information has a bearing on the expectation formation of the firms in our sample. Jarociński and Karadi (2020) propose such a measure by decomposing monetary policy surprises into monetary policy shocks and central bank information shocks, as explained in Subsection 1.2.2 above. We use these measures rather than overall monetary policy surprises and estimate the following variant of equation (1.1):

$$\Delta f(y)_{i,t} = \alpha + \beta D_{i,m}\varepsilon_m^{MP} + \gamma D_{i,m}\varepsilon_m^{CBI} + \delta_1 f(y)_{i,t-1} + \delta_2 Z_{i,t-1} + u_{i,t}. \tag{1.4}$$

Here, ε_m^{MP} and ε_m^{CBI} denote monetary policy and central bank information shocks, as identified by Jarociński and Karadi (2020). These shocks are generated regressors. Still, as pointed out by Coibion and Gorodnichenko (2015a), the standard errors on the generated regressors are asymptotically valid under the null hypothesis that the coefficient is zero (Pagan 1984). The dummy variable $D_{i,m}$ is specified as in the previous section: it captures whether a firm has filed response to the survey two working days before or after the monetary event.

Our sample runs from June 2004 to November 2016. For this period, Jarociński and Karadi (2020) identify 182 shocks. We keep these shocks in our sample to the extent

Table 1.6: Effect of 1-month OIS changes and Jarociński-Karadi shocks

	Dep. var.: change in the expectations for prices		production	
	(1) CBI	(2) CBI+MP	(3) CBI	(4) CBI+MP
OIS, 1-month	-0.002	-0.003*	-0.005***	-0.002
	(0.001)	(0.0016)	(0.002)	(0.002)
OIS, 1-month, cubic	0.137**	0.157**	0.261***	0.226***
(coeff. & s.e.$\times 10^{-4}$)	(0.069)	(0.071)	(0.079)	(0.082)
Central bank information shock (CBI)	0.004*	0.006**	0.003	0.0005
	(0.002)	(0.003)	(0.003)	(0.003)
Monetary policy shock (MP)		0.004		-0.006*
		(0.003)		(0.004)
Controls	X	X	X	X
Observations	56109	56109	54754	54754
Adjusted R^2	0.28	0.28	0.33	0.33
Observations before	26706	26706	26046	26046
Observations after	29403	29403	28708	28708

Notes: Results based on regression of changes in expectations on monetary policy surprises and monetary policy shocks (standardized) provided by Jarociński and Karadi (2020). Firm expectations are obtained from IBS. Monetary policy surprises are measured using high-frequency interest rate changes from Altavilla et al. (2019). Control variables included but not shown. For details, see Table A.12 in Appendix A2. Robust standard errors displayed in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

that we are able to define a four-working-day window around each shock.[19] As we estimate equation (1.4), we use, as before, the change in expectations regarding future prices and production as the dependent variable. We find that monetary policy shocks do not affect price and production expectations significantly. Central bank information shocks, instead, raise price and production expectations. Table A.7 in Appendix A1 shows the results.

Against this background, we explore whether central bank information shocks can explain the non-linear effect of monetary policy surprises on expectations, documented in Section 1.4 above. To this end, we include the identified shocks of Jarociński and Karadi (2020) as additional regressors in equation (1.3). Table 1.6 reports the results. We find that, while the central bank information shock has a significant and positive effect on price expectations, the cubic term capturing the monetary surprise as such

[19]There are more events in total despite the shorter sample period compared to the EA-MPD because Jarociński and Karadi (2020) also include speeches in addition to governing council meetings. For 159 of these events, we have sufficiently many firms responding in the four-working-days window around the event. For 127 events we have both OIS surprises from the EA-MPD as well as the shocks provided by Jarociński and Karadi (2020).

remains nevertheless highly significant. To illustrate this result further, we again plot the non-linear effects. The effects are very similar to the baseline case, as can be seen in Figure A.1 in Appendix A1.

We conclude that, while central bank information shocks play an important role for firm expectations, they cannot explain why the effect of monetary policy surprises on firm expectations depends on the size of the surprise. This might be a result of the identifying assumption maintained by Jarociński and Karadi (2020). Recall that their decomposition assumes that monetary policy shocks and central bank information shocks are orthogonal. Our results, instead, suggest that the amount of attention that firms pay to new information depends on the size of the monetary surprise. Under this hypothesis, we would expect a systematic link between monetary policy shocks and central bank information shocks. Jarociński and Karadi (2020) ruled out such a link.[20]

1.5.2 Central bank communication shocks

The EA-MPD provides distinct measures for interest rate surprises for each monetary policy event: one for a window around the press release and one for a window around the press conference of the ECB. The press release contains little more than the monetary policy decision as such. The market reaction in response to the press release thus provides a natural measure of the monetary policy surprise. Instead, the surprise captured by the interest rate change within the press conference window should more directly reflect the effect of central bank communication, net of the effect of the pure interest rate change (since this has been released earlier).

In principle, a systematically different communication in the press conferences after large surprises in central bank rates could explain the non-linear effects of monetary policy surprises on firm expectations. To give an example, if the ECB regularly justified large reductions in interest rates with a gloomy view on the economy, this would counteract the stimulating effects of the rate cuts. To explore this hypothesis systematically, we run the regression from equation (1.3) again but for different monetary policy windows. The first regression includes the change (linear and cubic term) in the 1-month OIS rate that occurs during the time window around the press release, while the second regression features the change in the same variable in the window around the press conference.

Table 1.7 shows the results for price and production expectations. Columns (1) and (3) report results for the press release, columns (2) and (4) for the press conference.

[20] In fact, we find that while monetary policy and central bank information shocks are uncorrelated by construction, they turn out to be correlated in a non-linear way: the correlation between the monetary policy shock cubed and the central bank information shock cubed is 0.68 and significant at the 1% level.

Table 1.7: Effect of changes in the 1-month OIS rate, press release and press conference window separately

	Dependent variable: change in the expectations for			
	prices		production	
	(1) Release	(2) Conference	(3) Release	(4) Conference
OIS, 1-month	-0.002	-0.007**	-0.007***	-0.0003
	(0.002)	(0.003)	(0.002)	(0.004)
OIS, 1-month, cubic	0.169**	2.980***	0.363***	1.110
(coeff. & s.e.$\times 10^{-4}$)	(0.077)	(0.881)	(0.088)	(1.139)
Controls	X	X	X	X
Observations	58779	58779	57379	57379
Adjusted R^2	0.28	0.28	0.33	0.33
Observations before	28761	28761	28058	28058
Observations after	30018	30018	29321	29321

Notes: Results for regression of changes in expectations on monetary policy surprises and surprises cubed in four-working-day windows around the respective events, separately for surprises from windows around ECB press releases and ECB press conferences. Firm expectations are obtained from IBS. Monetary policy surprises are measured using high-frequency interest rate changes from Altavilla et al. (2019). Control variables included but not shown. For details, see Table A.13 in Appendix A2. Robust standard errors displayed in parentheses. * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

The cubic term remains significantly positive in all specifications, except for production expectations in the case of the conference window. We also provide a graphical illustration in Appendix A1, see Figure A.2. The responses to the press release exhibit a very similar pattern as in our baseline findings, shown in Figure 1.4 above. In contrast, the information released during the press conference seems to trigger quite distinct reactions.[21] These results suggest that the information effect of monetary policy surprises is not confined to the press conference. Instead, it appears that monetary surprises as such may induce an information effect to the extent that they are large.

1.6 Conclusion

We have asked whether monetary policy announcements affect firm expectations and, if so, how. This issue is of particular importance because for policy announcements to be effective, theory requires them to affect expectations of price setters, that is, firms.

[21] Conrad and Lamla (2010) show that for a given monetary policy tightening or easing, financial-market reactions, as measured by the exchange-rate response, depend on the specific topic of communication during the ECB press conference. This might explain why we do not find a clear pattern for the conference window.

Yet, most evidence regarding the effect of monetary policy on expectations is confined to professional forecasters. In this analysis, we focus directly on firm expectations using a uniquely suited data set. It contains survey responses of several thousand firms in the German manufacturing sector. Our sample runs from 2004 to 2018. Observations are monthly, but we also know the calendar date at which firms file their responses. Our empirical strategy relies on this specific feature of the data set.

In a first step, we conduct an event study: we estimate the effects of the announcements of non-conventional policies by the ECB since the crisis. We compare the responses of firms in a four-working-day window around the announcement and find no significant effects on firms' price and production expectations, except for a few instances where expectations are revised *downwards*. Taken at face value, this result is surprising. However, similar findings for professional forecasters in the US have been rationalized on the ground that monetary policy announcements by the Federal Reserve may have revealed bad news about the economy. In the context of non-conventional policy announcements, this phenomenon has given rise to the notion of "Delphic" forward guidance, as opposed to "Odyssean" forward guidance. Delphic forward guidance, in turn, relates to the broader concept of information effects, which may at times offset the conventional effect of monetary policy measures.

In a second step, we explore this issue further as we rely on a broader measure of monetary policy surprises, compiled using of high-frequency data and available in the Euro Area Monetary Policy Event-Study Database. We estimate the responses of firms' expectations to monetary policy surprises and find that they affect firm expectations significantly. A second important finding is that the effect of policy surprises depends on their size. Moderate surprises affect firm expectations in line with standard theory: interest rate increases lower price and production expectations, while interest rate reductions raise them. Large changes, however, have no significant effect. This finding is consistent with information effects, provided that such effects operate along a specific margin, namely the size of the monetary policy surprise.

In the last part of our analysis, we shed more light on this possibility and investigate the effect of central bank information and communication on firm expectations. First, we consider the time series of central bank information shocks computed by Jarociński and Karadi (2020). We find that firms revise their price expectations upward in response to positive shocks, even if they appear in the context of surprise tightenings of the policy rate. However, once we include central bank information shocks in our baseline model as an additional control variable, the non-linear relationship between the size of a monetary policy surprise and its effect on firm expectations remains intact. Second, we distinguish between monetary surprises due to the ECB's press releases and the ECB's press conferences. We find that our main result continues to hold once

we consider only monetary policy surprises around the press release window.

Our results are likely to have a bearing on actual policy design. First, they suggest that it may be beneficial to separate interest rate decisions from central bank communication about the state of the economy as much as possible. Wiederholt (2015) provides a recent analysis of the effectiveness of central bank communication in an environment with dispersed information and conventional monetary policy constrained by the zero lower bound. Second, our results also underscore the specific challenges for the conduct of monetary policy in the presence of information frictions. Jia (2019) performs a model-based analysis and finds that stabilizing the economy becomes more difficult for monetary policy if private agents extract information about non-monetary fundamentals from policy decisions. Our results point in the same direction. Central banks may face a specific dilemma in the presence of information frictions: there might simply be no way to justify drastic policy measures without generating attention for the non-monetary fundamentals, which motivate the policy, rendering large policy changes ineffective. However, we stress that at this point our results are based on a purely empirical analysis and call for a further empirical and model-based analysis before they can inform actual policy making in a reliable manner.

Appendices to Chapter 1

A1 Additional figures and tables

Figure A.1: Effect of changes in the 1-month OIS rate, with Jarociński-Karadi shocks

(a) Price expectations

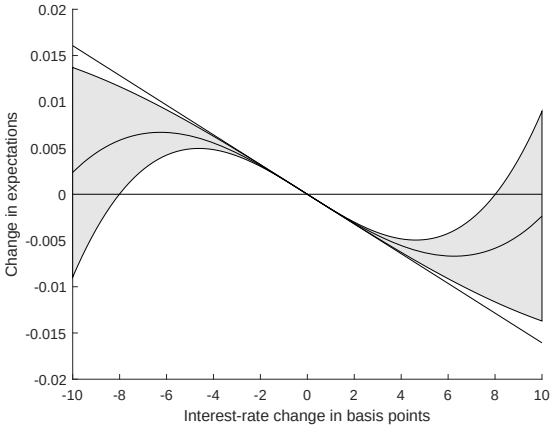

(b) Production expectations

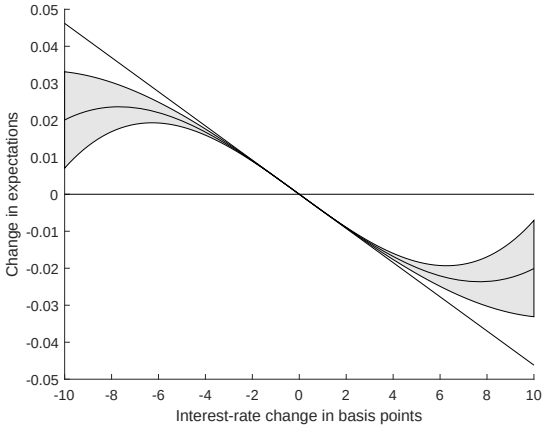

Notes: Graphical representation of results from the regression of changes in expectations on monetary policy surprises, surprises cubed, and the central bank information shock provided by Jarociński and Karadi (2020) in four-working-day windows around monetary events (see Table 1.6). Straight line represents estimate of linear term for changes in the 1-month OIS rate. Shaded area indicates 90% confidence interval around the cubic component. Firm expectations are obtained from IBS. Monetary policy surprises are measured using high-frequency interest rate changes from Altavilla et al. (2019).

Figure A.2: Effect of changes in the 1-month OIS rate, press release vs. press conference window

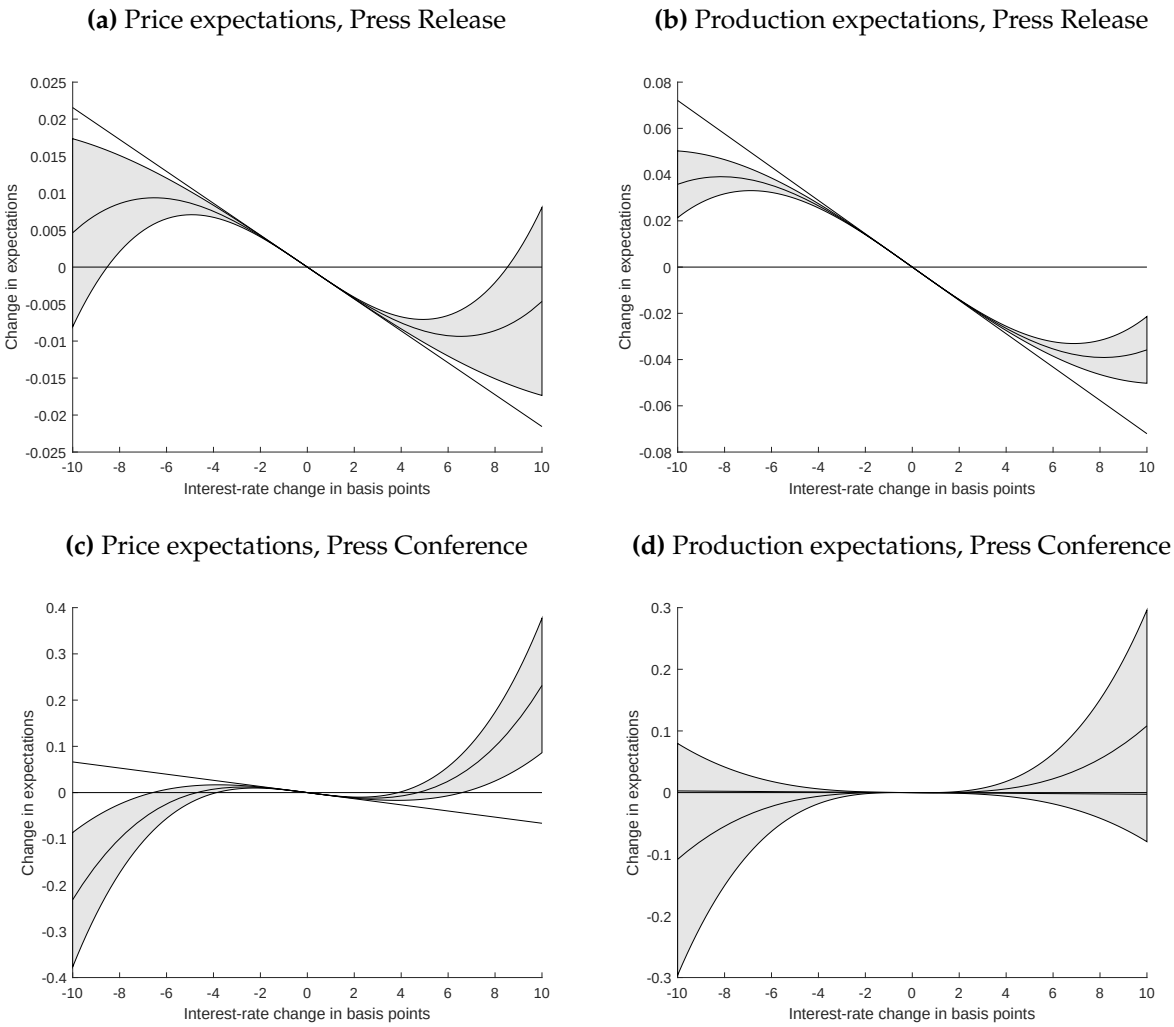

Notes: Graphical representation of results from the regression of changes in expectations on monetary policy surprises and surprises cubed in four-working-day windows around the events, separately for the press release and the press conference window (see Table 1.7). Straight line represents estimate of linear term. Shaded area indicates 90% confidence interval around the cubic component. Firm expectations are obtained from IBS. Monetary policy surprises are measured using high-frequency interest rate changes from Altavilla et al. (2019).

Table A.1: All questions from the IBS used in the estimations

Label	Name	Question	Possible answers
Q1	expected prices	Expectations for the next 3 months: Taking changes of terms and conditions into account, our domestic sales prices (net) for product XY will probably ...	increase [1] not change [0] decrease [-1]
Q2	expected production	Expectations for the next 3 months: Our domestic production activity regarding good XY will probably ...	increase [1] not change [0] decrease [-1]
Q3	prices	Tendencies in the previous month: Taking changes of terms and conditions into account, our domestic sales prices (net) for product XY have ...	increased [1] not changed [0] decreased [-1]
Q4	production	Tendencies in the previous month: Our domestic production activities with respect to product XY have ...	increased [1] not changed [0] decreased [-1]
Q5	demand	Tendencies in the previous month: The demand situation with respect to product XY is ...	better [1] not changed [0] worse [-1]
Q6	orders	We consider our order backlog to be ...	relatively high [1] sufficient [0] too small [-1]
Q7	foreign orders	We consider our order backlog for exports to be ...	relatively high [1] sufficient [0] too small [-1]
Q8	capacity utilization	The current utilization of our capacities for producing XY (standard utilization = 100%) is currently x%.	x is a value between 30 and 100 divisible by 10 OR if value $>$ 100, firms can write this value down
Q9	state of business (ordinal)	Current situation: We evaluate our state of business for XY to be...	good [1] satisfactory [0] bad [-1]
Q10	expected state of business (ordinal)	Expectations for the next 6 months: Our state of business for XY will ...	improve [1] stay the same [0] worsen [-1]
Q11	expected state of business (scale)	Expectations for the next 6 months: Our state of business for XY will x ...	x is a scalar between 0 and 100 chosen by moving a slider; the following values are labeled: 0–worsen, 50–stay the same, 100–improve
Q12	inventories	Current situation: we assess our stock of unsold amounts of good XY to be ...	too low [1] sufficient (for the season) [0] too large [-1]

Notes: Authors' translation of the most recent formulation of the question in German according to the EBDC Questionnaire manual. We only show those answer possibilities that we consider. Specifically, we exclude "no production" or similar answers which indicate that the question does not apply to the firm.

Table A.2: Descriptive statistics

	Full sample			Sample with part. date		
	Mean	Std. Dev.	Obs.	Mean	Std. Dev.	Obs.
Employees in production	489	3563.26	428790	548	3771.49	232267
Expected production, *t*	0.04	0.57	414486	0.06	0.58	224473
Expected prices, *t*	0.08	0.47	426451	0.08	0.47	231031
Production, *t-1*	-0.00	0.58	413784	0.01	0.58	224232
Prices, *t-1*	0.03	0.44	426706	0.04	0.43	231021
Demand, *t-1*	0.02	0.65	428220	0.03	0.66	231851
Orders, *t*	-0.14	0.65	426175	-0.12	0.66	231498
Foreign orders, *t*	-0.16	0.58	422043	-0.14	0.60	229778
Capacity utilization, *t*	81.08	16.57	366987	81.63	16.19	208385
State of business (ordinal), *t*	0.12	0.68	428291	0.15	0.69	231959
Exp. state of business (ordinal), *t*	0.02	0.60	427022	0.02	0.60	231297
Exp. state of business (scale), *t*	52.47	16.46	243925	52.64	16.44	213926
Inventories, *t*	-0.11	0.48	294251	-0.09	0.48	159477

Notes: Comparison of mean and variance for all variables we consider between the full sample and our sub-sample. Data from the IBS.

Table A.3: Important ECB announcements and SPF rounds

		Associated SPF survey		
Date	Announcement	Round	Start	End
05/07/2009	12-month Longer-term Refinancing Operations (LTROs) + other measures	2009Q3	07/15/2009	07/17/2009
08/04/2011	6-month LTROs + other measures	2011Q4	10/14/2011	10/18/2011
10/06/2011	12 and 13-month LTROs	2011Q4	10/14/2011	10/18/2011
12/08/2011	36-month LTROs	2012Q1	01/17/2012	01/20/2012
08/02/2012	Announcement of the Outright Monetary Transactions (OMT) program	2012Q4	10/16/2012	10/22/2012
09/06/2012	OMT implementation details	2012Q4	10/16/2012	10/22/2012
07/04/2013	First forward guidance announcement	2013Q3	07/16/2013	07/19/2013
06/05/2014	Targeted Longer-term Refinancing Operations (TLTROs)	2014Q3	07/17/2014	07/24/2014
09/04/2014	Announcement of the Asset-backed Securities Purchase Program (ABSPP) and the new Covered Bonds Purchase Program (CBPP3)	2014Q4	10/16/2014	10/23/2014
01/22/2015	Announcement of the expanded Asset Purchase Program (APP)	2015Q2	03/31/2015	04/07/2015
03/05/2015	APP implementation details	2015Q2	03/31/2015	04/07/2015
09/03/2015	Increase in public sector purchase program (PSPP) share limit	2015Q4	09/30/2015	10/06/2015
03/10/2016	Announcement of Corporate Sector Purchase Program (CSPP)	2016Q2	03/31/2016	04/06/2016
12/08/2016	First extension of the APP	2017Q1	01/04/2017	01/10/2017
10/26/2017	Second extension of the APP	2018Q1	01/08/2018	01/11/2018
06/14/2018	Announcement of the end of the APP	2018Q3	07/02/2018	07/06/2018

Notes: Dates are an extended version of the list provided by Dedola et al. (2018). We also show which round of the survey of professional forecasters (SPF) conducted by the ECB is associated with the events. This is needed for the robustness check using the SPF expectations, shown in Table A.4.

Table A.4: Effect of unconventional monetary policy on SPF expectations

	Current year		Next year	
	(1) GDP growth	(2) HICP inflation	(3) GDP growth	(4) HICP inflation
12-month SLTROs	-1.4795***	-0.5590***	-0.4456***	-0.2903**
	(0.12)	(0.10)	(0.10)	(0.12)
12/13-month SLTROs	0.0005	0.1564***	-0.6040***	0.0589
	(0.05)	(0.05)	(0.09)	(0.05)
36-month SLTROs	0.2275**	0.1927**	-0.0812	-0.0195
	(0.10)	(0.08)	(0.09)	(0.08)
OMT	-0.1390***	0.0950*	-0.4834***	0.1906***
	(0.04)	(0.05)	(0.07)	(0.04)
Forward Guidance	0.0868	0.0051	-0.0054	-0.2196***
	(0.10)	(0.07)	(0.09)	(0.07)
ABSPP+CBPP3	-0.1606***	-0.0061	-0.3764***	-0.2639***
	(0.04)	(0.05)	(0.05)	(0.05)
APP announcement+details	0.5741***	-0.2066**	0.3985***	-0.0956
	(0.06)	(0.10)	(0.06)	(0.06)
PSPP share limit	-0.0246	0.2727***	-0.2805***	-0.2714***
	(0.07)	(0.05)	(0.07)	(0.06)
CSPP	0.0708	-0.4221***	0.0271	-0.1864***
	(0.05)	(0.06)	(0.03)	(0.06)
First APP extension	-0.0256	0.1893***	-0.0960*	0.0272
	(0.04)	(0.06)	(0.05)	(0.04)
Second APP extension	0.2579***	-0.0494	0.2859***	0.0624
	(0.05)	(0.06)	(0.07)	(0.05)
APP end	-0.3658***	-0.0622*	0.1254***	0.0810**
	(0.03)	(0.04)	(0.04)	(0.03)
Previous SPF forecast for GDP growth in *t*	-0.0154*			
	(0.01)			
Previous SPF forecast for HICP inflation in *t*		-0.0188		
		(0.01)		
Previous SPF forecast for GDP growth in *t+1*			-0.1730***	
			(0.03)	
Previous SPF forecast for HICP inflation in *t+1*				-0.2300***
				(0.04)
Revision of expected oil price, average next 4 quarters (SPF)	0.0028***	0.0142***	0.0024**	0.0043***
	(0.00)	(0.00)	(0.00)	(0.00)
Revision of expected USD/EUR exchange rate, average next 4 quarters (SPF)	0.4114**	-0.2342	0.5644***	-0.1602
	(0.18)	(0.17)	(0.15)	(0.11)
Revision of expected main refinancing rate, average next 4 quarters (SPF)	0.1195*	0.1599***	0.0702	0.1196*
	(0.07)	(0.05)	(0.05)	(0.07)
Survey in first quarter	-0.2136***	0.0649*	-0.1855***	-0.0464
	(0.03)	(0.03)	(0.04)	(0.03)
Survey in second quarter	-0.1168***	0.1629***	-0.0995***	0.0093
	(0.03)	(0.02)	(0.03)	(0.03)
Survey in third quarter	-0.0272	0.1220***	-0.1663***	0.0011
	(0.02)	(0.02)	(0.03)	(0.03)
Constant	0.0337	-0.0352	0.2928***	0.3528***
	(0.02)	(0.02)	(0.05)	(0.07)
Observations	1217	1211	1170	1166
Within R^2	0.77	0.58	0.49	0.26
No. forecasters	50	50	49	49
Avg. obs/forecaster	24.3	24.2	23.9	23.8
Min. obs/forecaster	10	10	10	10
Incl. Reuters surprises	Yes	Yes	Yes	Yes

Notes: Results from regression of changes in the forecasts in the Survey of Professional Forecasters (SPF) conducted by the ECB on dummy variables indicating different monetary policy announcements. Forecasts are made for the current year and the next year. HICP inflation is measured as the year-on-year change in the HICP price index. GDP growth is measured as the annual real GDP growth rate. The timing of the announcements and the survey rounds is shown in Table A.3. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.5: Effect on realized prices and production

	Dependent variable:			
	Prices, *t*	Prices, *t+1*	Production, *t*	Production, *t+1*
	(1)	(2)	(3)	(4)
OIS, 1-month	-0.004228***	-0.006543***	-0.013134***	-0.012782***
	(0.00107)	(0.00122)	(0.00168)	(0.00184)
OIS, 1-month, cubic	0.000009*	0.000024***	0.000028***	0.000040***
	(0.00000)	(0.00001)	(0.00001)	(0.00001)
OIS, 1-month × after event	0.001812	0.002651	0.004510*	0.002126
	(0.00151)	(0.00169)	(0.00234)	(0.00250)
OIS, 1-month, cubic × after event	-0.000000	-0.000002	0.000005	0.000018
	(0.00001)	(0.00001)	(0.00001)	(0.00001)
Further Controls	X	X	X	X
Observations	55666	55199	54241	53820
Adjusted R^2	0.35	0.24	0.27	0.18
Observations before	27252	27004	26516	26304
Observations after	28414	28195	27725	27516

Notes: Results for regression of realized changes in prices and production on monetary policy surprises and surprises cubed for firms which responded in four-working-day windows around monetary events. Interactions with a dummy indicating whether response was recorded after an event are included. Firm responses are obtained from IBS. Monetary policy surprises are measured using high-frequency interest rate changes from Altavilla et al. (2019). Full set of additional control variables included but not shown. Robust standard errors in parentheses. * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Table A.6: Robustness checks, different dependent variables

	Dependent variable: change in			
	state of business	inventories	exp. state of business, ordinal measure	exp. state of business, scale measure
	(1)	(2)	(3)	(4)
OIS, 1-month	0.000358	-0.001707	-0.004467***	-0.080802**
	(0.00120)	(0.00139)	(0.00155)	(0.03264)
OIS, 1-month, cubic	-0.000002	0.000016*	0.000040***	0.000404**
	(0.00001)	(0.00001)	(0.00001)	(0.00019)
Further Controls	X	X	X	X
Observations	55327	36617	56989	52905
Adjusted R^2	0.32	0.52	0.31	0.17
Observations before	27102	17907	27916	27139
Observations after	28225	18710	29073	25766

Notes: Results for regression of changes in different dependent variables on monetary policy surprises and surprises cubed in four-working-day windows around the respective events. Answers in column (4) are measured on a scale from 0-100. Firm responses are obtained from IBS. Monetary policy surprises are measured using high-frequency interest rate changes from Altavilla et al. (2019). Full set of additional control variables included but not shown. Robust standard errors in parentheses. * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Table A.7: Effect of Jarociński-Karadi monetary policy shocks

	Dependent variable: change in the expectations for			
	Prices		Production	
	(1) Baseline	(2) Further controls	(3) Baseline	(4) Further controls
Monetary policy shock	0.0003	0.0010	-0.0003	-0.0015
	(0.002)	(0.002)	(0.003)	(0.002)
Central bank information shock	0.0068***	0.0061***	0.0072***	-0.0014
	(0.002)	(0.002)	(0.002)	(0.002)
Expected prices, *t-1*	-0.4531***	-0.5801***		
	(0.004)	(0.005)		
Expected production, *t-1*			-0.4936***	-0.6207***
			(0.004)	(0.005)
Average state of business, *t-1*	0.1447***	0.0838***	0.1449***	0.0936***
	(0.008)	(0.008)	(0.011)	(0.011)
Prices, *t-1*		0.2573***		
		(0.006)		
Prices, *t-2*		0.0159***		
		(0.006)		
Production, *t-1*				0.0868***
				(0.006)
Production, *t-2*				0.0097*
				(0.005)
Demand, *t-1*		0.0507***		0.2198***
		(0.003)		(0.004)
Demand, *t-2*		-0.0046*		-0.0129***
		(0.003)		(0.004)
Orders, *t-1*		0.0117***		0.0084*
		(0.004)		(0.005)
Foreign orders, *t-1*		-0.0021		0.0276***
		(0.004)		(0.005)
Capacity, *t-1*		-0.0001		-0.0006***
		(0.000)		(0.000)
Constant	0.0176***	0.0361***	-0.0017	0.0654***
	(0.002)	(0.009)	(0.002)	(0.012)
Observations	69121	62641	66931	61150
Adjusted R^2	0.23	0.29	0.24	0.33
Observations before	32571	29465	31535	28767
Observations after	36550	33176	35396	32383

Notes: Results based on regression of changes in expectations on monetary policy shocks provided by Jarociński and Karadi (2020) in four-working-day windows around monetary events. Monetary policy shocks are standardized. Firm expectations are obtained from IBS. Robust standard errors in parentheses. * p<0.1, ** p<0.5, *** p<0.01.

A2 Detailed versions of the main tables

Table A.8: Effect of unconventional monetary policy on firm expectations, detail

	Dependent variable: change in the expectations for prices			production		
	(1) Baseline	(2) Further controls	(3) Time FE	(4) Baseline	(5) Further controls	(6) Time FE
12-month LTROs	-0.1558***	-0.1010***	-0.0052	-0.1401***	-0.0661	-0.0561
	(0.032)	(0.031)	(0.038)	(0.041)	(0.041)	(0.051)
6-month LTROs	-0.0361	-0.0336	-0.0431	-0.0462	-0.0150	-0.0249
	(0.027)	(0.026)	(0.031)	(0.036)	(0.035)	(0.041)
12/13-month LTROs	-0.0292	-0.0639**	-0.0409	-0.1362***	-0.1528***	-0.0798*
	(0.026)	(0.025)	(0.028)	(0.038)	(0.040)	(0.044)
36-month LTROs	0.0699**	0.0859**	0.0562	-0.0027	0.0268	0.0696
	(0.035)	(0.035)	(0.046)	(0.042)	(0.040)	(0.056)
OMT details	-0.0537**	-0.0379	-0.0344	-0.1921***	-0.1345***	-0.1226***
	(0.026)	(0.026)	(0.029)	(0.039)	(0.040)	(0.044)
Forward Guidance	-0.0298**	-0.0187		-0.0047	0.0011	
	(0.013)	(0.012)		(0.019)	(0.018)	
TLTROs	-0.0702	-0.0552	-0.0227	-0.0423	0.0098	0.0482
	(0.052)	(0.052)	(0.056)	(0.067)	(0.069)	(0.074)
ABSPP+CBPP3	-0.0107	-0.0062		-0.0364*	0.0075	
	(0.013)	(0.013)		(0.021)	(0.021)	
APP details	0.0058	-0.0031		0.0279	0.0303	
	(0.020)	(0.020)		(0.026)	(0.027)	
PSPP share limit	-0.0267	-0.0190		0.0641**	0.1010***	
	(0.017)	(0.017)		(0.031)	(0.033)	
APP end	0.0337	0.0279	-0.0060	-0.0134	-0.0106	-0.0552
	(0.028)	(0.033)	(0.048)	(0.043)	(0.045)	(0.067)
Expected prices, *t-1*	-0.4389***	-0.5763***	-0.5818***			
	(0.002)	(0.003)	(0.003)			
Expected production, *t-1*				-0.4930***	-0.6072***	-0.6120***
				(0.002)	(0.003)	(0.003)
Prices, *t-1*		0.2664***	0.2617***			
		(0.004)	(0.004)			
Prices, *t-2*		0.0028	0.0020			
		(0.003)	(0.003)			
Production, *t-1*					0.0661***	0.0681***
					(0.003)	(0.003)
Production, *t-2*					0.0114***	0.0108***
					(0.003)	(0.003)
Demand, *t-1*		0.0490***	0.0455***		0.2021***	0.1980***
		(0.001)	(0.001)		(0.002)	(0.002)
Demand, *t-2*		-0.0026*	-0.0058***		-0.0096***	-0.0131***
		(0.001)	(0.002)		(0.002)	(0.002)
Orders, *t-1*		0.0088***	0.0066***		0.0079***	0.0044
		(0.002)	(0.002)		(0.003)	(0.003)
Foreign orders, *t-1*		0.0013	-0.0012		0.0244***	0.0212***
		(0.002)	(0.002)		(0.003)	(0.003)
Capacity utilization, *t-1*		0.0000	-0.0000		-0.0004***	-0.0004***
		(0.000)	(0.000)		(0.000)	(0.000)
Constant	0.0293***	0.0319***	-0.0270**	0.0228***	0.0561***	-0.0085
	(0.001)	(0.005)	(0.011)	(0.001)	(0.006)	(0.014)
Observations	236635	201212	201212	230028	197239	197239
Adjusted R^2	0.22	0.29	0.29	0.25	0.32	0.33
Monthly time FE	No	No	Yes	No	No	Yes

Notes: Detailed version of Table 1.1 in the main body, showing all control variables. Robust standard errors displayed in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table A.9: Effect of changes in the 1-month OIS rate, detail

	Dependent variable: change in the expectations for prices		Dependent variable: change in the expectations for production	
	(1) Baseline	(2) Further controls	(3) Baseline	(4) Further controls
OIS, 1-month	0.000712	0.001180*	0.001737*	0.000109
	(0.00070)	(0.00070)	(0.00093)	(0.00092)
Expected prices, *t-1*	-0.453926***	-0.577462***		
	(0.00442)	(0.00556)		
Expected production, *t-1*			-0.494913***	-0.621993***
			(0.00401)	(0.00493)
Average state of business, *t-1*	0.133941***	0.077335***	0.132691***	0.091695***
	(0.00847)	(0.00873)	(0.01141)	(0.01159)
Prices, *t-1*		0.255370***		
		(0.00656)		
Prices, *t-2*		0.012252**		
		(0.00571)		
Production, *t-1*				0.086247***
				(0.00588)
Production, *t-2*				0.010418*
				(0.00543)
Demand, *t-1*		0.049109***		0.220692***
		(0.00270)		(0.00455)
Demand, *t-2*		-0.004271		-0.013500***
		(0.00281)		(0.00447)
Orders, *t-1*		0.011549***		0.008696*
		(0.00384)		(0.00512)
Foreign orders, *t-1*		-0.003217		0.027333***
		(0.00385)		(0.00496)
Capacity utilization, *t-1*		-0.000096		-0.000619***
		(0.00011)		(0.00014)
Constant	0.019802***	0.036120***	0.004984*	0.066913***
	(0.00189)	(0.00905)	(0.00256)	(0.01227)
Observations	65003	58779	62968	57379
Adjusted R^2	0.23	0.28	0.24	0.33
Observations before	31978	28761	30960	28058
Observations after	33025	30018	32008	29321

Notes: Detailed version of Table 1.2 in the main body, showing all control variables and more digits for the coefficients. Robust standard errors displayed in parentheses. * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Table A.10: Effect of changes in the 1-month OIS rate, with cubic changes, detail

	Dependent variable: change in the expectations for prices			production		
	(1) Baseline	(2) Further controls	(3) Ex. largest OIS changes	(4) Baseline	(5) Further controls	(6) Ex. largest OIS changes
OIS, 1-month	-0.000516	-0.000683	-0.003538**	-0.001756	-0.003842**	-0.003887*
	(0.00116)	(0.00119)	(0.00171)	(0.00154)	(0.00152)	(0.00223)
OIS, 1-month, cubic	0.000007	0.000011*	0.000102***	0.000021***	0.000024***	0.000100***
	(0.00001)	(0.00001)	(0.00003)	(0.00001)	(0.00001)	(0.00004)
Expected prices, *t-1*	-0.453915***	-0.577454***	-0.575872***			
	(0.00442)	(0.00556)	(0.00565)			
Expected production, *t-1*				-0.495059***	-0.622135***	-0.622122***
				(0.00401)	(0.00493)	(0.00502)
Average state of business, *t-1*	0.134514***	0.078224***	0.082531***	0.134375***	0.093574***	0.094782***
	(0.00849)	(0.00875)	(0.00889)	(0.01143)	(0.01161)	(0.01177)
Prices, *t-1*		0.255362***	0.252039***			
		(0.00656)	(0.00668)			
Prices, *t-2*		0.012275**	0.012791**			
		(0.00571)	(0.00582)			
Production, *t-1*					0.086450***	0.084561***
					(0.00588)	(0.00598)
Production, *t-2*					0.010438*	0.011266**
					(0.00543)	(0.00552)
Demand, *t-1*		0.049103***	0.048136***		0.220601***	0.216653***
		(0.00270)	(0.00276)		(0.00455)	(0.00464)
Demand, *t-2*		-0.004357	-0.005429*		-0.013671***	-0.015195***
		(0.00281)	(0.00287)		(0.00447)	(0.00456)
Orders, *t-1*		0.011548***	0.012220***		0.008706*	0.011403**
		(0.00384)	(0.00391)		(0.00512)	(0.00521)
Foreign orders, *t-1*		-0.003237	-0.002813		0.027275***	0.026387***
		(0.00385)	(0.00392)		(0.00496)	(0.00505)
Capacity utilization, *t-1*		-0.000094	-0.000065		-0.000615***	-0.000621***
		(0.00011)	(0.00011)		(0.00014)	(0.00015)
Constant	0.019678***	0.035749***	0.034343***	0.004634*	0.066216***	0.072745***
	(0.00190)	(0.00905)	(0.00918)	(0.00256)	(0.01227)	(0.01247)
Observations	65003	58779	56491	62968	57379	55155
Adjusted R^2	0.23	0.28	0.28	0.24	0.33	0.33
Observations before	31978	28761	27395	30960	28058	26731
Observations after	33025	30018	29096	32008	29321	28424
Excl. largest surprises			X			X

Notes: Detailed version of Table 1.3 in the main body, showing all control variables and more digits for the coefficients. Robust standard errors displayed in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table A.11: Effect of changes in the 1-month OIS rate, binarized dependent variable, detail

	Dependent variable: change in the expectations for prices		Dependent variable: change in the expectations for production	
	(1) Upward revision	(2) Downward revision	(3) Upward revision	(4) Downward revision
OIS, 1-month	0.000139	0.000796	-0.001248	0.002247**
	(0.00078)	(0.00086)	(0.00097)	(0.00105)
OIS, 1-month, cubic	0.000002	-0.000010**	0.000007	-0.000016***
	(0.00000)	(0.00000)	(0.00000)	(0.00001)
Expected prices, *t-1*	-0.233636***	0.320257***		
	(0.00352)	(0.00425)		
Expected production, *t-1*			-0.262166***	0.305374***
			(0.00295)	(0.00325)
Average state of business, *t-1*	0.034385***	-0.036316***	0.013427*	-0.075074***
	(0.00623)	(0.00620)	(0.00786)	(0.00758)
Prices, *t-1*	0.129232***	-0.109666***		
	(0.00420)	(0.00486)		
Prices, *t-2*	-0.014247***	-0.033737***		
	(0.00377)	(0.00416)		
Production, *t-1*			0.037697***	-0.045633***
			(0.00377)	(0.00367)
Production, *t-2*			0.003718	-0.001646
			(0.00344)	(0.00360)
Demand, *t-1*	0.021587***	-0.026210***	0.095279***	-0.106071***
	(0.00192)	(0.00187)	(0.00292)	(0.00296)
Demand, *t-2*	-0.005468***	-0.001109	-0.010343***	0.000381
	(0.00201)	(0.00194)	(0.00291)	(0.00296)
Orders, *t-1*	0.006674**	-0.003678	-0.001640	-0.010240***
	(0.00272)	(0.00268)	(0.00329)	(0.00342)
Foreign orders, *t-1*	-0.005283*	-0.003187	0.009878***	-0.015610***
	(0.00274)	(0.00267)	(0.00323)	(0.00328)
Capacity utilization, *t-1*	-0.000185**	-0.000106	-0.000672***	-0.000046
	(0.00008)	(0.00007)	(0.00010)	(0.00009)
Constant	0.117578***	0.083770***	0.208047***	0.146019***
	(0.00656)	(0.00627)	(0.00814)	(0.00787)
Observations	58779	58779	57379	57379
Adjusted R^2	0.12	0.20	0.17	0.21
Observations before	28761	28761	28058	28058
Observations after	30018	30018	29321	29321

Notes: Detailed version of Table 1.4 in the main body, showing all control variables and more digits for the coefficients. Robust standard errors displayed in parentheses. * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Table A.12: Effect of 1-month OIS changes and Jarociński-Karadi shocks, detail

	Dependent variable: change in the expectations for prices		production	
	(1) CBI	(2) CBI+MP	(3) CBI	(4) CBI+MP
OIS, 1-month	-0.001607	-0.002871*	-0.004620***	-0.002371
	(0.00130)	(0.00164)	(0.00166)	(0.00207)
OIS, 1-month, cubic	0.000014**	0.000016**	0.000026***	0.000023***
	(0.00001)	(0.00001)	(0.00001)	(0.00001)
Central bank information shock (CBI)	0.004250*	0.005834**	0.003364	0.000541
	(0.00229)	(0.00267)	(0.00294)	(0.00328)
Monetary policy shock (MP)		0.003620		-0.006430*
		(0.00299)		(0.00362)
Expected prices, *t-1*	-0.577985***	-0.578014***		
	(0.00571)	(0.00571)		
Expected production, *t-1*			-0.622090***	-0.622083***
			(0.00505)	(0.00505)
Average state of business, *t-1*	0.078967***	0.078520***	0.091358***	0.092153***
	(0.00915)	(0.00916)	(0.01209)	(0.01209)
Prices, *t-1*	0.254644***	0.254637***		
	(0.00673)	(0.00673)		
Prices, *t-2*	0.014408**	0.014392**		
	(0.00588)	(0.00588)		
Production, *t-1*			0.087733***	0.087714***
			(0.00601)	(0.00601)
Production, *t-2*			0.011033**	0.011058**
			(0.00556)	(0.00556)
Demand, *t-1*	0.050929***	0.050972***	0.219238***	0.219172***
	(0.00277)	(0.00277)	(0.00465)	(0.00465)
Demand, *t-2*	-0.005250*	-0.005253*	-0.014850***	-0.014851***
	(0.00288)	(0.00288)	(0.00457)	(0.00458)
Orders, *t-1*	0.011596***	0.011589***	0.007719	0.007738
	(0.00396)	(0.00396)	(0.00526)	(0.00526)
Foreign orders, *t-1*	-0.002459	-0.002447	0.029101***	0.029080***
	(0.00397)	(0.00397)	(0.00510)	(0.00510)
Capacity utilization, *t-1*	-0.000130	-0.000130	-0.000596***	-0.000595***
	(0.00011)	(0.00011)	(0.00015)	(0.00015)
Constant	0.039176***	0.039209***	0.064532***	0.064498***
	(0.00926)	(0.00926)	(0.01254)	(0.01254)
Observations	56109	56109	54754	54754
Adjusted R^2	0.28	0.28	0.33	0.33
Observations before	26706	26706	26046	26046
Observations after	29403	29403	28708	28708

Notes: Detailed version of Table 1.6 in the main body, showing all control variables and more digits for the coefficients. Robust standard errors displayed in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table A.13: Effect of changes in the 1-month OIS rate, press release and press conference window separately, detail

	Dependent variable: change in the expectations for prices		production	
	(1) Release	(2) Conference	(3) Release	(4) Conference
OIS, 1-month	-0.002156	-0.006632**	-0.007207***	-0.000271
	(0.00154)	(0.00278)	(0.00191)	(0.00350)
OIS, 1-month, cubic	0.000017**	0.000298***	0.000036***	0.000111
	(0.00001)	(0.00009)	(0.00001)	(0.00011)
Expected prices, *t-1*	-0.577428***	-0.577458***		
	(0.00556)	(0.00556)		
Expected production, *t-1*			-0.622094***	-0.621948***
			(0.00493)	(0.00493)
Average state of business, *t-1*	0.077271***	0.082035***	0.089637***	0.092029***
	(0.00874)	(0.00896)	(0.01161)	(0.01181)
Prices, *t-1*	0.255341***	0.255253***		
	(0.00656)	(0.00656)		
Prices, *t-2*	0.012268**	0.012191**		
	(0.00571)	(0.00571)		
Production, *t-1*			0.086461***	0.086208***
			(0.00588)	(0.00588)
Production, *t-2*			0.010366*	0.010427*
			(0.00543)	(0.00543)
Demand, *t-1*	0.049070***	0.049115***	0.220432***	0.220665***
	(0.00270)	(0.00270)	(0.00455)	(0.00455)
Demand, *t-2*	-0.004426	-0.004175	-0.013855***	-0.013532***
	(0.00281)	(0.00281)	(0.00447)	(0.00447)
Orders, *t-1*	0.011570***	0.011543***	0.008783*	0.008717*
	(0.00384)	(0.00384)	(0.00512)	(0.00512)
Foreign orders, *t-1*	-0.003247	-0.003259	0.027256***	0.027320***
	(0.00385)	(0.00385)	(0.00496)	(0.00496)
Capacity utilization, *t-1*	-0.000094	-0.000097	-0.000614***	-0.000619***
	(0.00011)	(0.00011)	(0.00014)	(0.00014)
Constant	0.035939***	0.035265***	0.067156***	0.067040***
	(0.00905)	(0.00906)	(0.01227)	(0.01227)
Observations	58779	58779	57379	57379
Adjusted R^2	0.28	0.28	0.33	0.33
Observations before	28761	28761	28058	28058
Observations after	30018	30018	29321	29321

Notes: Detailed version of Table 1.7 in the main body, showing all control variables and more digits for the coefficients. Robust standard errors displayed in parentheses. * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Chapter 2

Household Expectations, Monetary Policy and the Media*

*I thank Hassan Afrouzi, Rüdiger Bachmann, Gerhard Illing, Stefan Lautenbacher, Alexander Schramm, Alexander Schwemmer, and Peter Zorn for helpful comments. In addition, I thank Alexander Schramm for great discussions leading to the initial idea for this chapter, and Brendan Shanks for sharing his materials about working with R. Finally, I thank Omari Teka for excellent research assistance.

2.1 Introduction

In recent years, central banks have increasingly focused on managing firm and household expectations. This is mainly due to the constraints of the effective lower bound forcing central banks to rely on additional instruments (den Haan 2013). Given the central banks' focus on price stability, inflation expectations are particularly relevant. Inflation expectations are a key transmission channel from monetary policy to inflation, both directly and indirectly through the real interest rate. It is therefore important to understand how monetary policy affects inflation expectations of all actors in the economy. Given the scarce empirical evidence so far, this chapter focuses on the role of monetary policy for households' inflation expectations.

I show that US households' quantitative inflation expectations increase after a surprise tightening of monetary policy, i.e. a surprise increase in interest rates after a federal open market committee (FOMC) announcement. For the inflation expectations I use data from the Michigan Survey of Consumers. Monetary policy surprises are measured using the high-frequency identified changes in interest rates in a small window around FOMC announcements provided by Jarociński and Karadi (2020). The effect is robust to controlling for unobserved household heterogeneity and to using different measures of monetary policy. The results are not driven by outlier observations or special episodes of US monetary policy. They are also robust across different demographic groups. In addition, results do not change when including other economic news released in the same month, in which the monetary policy announcement occurred.

The positive response to surprise increases is in line with a Delphic interpretation of monetary policy announcements (Campbell et al. 2012) or the information effect of monetary policy (Nakamura and Steinsson 2018). These theories propose that agents in the economy focus on the implication of the monetary policy announcement for the current inflation outlook, not on the effect of the changed monetary policy stance on the future development of inflation. This view of monetary policy goes back to the argument by Romer and Romer (2000) that the central banks may be (perceived to be) better informed about the future state of the economy than households and other economic agents. Because agents are aware of this informational advantage, they look to central bank announcements not only for information about the decision but also for information about the current economic outlook.

In order to verify that indeed this information effect is at play here, I analyze newspaper reporting after FOMC meetings. By now, it is commonly understood that most households do not directly follow monetary policy announcements (Binder 2017b). However, in order to form expectations, households rely on various information sources

including media reporting (D'Acunto et al. 2019b). I therefore use text analysis tools to classify articles published in US newspapers after monetary policy announcements into articles referring to high or increasing inflation and articles referring to low or decreasing inflation. After tightening monetary policy surprises more articles referring to high or increasing inflation are published. Articles usually refer to high or increasing inflation as the explanation for the decision made. The newspaper reporting in turn influences inflation expectations. I conclude that newspaper reporting provides a transmission channel of the information effect from the central bank to households.

This chapter relates to different strands of the literature studying expectations, namely studies of expectation formation generally, of the interaction between expectations and news reporting as well as monetary policy, and studies of the effect of expectations on outcomes. In the following, I discuss the most important papers from these strands.

The study of expectations has been quickly expanding in recent years with a focus on using micro level data. In part, this is due to the central banks' increased focus on expectations management. Another reason for the increased interest are the missing disinflation puzzle after the financial crisis and the missing inflation puzzle during the recovery (Coibion and Gorodnichenko 2015b). While earlier work already emphasized that the full information, rational expectations hypothesis is not in line with the empirical evidence (see e.g., Carroll 2003), it focused on time series data. More recently, an emphasis is put on micro-level data and constructing moments from these data to inform models (Coibion and Gorodnichenko 2012, 2015a). Furthermore, the literature now goes beyond testing for rationality or specific models of deviation from rationality.

Among others, Binder (2017a) and Das, Kuhnen, and Nagel (2020) show that expectations depend on demographic factors and the socio-economic status of households. Using a Finnish data set, D'Acunto, Hoang, Paloviita, and Weber (2019a) show that the IQ of a person has significant effects on their expectations, beyond education and income. In addition, individuals' experiences matter. Malmendier and Nagel (2016) find that households extrapolate from their own experiences over their lifetime when forming expectations. Therefore younger people react more strongly to new information arriving (because their accumulated experiences are shorter), and older people who experienced the high inflation periods in the US, generally have higher inflation expectations. Recent research using a novel data set of scanner data by D'Acunto et al. (2019b) highlights the importance of prices of frequently purchased goods. Their work indicates that observed price changes of groceries are one major source of information for inflation expectations. Previous literature already showed that inflation expectations react strongly to aggregate changes in food and gas prices (see e.g., Coibion and Gorodnichenko 2015b). I will take the evidence on household heterogeneity into ac-

count by considering whether different households react differently to monetary policy announcements. In addition, I control for current developments in food and gas price inflation given the relevance of these particular prices for households.

Another strand of literature analyzes whether news reporting affects the expectation formation process. Most work on this is related to tests of the epidemiological model of inflation expectations introduced by Carroll (2003). While earlier research focused on whether there are any effects of news, more recent contributions look at the determinants of these effects. Ehrmann, Pfajfar, and Santoro (2017) show that more pessimistic households have a larger initial bias in inflation expectations but also react more strongly to news. The content of news also matters. Lamla and Lein (2014) and Pfajfar and Santoro (2013) show that households react more strongly to news with a positive tone.[1] Lamla and Maag (2012) find that households respond more strongly to news when it is about higher prices or inflation. Generally, hearing news about inflation reduces the bias in inflation expectations.[2] This chapter provides additional evidence on the role of news by classifying articles based on their information about inflation. Furthermore, I add a new dimension to the analysis by also considering how monetary policy affects news reporting, and how monetary policy is transmitted to expectations via news.

The relevance of analyzing survey expectations of course depends on the assumption that expectations reported in surveys are actually relevant for households' decisions. Due to a lack of data on the individual level and general skepticism towards surveys, this has been difficult to do empirically for a long time. Early work by Bachmann, Berg, and Sims (2015) found only a weak relationship between inflation expectations and the willingness to spend on durable goods in the Michigan Survey of Consumers. Using the new and more detailed New York Fed Survey of Consumer Expectations, Crump, Eusepi, Tambalotti, and Topa (2015) find higher inflation expectations lower expected consumption in the future, in line with the Euler equation. Ichiue and Nishiguchi (2015) show that Japanese households increase current spending and lower future spending in response to higher inflation expectations. Dräger and Nghiem (2020) find a similar effect in Germany. Combining survey data with administrative data in the Netherlands, Vellekoop and Wiederholt (2019) show that households with higher inflation expectations lower their net worth, i.e. save less, and are more likely to purchase a car, as a proxy for spending on durable goods.

However, these studies may still suffer from endogeneity issues due to omitted variables. Therefore, several authors have explored natural experiments. D'Acunto,

[1] The tone is determined by hand coding articles based on methodologies from media analysis.

[2] The bias is usually measured by computing the difference between households' inflation expectations and some benchmark expectations series (such as professional forecasters) or realized inflation in the forecasting period.

Hoang, and Weber (2018) show that the announcement of a VAT increase in Germany led to an increase in inflation expectations and an increase in spending before the VAT increase took place. Carrillo and Emran (2012) exploit an error in reporting of inflation in Ecuador. Officially reported inflation rates were higher than actual inflation, which led households to increase inflation expectations and lower savings at the same time. Other authors have introduced exogenous variation into survey expectations. Coibion, Georgarakos, Gorodnichenko, and van Rooij (2019a) provide information treatments to households in a Dutch survey, which change the households' inflation expectations. In reaction to higher inflation expectations induced by the information treatment, households slightly increase their non-durable spending, but decrease their durable spending, so that overall spending falls. This reaction is due to households expecting their real income to fall when inflation increases, i.e. they associate higher inflation with worse economic outcomes. In a related study, Kamdar (2019) argues consumers' expectations can be explained by a single factor, which she labels sentiment. If households are generally optimistic about the future, they expect the economy to expand but inflation to decline. This result is in line with households expecting less real income when their inflation expectations increase. The recent evidence from survey experiments and natural experiments shows that expectations (also as reported in surveys) do matter for actions. Thus, the analysis of micro-level survey data can indeed provide insights into 'actual' behavior of agents in the economy.

Finally, regarding the empirical relationship between monetary policy and expectations, much of the literature so far has focused on the perception of central banks in general as well as the consistency of expectations with certain relationships deemed relevant for monetary policy. For a detailed discussion of the studies, see Binder (2017b). In addition, some authors have compared different communication channels and tools. Coibion, Gorodnichenko, and Weber (2019b) set up a randomized control trial within a survey of consumers where they provide different information treatments, ranging from information about actual inflation to newspaper articles about FOMC meetings and actual FOMC statements. The authors find households update inflation expectations most strongly in response to being provided with actual figures and with FOMC statements. Bholat, Broughton, Ter Meer, and Walczak (2019) show that simplifying text and relating information to people's experiences in central bank communication improves comprehension of the information and increases trust in the central bank.

The direct effects of monetary policy on household expectations are now being considered as well. Lamla and Vinogradov (2019) conduct an online survey within a few days around several FOMC meetings in 2016 and 2017 asking people about their inflation expectations, among other things. They find only weak effects of monetary policy

announcements on inflation expectations, largely related to increased news coverage around FOMC meetings. However, the study lacks a panel dimension so the effect may be due to differences between the households interviewed before and after the announcements. In contrast to this result, Lewis, Makridis, and Mertens (2019) find a significant and negative effect of tightening monetary policy surprises on expectations about the general economic situation. The authors also have access to daily expectations data but for a much longer time period than Lamla and Vinogradov (2019). In addition, they use high-frequency identified monetary policy surprises to quantify the announcement beyond the announced target rate change. Finally, Claus and Nguyen (2020) use aggregate Australian survey data and a latent factor model to identify the effects monetary policy. They find inflation expectations increase both after easing and tightening announcements.[3] The initial response depends on the income level: richer households respond immediately while poorer households respond only a few months after the announcements.

This chapter directly relates to the set of papers discussed in the last paragraph. Similar to Lewis et al. (2019), I study the effect of monetary policy using high-frequency changes in interest rates. However, I focus on inflation expectations instead of the economic outlook. While I lack the detailed information on the timing of the responses as available to Lamla and Vinogradov (2019) and Lewis et al. (2019), I cover a much larger sample of monetary policy announcements, namely all meetings since 1994. In addition, I can control for individual-specific effects using the panel dimension in my household data. I also explicitly control for the role of news by analyzing actual reporting after monetary policy announcements. For the analysis of news or other text data in this context, the papers I am aware of, have so far relied on human coding or a general method of tone analysis to distinguish different types of reporting. In this chapter, I develop a new approach to classify newspaper articles about inflation based on a dictionary method.

To summarize, this chapter is the first to consider the direct effects of quantitative monetary policy surprises on households' inflation expectations in a large micro-level sample with a long time dimension as well as a panel dimension. Additionally, it is the first to systematically link the effects of monetary policy and its transmission via news reporting for households.

The remainder of this chapter is structured as follows. The next section describes the data used. Section 2.3 describes the main results for the effect of monetary policy on expectations. Section 2.4 presents the text analysis approach and the results for the transmission of monetary policy to households via news reporting. Section 2.5 concludes.

[3]The authors distinguish announcements based on the announced target rate change.

2.2 Data

This study relies on three different types of data: expectations data from the Michigan Survey of Consumers (MSC), monetary policy surprises, and articles published by US newspapers. This section briefly describes the data sources and adjustments to the data.

2.2.1 Expectations data

The MSC is a regular survey of US consumers, which is representative for the US population,[4] and has been running since 1946. Since 1978, the survey is conducted at a monthly frequency. In each month, at least 500 households are interviewed via telephone.[5] In addition, about one half of the households is interviewed again after six months. In this study, I will exploit this limited panel dimension in order to control for idiosyncratic household effects. The analysis will therefore be limited to households responding to all questions needed for the estimation in both interviews.

The sample in the baseline case starts in January 1994. I choose this starting point due to the changes, which occurred in monetary policy communication over time. Since I am interested in the transmission via newspaper reporting, it makes sense to consider only the period since the FOMC regularly publishes press statements. Before 1994, this was not common practice and the release of a statement may have indicated that this meeting was different from a standard policy meeting (Binder 2017b). The communication practices since 1994 are also more in line with the way monetary policy is conducted nowadays. However, for robustness I also consider earlier periods as far as the measures of monetary policy allow. The baseline sample ends in December 2016 due to the availability of the monetary policy surprise series.

In total 85,146 different households responded to the survey between January 1994 and December 2016. 54,165 households in this period were interviewed twice and 30,981 were only interviewed once. This yields 139,311 observations in total and 108,333 observations with a panel dimension. On average, the number of observations per month is 505 in the full sample and 400 in the panel sample. The latter reflects that on average 200 respondents are being re-interviewed in each month.[6]

[4] Except for Hawaii and Alaska.

[5] Since the 1970s sampling was based on landline phones. Since 2014 the MSC has completely shifted to cell phones (Curtin and Dechaux 2015).

[6] I interchangeably use the terms household and respondent due to the mixed manner in which the MSC treats its respondents. The sampling is focused on representing US households and many question refer to household level variables. At the same time the expectations and perceptions questions of course focus on the individual responding to the survey. For details on the sampling see Curtin (1982) and more recently Curtin and Dechaux (2015).

The focus of this chapter are inflation expectations. In the MSC, inflation expectations are elicited in a two-step manner.[7] First, the respondents are asked about the direction of the expected change in prices:

Q1 During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?

If the respondents answer 'stay where they are', they will be asked whether they actually mean constant prices or whether they mean a constant inflation rate. If they confirm that they mean constant prices, their answer is recorded as 0. In case they mean a constant inflation rate, they are asked the same follow-up question as those respondents who are expecting a change, namely:

Q2 By about what percent do you expect prices to go *(up/down)* on the average, during the next 12 months?

Whether the questions mentions *up* or *down* depends on the answer in Q1. Answers are recorded as integer values. If the households provide an answer larger than 5% in either case, they are prompted to confirm this answer. If they say, that they do not know in response to Q2, they are asked again referring to 'cents on the dollar' instead of percent. Appendix B2.1 shows the detailed set-up of the questions.

From now on, I focus on the subset of households relevant for my analysis, namely those who are interviewed twice and report quantitative inflation expectations in both interviews: around 85% of all households with a panel dimension (45,874) in fact report inflation expectations twice. Table 2.1 shows summary statistics for these inflation expectations. I consider both the full sample as well as a breakdown by demographic groups. The average over the full sample is 3.4% and the median response is 3%.[8] These values are substantially higher than actual CPI inflation in the United States, which was around 2.2% during the sample period. Panel (a) of Figure 2.1 shows that the difference has grown larger since the financial crisis. This observation has been linked to energy prices and used to explain the missing disinflation puzzle (Coibion and Gorodnichenko 2015b). Panel (b) highlights the strong correlation between consumers' expected inflation and the level occurring in the following six months. Despite the level difference, inflation expectations and actual inflation co-move not only contemporaneously but also for leads of inflation. The strong contemporaneous correlation is in line with the fact that households extrapolate from observed price changes (D'Acunto et al. 2019b), while the correlation with future actual inflation indicates some predictive power of these expectations.

[7] The full questionnaire is available at `https://data.sca.isr.umich.edu/fetchdoc.php?docid=24776`.

[8] Remember the MSC only reports integer values.

Table 2.1: Inflation expectations, balanced sample, 1994 to 2016

	Observ.	Share	Mean	Std.Dev.	P10	P25	P50	P75	P90
All	91748		3.40	4.08	0	1	3	5	9
Gender									
Male	46014	0.50	3.05	3.59	0	1	3	5	7
Female	45734	0.50	3.74	4.50	0	1	3	5	10
Renter/owner									
Owns home	72846	0.80	3.35	3.93	0	1	3	5	8
Rents home	18684	0.20	3.55	4.63	0	1	3	5	10
Cohorts									
Born before 1970	75848	0.83	3.42	4.05	0	1	3	5	9
Born 1970 and later	15574	0.17	3.27	4.23	0	1	3	5	9
Age									
18-39 years	26047	0.28	3.29	4.29	0	1	3	5	8
40-59 years	37197	0.41	3.38	3.96	0	1	3	5	8
60 years and older	28179	0.31	3.52	4.04	0	1	3	5	10
Education									
High school or less	27063	0.30	3.85	4.67	0	1	3	5	10
Some college, no degree	24058	0.26	3.56	4.24	0	1	3	5	10
College degree	23770	0.26	3.05	3.59	0	1	3	5	6
Graduate studies	16667	0.18	2.90	3.32	0	1	3	4	5
Income									
Bottom 33%	19938	0.23	4.08	4.89	0	1	3	5	10
Middle 33%	30287	0.35	3.47	4.04	0	1	3	5	10
Top 33%	36456	0.42	2.93	3.48	0	1	3	4	6
Stock market participation									
No stocks	24696	0.37	3.76	4.34	0	1	3	5	10
Investment in bottom 33%	11673	0.18	3.49	4.00	0	1	3	5	10
Investment in middle 33%	13555	0.21	3.22	3.73	0	1	3	5	8
Investment in top 33%	16178	0.24	2.85	3.21	0	1	3	4	5

Notes: Table shows summary statistics for inflation for the balanced sample of households, i.e. those reporting quantitative inflation expectations (Q2) in both interviews. The first row shows all households in the balanced sample. The remaining rows show statistics for different demographic groups. The second column ('Share') shows how the demographic traits are distributed within the sample by providing the share in all observations for the respective subgroup. The distribution of demographic groups in the balanced sample differs slightly from the full sample, as shown in Table B.1 in Appendix B1.

In addition to the aggregate time series, I also consider the cross sectional heterogeneity of reported expectations. Table 2.1 shows that the interquartile range is 1% to 5%, and the bottom and top 10 percent responses are 0% and 9%, respectively. Furthermore, 83% of all answers fall within -2% and 5%, which is the range of actual inflation in the sample period. However, Figure B.1 in Appendix B1 shows that some house-

Figure 2.1: One-year ahead inflation expectations and actual inflation, 1994 to 2016

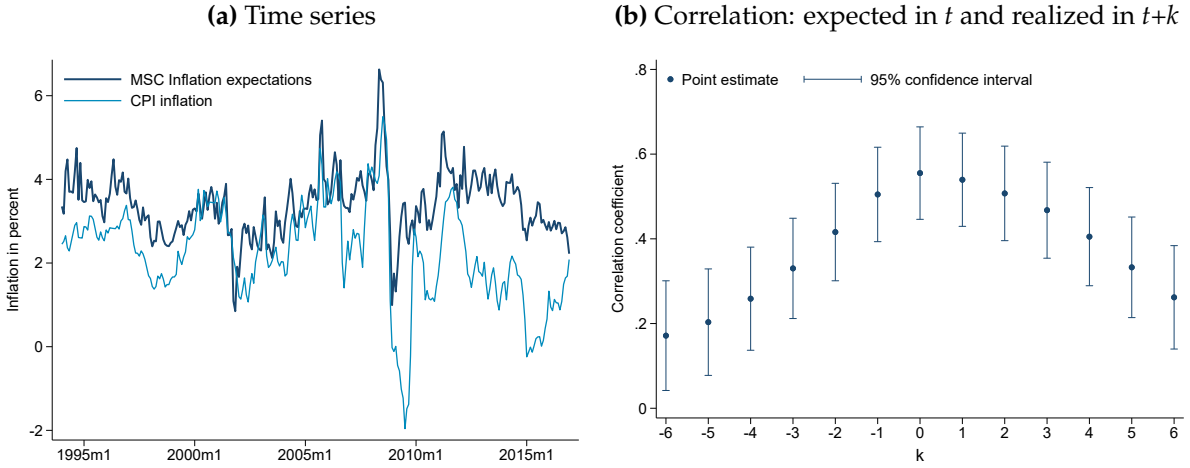

Notes: Expectations are measured as the mean of all answers in each month from the MSC (only households who provided expectations in both their interviews). Actual inflation is the seasonally adjusted year-on-year change in the CPI for All Urban Consumers taken from the FRED database. All data are monthly.

holds also report very large numbers, up to -50% and +50%. In addition, bunching at multiples of five is noticeable. Rounding to multiples of five or ten is considered to indicate uncertainty about the forecast (Binder 2017c). I will take both the rounding and the very large responses into account in the analysis later on.

The remainder of Table 2.1 shows some details on the demographic composition of the sample, and inflation expectations by the different groups.[9] The variables displayed are those used later on to analyze potential heterogeneity of the effects. I select them based on what previous literature showed to be relevant for expectation formation.[10]

The genders are fully balanced in the sample. In line with previous findings, women have higher inflation expectations than men. This difference has recently been linked to differences in the contribution to grocery shopping at the household level (D'Acunto, Malmendier, and Weber 2020). The vast majority of households (80%) owns a house. There is again a slight difference is expectations: renters have somewhat higher expectations than owners. Only 17% of all respondents were born after 1970, which means they did not experience the high inflation periods in the 1970s. In line with the results by Malmendier and Nagel (2016), they report lower expectations than households born earlier. All respondents in the MSC are at least 18 years old. They can be split into

[9]The questionnaire provided by the MSC does no contain the questions about the demographic variables, in particular not the one about income, only the questions about home and stock ownership are shown. Details on these can be found in Appendix B2.2.

[10]Note that the shares of the different groups are computed only for household reporting inflation expectations and responding to two interviews. This sample differs slightly from the full MSC sample, as shown in Table B.1 in Appendix B1.

roughly equal groups by choosing 40 years and 60 years as cut-off points. Inflation expectations increase with age. Education levels are recorded in the MSC at a detailed grade level. To achieve roughly equal sized groups for the analysis later on, I reclassify this data into four categories: high school degree or less, some college but no degree, undergraduate college degree, and graduate studies. For education, we see a strong decrease in inflation expectations with the length of education. The negative correlation is in line with results on the role of the socio-economic status for expectations by authors such as Das et al. (2020).

Finally, the MSC collects information on household income and stock market participation. For income, I consider the tercile split.[11] For stock market participation, I separate those who invest from those who do not. The former are then further split into terciles based on the amount they invest. In total only 63% of all respondents invest in the stock market. Also for income and stock ownership inflation expectations decrease with higher levels of income and investment, in line with the results for education.

2.2.2 Monetary policy measures

The main measure of monetary policy surprises I use are high-frequency changes in the 3-months fed funds futures provided by Jarociński and Karadi (2020). Their data is an updated version of the dataset by Gürkaynak et al. (2005). The changes are computed by taking the difference between the interest rate 10 minutes before the respective FOMC announcement and 30 minutes after.[12] Jarociński and Karadi (2020) use 3-months fed funds futures to measure monetary policy surprises because they capture both current changes in the policy stance as well as very short run forward guidance. In addition, they are not affected by 'timing surprises' if the Fed postpones the expected rate decision by one meeting (Jarociński and Karadi 2020).

The monetary surprises are available from January 1990 to December 2016. The surprises are generally small, with a mean of -0.91 basis points and a standard deviation of only 5 basis points. The majority of shocks is smaller than 10 basis points in absolute value. About one quarter of surprises is actually exactly zero. There are, however, also some very large surprises, as can be seen from the histograms in Figure 2.2. The histograms also show that over the sample period more negative than positive surprises occurred and that most very large surprises were negative.

Jarociński and Karadi (2020) also provide a split of the monetary policy surprises into a pure monetary policy shock and an information shock. This split is based on the

[11] The terciles are defined every month to ensure comparability over time.

[12] Between 1994 and 2016, the announcement is usually the press release. Before 1994 it is usually the first open market operations after the meeting which informed market participants about the decision taken (Jarociński and Karadi 2020).

Figure 2.2: Monetary policy surprises, 1994 to 2016

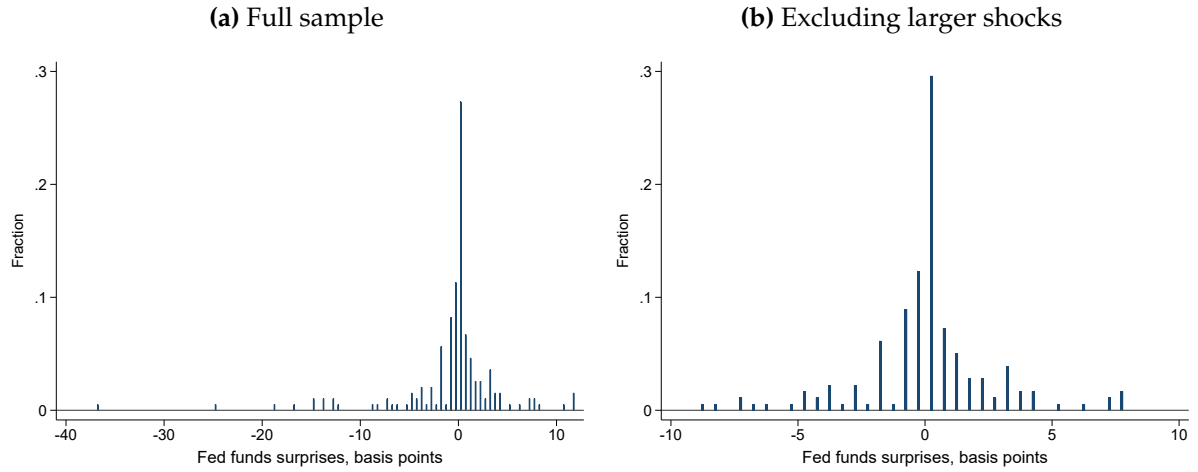

Notes: Histograms show all monetary policy surprises between January 1994 and December 2016 as provided by Jarociński and Karadi (2020). Monetary policy surprises are difference between the 3-months fed fund futures measured in a 10-minutes window before the respective FOMC announcement and a 30-minutes window after.

direction of change of the interest rate and the stock market in the tight window around the meeting. A pure monetary policy surprise moves interest rates and stock markets in opposite directions, and a central bank information shock moves both variables in the same direction.

The main part of my analysis focuses on the observed surprises in fed fund futures without considering the split provided by Jarociński and Karadi (2020). I choose to use the total surprises because households' may interpret the observed surprises differently than financial markets. However, as a robustness check I use the two separate shocks. The results confirm the finding of an information effect.

In addition, I use monetary policy surprises provided by Wieland and Yang (2020) which is an extension of the series by Romer and Romer (2004) for robustness checks. These surprises are derived by comparing the intended fed funds target from before a FOMC meeting to the one decided at the meeting. The targets are constructed using a narrative approach based on materials provided by the Fed. In a second step, the difference between the intended and the actual target is regressed on the Fed's forecasts of inflation, unemployment and real output growth. The residuals from this regression are estimates of monetary policy surprises purged of anticipatory actions by the Fed. Since the Fed releases its forecasts only with a long lag, private agents may still perceive theses surprises as containing information about the Fed's economic outlook.

Figure 2.3: Average number of articles per day with keywords around FOMC meetings, 1994 to 2016

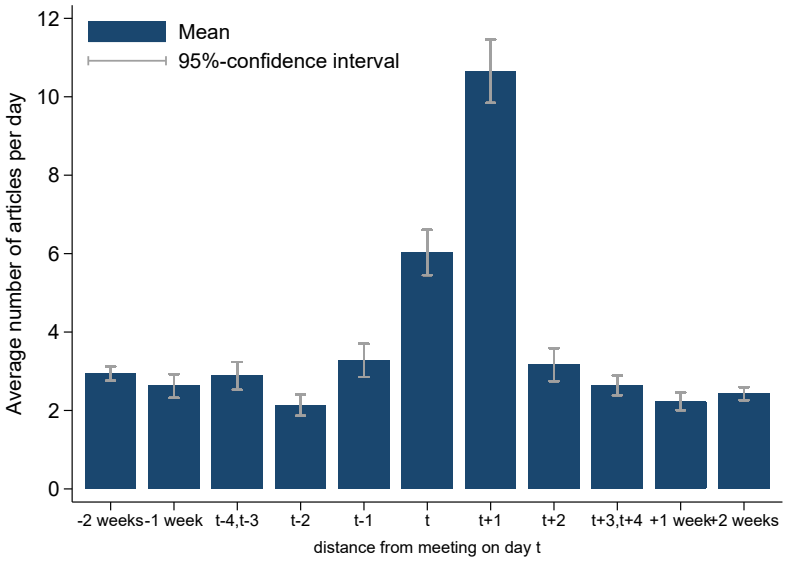

Notes: Average number of articles per day found around monetary policy meetings by searching the LexisNexis database for all articles published in US newspapers containing the three keywords 'inflation', 'central bank' and 'fed'.

2.2.3 News data

The primary source of news data are newspaper articles published in the United States. The articles are downloaded from the LexisNexis database, which contains a large set of US newspapers. Since I use the newspaper articles to verify the information effect hypothesis, they should be about monetary policy and inflation. Therefore, I focus only on articles, which contain the following three keywords: 'inflation', 'central bank' and 'fed'. Figure 2.3 shows the distribution of articles around FOMC meeting dates. There are clear spikes on the day of the meeting and the day after indicating that the search terms indeed yield relevant articles. In the analysis later on, I focus only on articles published within seven days following the meeting. I exclude the day of the meeting because most articles, especially before online reporting became more relevant, will be about what to expect from the meeting instead of the actual outcome. I limit my analysis to seven days because most articles are published within in this period and articles published later should not contain much more additional information. This also lowers the computational burden for the text analysis. Details on the text analysis approach are described in Section 2.4.

In total, the analyzed articles are from 69 different local and national newspapers. Table B.2 lists these newspapers. The most common are The New York Times and The

Washington Post[13]. The remaining top 10 newspapers are all local. Unfortunately, the LexisNexis database does not include other important local newspapers such as the Los Angeles Times or the Chicago Tribune. Generally, the resulting newspaper coverage depends both on how strongly the newspaper is represented in the LexisNexis database and on how often it publishes articles about monetary policy and inflation. Finally, Table B.2 shows that the total number of articles included in the analysis for the period from January 1994 to December 2016 is 3750. Given there are 191 meetings in the sample period, the average number of articles per meeting is about 20.

In addition to the actual newspaper articles, I also rely on the self-reported news heard by respondents in the Michigan survey. The survey asks respondents whether they heard any favorable or unfavorable news about business conditions. If they heard news, they can report what they heard in an open format. Two of these answers are provided to the researcher by classifying them based on a large set of categories (for details see questions Q5 and Q5a in Appendix B2.2). Based on the categories, I generate three measures of news heard, which relate more directly to monetary policy: news about the economic situation (good/better or bad/worse),[14] news about inflation (increasing/more or decreasing/low) and news about interest rates (increasing/high or decreasing/low). Table B.7 and Table B.8 in Appendix B2 show the mapping between the MSC categories and my categories.

Beyond these main data, I also use information on CPI inflation, food price inflation and gas price inflation as well as the federal funds target. All these are obtained from the FRED database maintained by the Federal Reserve Bank of St. Louis.

2.3 Effect of monetary policy on expectations

The focus of this chapter is the analysis of how monetary policy affects households' inflation expectations. We know that there is a lot of heterogeneity in these expectations. Therefore, I exploit the panel dimension of the MSC to control for unobserved idiosyncratic heterogeneity. In addition, observed price changes of frequently purchased goods are an important driver of inflation expectations (Coibion and Gorodnichenko 2015b; D'Acunto et al. 2019b). To take this into account, I include the current month-on-month inflation rate for food and gasoline. Specifically, I estimate the following model using the fixed effects (within) estimator:

$$\pi^{e}_{i,t} = \alpha_i + \beta \epsilon^{m}_{t-1} + \delta Z_t + u_{it}, \tag{2.1}$$

[13]Together with the online articles from WashingtonPost.com which is listed separately, The Washington Post has 485 articles.

[14]If both items reported fall into this group but refer to opposite directions of change, the households are classified as reporting 'contradicting' news.

where $\pi^e_{i,t}$ is the expected rate of inflation over the next 12 months reported by household i in month t measured in percent, ϵ^m_{t-1} is the monetary policy surprise in the previous month in percentage points,[15] and Z_t are control variables (food and gas price inflation in the baseline specification). α_i is the individual specific effect which is removed by the within transformation.[16] The error terms u_{it} are clustered at the individual level and robust to heteroscedasticity.

The coefficient of interest is β, which captures how the monetary policy surprise last month affects inflation expectations in the current month. I chose this timing because no information on the timing of the interviews for the MSC in is available. Using the previous month's surprise ensures that the households respond after the respective meeting. The drawback is that the surprise may affect other variables which than affect expectations. However, most variables except for expectations and interest rates respond slowly to monetary policy surprises. Therefore, this should not affect results strongly.[17] In addition, I show in Appendix B3 that other news released in the same month do not affect results.

2.3.1 Results

Table 2.2 displays the results of estimating the model from equation (2.1) for different samples. In all cases, positive monetary policy surprises, i.e. interest rate increases, significantly increase inflation expectations. However, the magnitude of the effect differs depending on the sample. Columns (1) and (2) include all observations with a panel dimension from January 1994 to December 2016. Comparing columns (1) and (2) highlights the importance of including current inflation rates as control variables. Both inflation measures are highly significant, and their inclusion increases the estimated effect of monetary policy on expectations.

In column (3), I exclude the largest and smallest 1% of reported inflation expectations in each month. To exclude outlier observations is standard practice when analyzing inflation expectations (see e.g., Pfajfar and Santoro 2013).[18] This exclusion does not affect results much.

[15] In months with several meetings, I use the one that was scheduled or the one that was not a conference call because those are usually the regular meetings.

[16] For simplicity, I combine the aggregate constant effect with the constant individual specific effect.

[17] In Enders, Hünnekes, and Müller (2019b), we have data on the day of response for a survey of German firms. We exploit this timing to only focus on firms responding closely around monetary policy meetings. However, we also show results using all firms responding in the month following the meeting and find qualitatively similar results.

[18] I prefer to exclude observations based on the distribution in each month instead of excluding fixed values because changes over time may make values which were previously considered unrealistic more plausible. In addition, using relative cut-offs implies that in each month roughly the same number of households are excluded which does not affect the overall sample as much.

Table 2.2: Baseline results, monetary policy and inflation expectations

	Dependent variable: inflation expectations in *t*					
	(1) All	(2) Add infl.	(3) Ex. largest exp.	(4) Ex. largest neg. FF	(5) Ex. largest pos. FF	(6) Ex. 01/2009
FF surprise, *t-1*	1.14***	1.96***	2.07***	3.66***	3.59***	1.00**
	(0.43)	(0.43)	(0.38)	(0.53)	(0.54)	(0.50)
Gas price inflation, *t*		0.04***	0.04***	0.04***	0.04***	0.04***
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Food price inflation, *t*		0.77***	0.74***	0.70***	0.67***	0.63***
		(0.07)	(0.06)	(0.06)	(0.06)	(0.06)
Observations	91748	91748	89964	88612	87630	86988
Households	45874	45874	44982	44306	43815	43494
Within R^2	0.00	0.01	0.01	0.01	0.01	0.01

Notes: FF surprise: surprise change in the 3-month federal funds future around FOMC meetings. Inflation expectations are from the Michigan Survey of Consumers, and refer to the next 12 months. Realized inflation is measured as the month-on-month change in the price level for food and gasoline products, respectively. Cluster-adjusted robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Columns (4) to (6) consider the effect of excluding large monetary policy surprises. The exclusion of the largest two negative surprises leads to a significant increase in the effect.[19] Excluding the largest two positive surprises does not have a big effect, as shown in column (5). However, one relatively large negative surprise has a large effect on the estimation. In column (6), I exclude the surprise from December 2008, i.e. survey observations from January 2009. This is the month the Fed first reached the zero lower bound (ZLB) and changed to a target range. The measured surprise in this month is -0.185 percentage points. Excluding it reduces the estimated coefficient from 3.59 to 1.00. This change highlights the importance of controlling for outliers. Therefore, I choose the specification from column (6) as the baseline.

Quantitatively the results imply a positive monetary policy surprise leads to a one for one increase in inflation expectations, i.e. a one standard deviation surprise (5 basis points or 0.05 percentage points) increases inflation expectations by 0.05 percentage points. This is a small effect, which is not surprising given the size of the observed surprises. However, the fact that these small surprises measured in a small window around the meetings impact expectations significantly at all is meaningful as such. In addition, the observed effect is probably a mixture of effects from households who do not react at all and those that do react. In Subsection 2.3.2, I further analyze the effects for different demographic groups. In fact, the effect is about twice as large as

[19] The largest two negative surprises -0.37 percentage points in April 2004 and -0.25 percentage points in October 1998.

the baseline effect for some of these groups.

The sign of the effect implies that a surprise policy tightening leads households to revise their inflation expectations upward, not downward as the standard channel of monetary policy would imply. This result is in line with households reacting more strongly to the perceived information content of the announcement than to the effect of the announced policy: if the central bank announces a policy that is more tightening than expected by financial markets, this implies that inflation may increase more than previously thought. Therefore, households revise their inflation expectations upwards. They thus pay less attention to the fact that higher interest rates should lower inflation in the future in line with standard monetary policy transmission. I support this interpretation by showing that also newspaper articles refer to high or increasing inflation more frequently after tightening surprises in Section 2.4.

In addition to considering outlier observations regarding the size of surprises, it is useful to control for special episodes of monetary policy over the course of the sample. Table 2.3 shows estimations allowing for different effects at the ZLB, during the great financial crisis of 2008/09, due to unconventional policies and given the level of the current federal funds target. In addition, I check for possible non-linearities due to the sign of the surprise as well as for the effects of rounding.

The first three columns highlight that the effects are weaker at the ZLB. At the ZLB, a positive surprise actually lowers inflation expectations. The same is true for the period of the financial crisis. Given these two periods partially overlap, this is plausible. However, the fact that at the ZLB information effects are less relevant seems somewhat surprising. This may partially be due to many different types of announcements falling into this period. Column (3) highlights that meetings with a quantitative easing (QE) announcement have very different effects compared to meetings with a forward guidance announcement (dates from Swanson 2017). In particular, forward guidance announcements have much stronger information effects. Given that most Fed forward guidance announcements can be considered Delphic (see e.g., Moessner, Jansen, and de Haan 2017), this is not surprising. Delphic forward guidance announcements reveal additional information about the future and therefore lead to positive co-movement of monetary policy and inflation expectations (Campbell et al. 2012). Beyond the simple dummy indicator for the ZLB, one can also control for the level of the federal funds rate in general. In column (4), the current fed funds target level and as well as an interaction between the fed funds target and the monetary policy surprise are included.[20] These results show that generally a lower fed funds target is associated with slightly lower inflation expectations while the reaction to the surprise is stronger the lower the

[20]For the period in which the Fed is using a target range, the midpoint of the range is used.

Table 2.3: Financial crisis, unconventional policy and state dependence

	Dependent variable: inflation expectations in *t*					
	(1) ZLB	(2) Fin. crisis	(3) QE+FG	(4) State dep.	(5) Sign	(6) Rounding
FF surprise, *t-1*	1.14**	1.51***	1.04**	4.68***	0.80	1.81***
	(0.51)	(0.57)	(0.50)	(1.34)	(0.64)	(0.53)
ZLB × FF surprise, *t-1*	-6.30**					
	(2.56)					
GFC × FF surprise, *t-1*		-3.21*				
		(1.79)				
Post GFC × FF surprise, *t-1*		-0.63				
		(2.45)				
QE × FF surprise, *t-1*			-15.07*			
			(9.05)			
Forward Guidance × FF surprise, *t-1*			22.08***			
			(7.83)			
FF target, *t-1* × FF surprise, *t-1*				-0.95***		
				(0.31)		
FF target, *t-1*				-0.03***		
				(0.01)		
Pos. FF surprise × FF surprise, *t-1*					0.71	
					(1.27)	
Exp. multiple of 5 × FF surprise, *t-1*						-1.94*
						(1.05)
Gas price inflation, *t*	0.04***	0.04***	0.04***	0.04***	0.04***	0.04***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Food price inflation, *t*	0.62***	0.63***	0.62***	0.66***	0.62***	0.63***
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)
Observations	86988	86988	86988	86988	86988	86988
Households	43494	43494	43494	43494	43494	43494
Within R^2	0.01	0.01	0.01	0.01	0.01	0.01

Notes: FF surprise: surprise change in the 3-month federal funds future around FOMC meetings. Inflation expectations are from the Michigan Survey of Consumers, and refer to the next 12 months. Realized inflation is measured as the month-on-month change in the price level for food and gasoline products, respectively. ZLB: 12/2008-11/2015. GFC = great financial crisis: 01/2008-06/2009. QE and FG dates from Swanson (2017). Cluster-adjusted robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

target. The coefficients imply that at the ZLB with a target of 0.125%,[21] a one standard deviation monetary policy surprise leads to a 0.23 percentage point increase in inflation expectations, compared to an increase of only 0.09 percentage points at a target

[21]The midpoint of the 0-0.25% range the Fed used at ZLB.

level of 3%.[22]

Finally, Table 2.3 explores two additional non-linearities. Column (5) shows that the sign of the surprise does not lead to significantly different effects. Column (6) shows that households who report a multiple of five as their inflation expectation do not respond to the monetary policy surprise (the sum of the two coefficients is not significantly different from zero). Reporting a multiple of five is considered to indicate a high degree of uncertainty about the forecast (Binder 2017c). It is plausible that individuals who are uncertain about their forecast are also individuals who are not aware of monetary policy announcements or do not pay attention to them, and thus do not react to them.[23]

An additional issue to consider is other economic news being released in the same month, which may be correlated with the surprise in monetary policy. While the small window around the meeting ensures that the interest rate surprise does not pick up any information beyond the one revealed by the central bank announcements, it is still possible that the surprise correlates with other economic news. One way to control for this is to include measures of other economic news surprises in the analysis explicitly (see e.g., Del Negro et al. 2012). Because data availability is limited on these data, I refer this additional check to the appendix. Table B.9 shows that the inclusion of other economic news does not affect results much. See Appendix B3 for details.

2.3.2 Heterogeneity

In a next step, I explore the heterogeneity across households in more detail. The choice of variables largely follows other authors who considered household heterogeneity in expectations (Binder 2017a; Coibion et al. 2019b; Das et al. 2020; Dräger and Lamla 2018). Table 2.4 displays the effect of monetary policy surprises on household inflation expectations for different demographic groups. There are no significant differences between any of the groups considered. However, the coefficients differ in the expected direction. Male respondents react more strongly than female respondents (column (1)). Households in the upper two terciles of the income distribution react more strongly than the bottom third. The effect is largest for the middle tercile though, with 1.76 (significant at the 10% level). For education, results are more mixed: respondents with at least a college degree react more strongly than those with no degree. However, those with only high school or less show the largest response (1.89 also significant at the 10%

[22] A one standard deviation surprise is 0.05 percentage points. Thus the effects are $4.68 \cdot 0.05 - 0.95 \cdot 0.05 \cdot 0.125 = 0.23$ and $4.68 \cdot 0.05 - 0.95 \cdot 0.05 \cdot 3 = 0.09$

[23] Binder (2017c) shows that rounding behavior for inflation expectations is a bit more complex and goes on to develop an uncertainty index based on time varying rounding behavior. However, for the purpose of this study the multiples of five should provide a sufficient proxy. This is supported by the fact that most extreme values are also multiples of five (see Figure B.1).

Table 2.4: Effects for different demographic groups

	Dependent variable: inflation expectations in *t*				
	(1) Gender	(2) Income	(3) Education	(4) Homeowner	(5) Stockowner
FF surprise, *t-1*	1.15*	0.55	0.82	1.37**	1.70
	(0.61)	(0.70)	(0.94)	(0.56)	(1.20)
Female × FF surprise, *t-1*	-0.30				
	(1.01)				
Income, bottom 33% × FF surprise, *t-1*		-0.29			
		(1.49)			
Income, middle 33% × FF surprise, *t-1*		1.21			
		(1.16)			
HS degree or less × FF surprise, *t-1*			1.07		
			(1.39)		
Some college, no degree × FF surprise, *t-1*			-1.14		
			(1.34)		
College degree × FF surprise, *t-1*			0.32		
			(1.30)		
Rent × FF surprise, *t-1*				-1.91	
				(1.28)	
No stocks × FF surprise, *t-1*					0.12
					(1.64)
Stocks, bottom 33% × FF surprise, *t-1*					-1.11
					(2.05)
Stocks, middle 33% × FF surprise, *t-1*					1.12
					(1.81)
Gas price inflation, *t*	0.04***	0.04***	0.04***	0.04***	0.04***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Food price inflation, *t*	0.63***	0.64***	0.63***	0.63***	0.89***
	(0.06)	(0.07)	(0.06)	(0.06)	(0.08)
Observations	86988	80152	86672	86594	56614
Households	43494	40076	43336	43297	28307
Within R^2	0.01	0.01	0.01	0.01	0.01

Notes: FF surprise: surprise change in the 3-month federal funds future around FOMC meetings. Inflation expectations are from the Michigan Survey of Consumers, and refer to the next 12 months. Realized inflation is measured as the month-on-month change in the price level for food and gasoline products, respectively. Baseline categories for the demographic variables are (1) male, (2) income in top 33%, (3) graduate degree, (4) homeowner, (5) stockowner with inv. amount in top 33%. Cluster-adjusted robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

level).[24]

[24]I compute the coefficients for the subgroups as the sum of the the average effect and the interaction effect of the subgroup. The significance level is obtained using the two standard errors and the covariance of the estimated coefficients.

In addition, I differentiate based on home ownership and stock ownership. These two variables are of course correlated with the income and education levels. Homeowners significantly increase expectations after a tightening surprise while renters do no react. Furthermore, stockowners with an investment amount in the top tercile react more strongly than those with investment in the bottom tercile. Similar to the income result the middle tercile reacts even more strongly than the top (coefficient of 2.82, significant at the 5% level).

These results are in line with male, richer and more educated respondents being more informed in general. Table B.3 in Appendix B1 shows that indeed these households are more likely to report to have heard news about the economy and inflation in particular.[25] Homeowners and stockowners are more likely to report hearing news as well. In addition to these demographic variables, individual experiences over time may also matter (Malmendier and Nagel 2016). Therefore, Table B.4 in Appendix B1 provides results by age groups and cohort (following Dräger and Lamla (2018) simply split based on whether they experienced the high inflation episodes in the 1970s). Difference between the groups are relatively small. For completeness, the table also reports some joint specifications with all demographic variables.

2.3.3 Other monetary policy measures

So far, all results are based on one measure of monetary policy and on a sample starting in 1994. Table 2.5 shows results focusing on a longer sample and additional measures. In columns (1) and (2), I extend the sample to 1990 but still use the high-frequency surprises in the 3-months federal funds futures from Jarociński and Karadi (2020).[26] In the years 1990 to 1993, the effect is also positive but not significant. In the full sample from 1990 to 2016, the effect is very similar to the baseline case (1.10 compared to 1.00, see Table 2.2).

In the next two columns, monetary policy surprises are split into the pure monetary policy (MP) shock and the central bank information (CBI) shock, as discussed in Section 2.2. Jarociński and Karadi (2020) provide two ways of splitting the surprises into these two shocks: in one case, they implement sign restrictions in a VAR. This approach allows for both types of shocks to occur at once. In the second case, they simply call the interest rate surprise a MP shock if interest rates and the stock market move in opposite directions. If they move in the same direction, the surprise is called CBI shock. This approach is labeled the 'poor man's identification'. Both shocks resulting from the VAR with sign restrictions significantly increase expectations (column (3)). When using the

[25]The results are obtained by estimating a probit estimation for the dummy variables indicating whether the hosuehold heard news about the topic specified.

[26]The data for the monetary policy surprises is not available before 1990.

Table 2.5: Different samples and monetary policy measures

	Dependent variable: inflation expectations in *t*						
	(1) 1990-1993	(2) 1990-2016	(3) JK sign	(4) JK poor man	(5) Romer, 1994-2016	(6) Romer, 1980-1993	(7) Romer, 1980-2016
FF surprise, *t-1*	0.99	1.10**					
	(1.02)	(0.49)					
MP, sign, *t-1*			1.18***				
			(0.44)				
CBI, sign, *t-1*			1.25*				
			(0.65)				
MP, poor man, *t-1*				0.79			
				(0.48)			
CBI, poor man, *t-1*				2.31***			
				(0.80)			
Romer, *t-1*					0.34***	0.19	0.21*
					(0.13)	(0.13)	(0.11)
Gas price inflation, *t*	0.05***	0.04***	0.04***	0.04***	0.03***	0.05***	0.04***
	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)
Food price inflation, *t*	0.03	0.24***	0.62***	0.61***	0.06	0.01	0.06
	(0.06)	(0.05)	(0.06)	(0.06)	(0.07)	(0.06)	(0.05)
Observations	69204	158068	86988	86988	53762	69956	127662
Households	34602	79034	43494	43494	26881	34978	63831
Within R^2	0.00	0.00	0.01	0.01	0.00	0.00	0.00

Notes: FF surprise: surprise change in the 3-month federal funds future around FOMC meetings. Inflation expectations are from the Michigan Survey of Consumers, and refer to the next 12 months. Realized inflation is measured as the month-on-month change in the price level for food and gasoline products, respectively. MP: monetary policy shock from Jarociński and Karadi (2020). CBI: central bank information shock from Jarociński and Karadi (2020). 'sign' and 'poor man' refer to the different identification strategies. Romer: monetary policy shocks from Wieland and Yang (2020), based on Romer and Romer (2004). Cluster-adjusted robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

shocks based on the sign of the changes in interest rates and the stock market (poor man's identification, column (4)), only the CBI shock has a significant effect, although both coefficients are positive. The CBI shocks is precisely constructed to measure information provided by the central bank. In particular, a positive CBI shock implies information about more expansionary economic developments, i.e. more growth and higher inflation. This supports the interpretation that the overall response to the monetary policy surprise is due to information effects. In addition, the positive coefficient of the MP shock and the difference in the results with the two shock versions points to potential differences in how financial markets and other actors perceive monetary policy surprises.

The last three columns of Table 2.5 show results using the monetary policy surprises provided by Wieland and Yang (2020) based on Romer and Romer (2004). For the baseline sample in column (5), a positive, i.e. tightening, Romer shock also has a positive and significant effect on inflation expectations. This in line with the response to the high-frequency monetary policy surprises. The magnitude cannot be compared directly because the Romer shocks are residual measures. One way to think about magnitudes is to compare the effect of a one standard deviation change. In the baseline sample the standard deviation of the Romer shocks is 0.16 percentage points, implying an effect of $0.34 \cdot 0.16 = 0.054$ percentage points. This is very close to the effect of a one standard deviation high-frequency surprise, which is 0.05 percentage points (see above). The Romer shock series only measures surprise movements by the Fed based on the Fed's internal forecasts (for details see Subsection 2.2.2). However, these forecasts are not available to the general public at the time. Therefore, it is plausible that households also react to these surprises.

The Romer series is available since 1969. It therefore allows for an estimation with a longer time dimension than the data by Jarociński and Karadi (2020). Since the panel dimension of MSC data starts in July 1980, columns (6) and (7) show results for samples from July 1980 to December 1993 and December 2016 respectively. Effects are weaker and less significant when including these earlier observations indicating that information effects were less strong before 1994. This result makes sense in light of the less frequent and more opaque communication by the Fed in this earlier period.

To summarize, the results in this section show that monetary policy surprises increase inflation expectations of US households. The size of the effect depends on the sample and on current conditions. However, the direction is quite robust to various specifications. I interpret this finding as evidence for the information effect being present for households. Similar findings already exist for financial markets and professional forecasters (Campbell et al. 2012; Nakamura and Steinsson 2018) and firms (Enders et al. 2019b). In the following section, I will use information about actual newspaper reporting around monetary policy meetings to corroborate this interpretation.

2.4 Transmission via newspaper reporting

So far, we have seen that monetary policy surprises increase inflation expectations. However, it is commonly understood that most households do not know a lot about central banks and do not follow monetary policy decisions closely (Binder 2017b). Instead, they use other sources of information such as shopping experiences but also media reporting when forming expectations (D'Acunto et al. 2019b). Therefore, one

potential transmission channel of monetary policy to households is reporting by newspapers.

The basic idea is that households may not actively follow Federal Reserve policy but do regularly read newspapers. After central bank decisions, newspapers report about the decision and, in particular, report the perceived view of the central bank on inflation. Specifically, after decisions, which were more tightening than expected, the articles refer more to the risk or expectation of high inflation. When seeing these articles, households notice the information about current and future inflation but may not pay much attention to the implied effects of monetary policy on future inflation. If this is the case, the information will lead them to increase their own inflation expectations after tightening announcements and vice versa after easing announcements. In this section, I will provide evidence for this channel.

2.4.1 Text analysis

In order to analyze how newspaper reporting transmits monetary policy surprises to households, I first need to develop an appropriate measure of reporting. Since the focus of this chapter are inflation expectations, I will use a simple text analysis method, called dictionary or lexicon method, to classify articles into those referring to high or increasing inflation versus those referring to low or decreasing inflation. No off-the-shelf dictionary is available for this kind of classification. Therefore, I develop my own dictionary.[27] For this purpose, I first separate all articles published within a week of the FOMC meeting and selected based on the keywords described in Subsection 2.2.3, into sentences.[28] Then I locate all instances of the word 'inflation' and tabulate all words occurring close to this keyword. From this list, I choose words indicating either high/increasing inflation or low/decreasing inflation. Table B.10 in Appendix B4 shows the resulting lists. In a final step, I classify each sentence to refer to one of the two categories if one of the keywords occurs within five words before or after 'inflation'. The choice of five words is based on preliminary analysis weighting the probability of creating false positives and false negatives. Five words allow sufficiently many words to be considered to capture language that is more complex but ensure the keyword still refers to 'inflation'. In addition, I adjust for negations and conjunctions as well as some recurring phrases. I also manually check the classification of the sentences

[27] There are several dictionaries for classifying text into positive and negative based on the sentiment expressed therein (see Shapiro, Sudhof, and Wilson 2019, for an overview and discussion). Picault and Renault (2017) develop a dictionary for classifying all ECB statements into hawkish, dovish or neutral by manually classifying sentences within statements and the searching for phrases and words which best predict the human coding.

[28] Before separating articles into sentences, I use the LexisNexisTools package provided by Johannes Gruber at `https://github.com/JBGruber/LexisNexisTools` to convert the articles downloaded from LexisNexis into a format suitable for text analysis.

Table 2.6: News classification at the article level

	Group of articles considered:				
	'up' only	'down' only	'up'+'down'	neither	all
Number of articles	1018	627	475	1630	3750
Share in all articles	0.271	0.167	0.127	0.435	1
Avg. share of 'up' sentences	0.054	0	0.055	0	0.021
Avg. share of 'down' sentences	0	0.049	0.049	0	0.014
Avg. diff. in sentence shares	0.054	-0.049	0.006	0	0.007

Notes: Descriptive statistics for the different news reporting measures. Based on articles from LexisNexis database published in US newspapers within seven days of FOMC meetings in the sample period, and containing the three keywords 'inflation', 'central bank' and 'fed'. 'up' means sentence/article refers to high or increasing inflation. 'down' means sentence/article refers to low or decreasing inflation. Different columns refer to classification at the article level.

for randomly drawn subsets. For details on the procedure and the checks, please see Appendix B4.

This procedure provides me with a count of sentences about high/increasing inflation and low/decreasing inflation for each article. There are different ways to aggregate this measure to the frequency of the FOMC meetings. I use two different types of measures. The first is constructed at the article level. I define two binary indicators for each article: an article is defined to be an 'up' article if it includes at least one sentence referring to high or increasing inflation, and it is defined to be a 'down' article if it includes at least one sentence referring to low or decreasing inflation. This means articles can also be part of both categories. Table 2.6 shows the resulting distribution for all articles published within seven days after the FOMC meetings. In total 39.8% of all articles are classified as 'up': 27.1% are only 'up' and another 12.7% are both 'up' and 'down'. 'Down' articles are less frequent: in total 29.4% are 'down' articles, and 16.7% are only 'down'. Finally, these measures can be aggregated to the meeting level by simply computing the share of each article type in all articles published after the meeting. In addition, I also construct a difference measure by subtracting the total share of 'down' articles from the total share of 'up' articles for each meeting.

The second measure is constructed at the sentence level. For this, I simply compute the share of 'up' sentences (i.e. referring to high or increasing inflation) in all sentences of each article, and the share of 'down' sentences in all sentences. I can again construct a difference measure by subtracting the two shares. Table 2.6 displays the resulting shares. The average share of 'up' sentences is 2.1% for all articles and 5.4% for articles with at least one 'up' sentence. The average share of 'down' sentences is slightly lower in the full sample with 1.4%. This is partially due to fewer articles containing any

Table 2.7: Newspaper reporting and monetary policy

	Dependent variable: news reporting index, sentence-based measure			article-based measure		
	(1) 'up'	(2) 'down'	(3) diff.	(4) 'up'	(5) 'down'	(6) diff.
FF surprise, t	0.027**	0.004	0.023	0.356**	0.188	0.168
	(0.012)	(0.013)	(0.019)	(0.161)	(0.214)	(0.271)
Constant	0.021***	0.015***	0.006***	0.393***	0.314***	0.079***
	(0.001)	(0.001)	(0.001)	(0.015)	(0.015)	(0.023)
Observations	188	188	188	188	188	188
Adjusted R^2	0.00	-0.01	-0.00	0.00	-0.00	-0.00

Notes: FF surprise: surprise change in the 3-month federal funds future around FOMC meetings. 'up' means sentence/article refers to high or increasing inflation. 'down' means sentence/article refers to low or decreasing inflation. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

'down' sentences and partially due to a lower share of 'down' sentences in articles, which do contain 'down' sentences (4.9%). Finally, the last row of Table 2.6 shows the difference measure at the sentence level. In case of the 'up only' and 'down only' articles it simply equals the share of the respective group (with a negative sign for 'down' articles). In the case of the categories where both 'up' and 'down' articles exists, it equals the difference between the shares. For all articles, the difference measure equals 0.7 percentage points. For the group of articles, which contain both types of sentences it is 0.6 percentage points. To get a measure at the meeting level, I again average over all articles published after each meeting.

2.4.2 Results

In order to establish how news reporting transmits monetary policy to households, I proceed in two steps. First, I show that a tightening monetary policy surprise increases the amount of reporting referring to high or increasing inflation. In a second step, I analyze how monetary policy and news reporting jointly affect inflation expectations.

Table 2.7 shows how monetary policy surprises affect newspaper reporting. I obtain the results by simply regressing the different measures of news reporting for articles published after a FOMC meeting on the high-frequency identified monetary policy surprises of the respective meeting.[29] The effect is only significant for the measures of reporting about high or increasing inflation ('up' measure). The coefficient is positive, both for the sentence-based and the article-based measure. This result is in line with the hypothesis that after tightening surprises reporting about high inflation increases.

[29]The regression model is a simple OLS model with robust standard errors.

However, the effect is insignificant for the measures of 'down' sentences as well as the difference measures. This indicates possible problems with the relatively simple text analysis methodology. A more complex approach such as proposed by Shapiro et al. (2019) or Picault and Renault (2017) may provide results that are more precise. However, this is beyond the scope of this chapter.

To support the results found, I also rely on the self-reported news from the Michigan Survey. The advantage of this measure is that we can be sure households actually heard the news. The drawback is the relatively low coverage. As already discussed above, only two news items are recorded per household, limiting the total number of specific news items reported. Nevertheless, Table B.5 in Appendix B1 shows that tightening surprises increase the probability of households hearing news about increasing inflation and higher interest rates.[30] These results are in line with the effect on actual reporting.

Finally, I want to understand how reporting affects inflation expectations and how much of the overall effect of monetary policy is due to reporting. For this purpose, I resort to a simple form of mediation analysis as used in Das et al. (2020). This simple approach relies on jointly estimating the effect of the explanatory variables (monetary policy in this case) and the mediator (news reporting). I therefore estimate the model from equation (2.1) again, but now also include the different measures of newspaper reporting. The timing is the same as before. I include the reporting associated with the monetary policy surprise from the previous month. The sample is the same as the baseline sample in Table 2.2.

Table 2.8 displays the results. In columns (1) and (3), I jointly include the two 'up' and 'down' measures due to the overlap between 'up' and 'down' articles for some meetings. As expected, more reporting about high or increasing inflation increases inflation expectations while reporting about low or decreasing inflation significantly lowers expectations. Interestingly, the effect of reporting about low inflation is almost twice as large as the effect of reporting about more inflation. This may be related to the fact that reporting about low inflation is more infrequent so it might be more noticeable when it does occur. The effects are independent of the measure used. The quantitative difference between the coefficients is related to the different quantitative levels of the two measures. Finally, the difference measures both increase inflation expectations, which is no surprise given the effects of the two separate measures.[31]

Based on these results one can discuss the direct and indirect effects of monetary

[30]The dependent variables are dummies whether the household heard something about higher prices/increasing inflation vs. lower prices/decreasing inflation and higher interest rates vs. lower interest rates respectively. Details on these questions can be found in Appendix B2.

[31]For completeness, Table B.6 in the appendix provides similar results using the self-reported news again. The direction of the effects is in line with results found for actual reporting.

Table 2.8: Expectations, monetary policy and newspaper reporting

	Dependent variable: inflation expectations in *t*			
	Sentence-based measure		Article-based measure	
	(1) 'up'/'down'	(2) diff.	(3) 'up'/'down'	(4) diff.
FF surprise, *t-1*	0.84*	0.95*	1.14**	1.22**
	(0.51)	(0.50)	(0.51)	(0.50)
Share of 'up' sentences, *t-1*	4.04***			
	(1.06)			
Share of 'down' sentences, *t-1*	-7.07***			
	(1.14)			
Diff. in sentences shares, *t-1*		5.37***		
		(0.87)		
Share of 'up' articles, *t-1*			0.28***	
			(0.07)	
Share of 'down' articles, *t-1*			-0.50***	
			(0.07)	
Diff. in articles shares, *t-1*				0.38***
				(0.06)
Gas price inflation, *t*	0.04***	0.04***	0.04***	0.04***
	(0.00)	(0.00)	(0.00)	(0.00)
Food price inflation, *t*	0.63***	0.62***	0.64***	0.62***
	(0.06)	(0.06)	(0.06)	(0.06)
Observations	86988	86988	86988	86988
Households	43494	43494	43494	43494
Within R^2	0.01	0.01	0.01	0.01

Notes: FF surprise: surprise change in the 3-month federal funds future around FOMC meetings. Inflation expectations are from the Michigan Survey of Consumers, and refer to the next 12 months. Realized inflation is measured as the month-on-month change in the price level for food and gasoline products, respectively. 'up' means sentence/article refers to high or increasing inflation. 'down' means sentence/article refers to low or decreasing inflation. Cluster-adjusted robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

policy. When including the sentence-based measures, the direct effect of monetary policy is somewhat smaller than without news reporting (0.84 and 0.95 compared to 1.00 in Table 2.2). When including the article-based measure, the effect actually increases slightly. The indirect effect of monetary policy on expectations can be computed by combining the effect of monetary policy on reporting from Table 2.7 with the effect of reporting on expectations in Table 2.8. For the case of the share of 'up' sentences, the indirect effect of monetary policy on inflation expectations is $0.027 \cdot 4.04 = 0.11$, i.e. 11% of the total effect in the baseline specification. Using the article-based 'up' measure it is: $0.356 \cdot 0.28 = 0.10$ or 10% of the total effect.

Quantitatively the news reporting measures thus explain only a small share of the overall effect of monetary policy. This may be due to several reasons. We have already seen above that the measures derived here may be imprecise. In addition, the newspaper coverage is limited and does not include all major (regional) US newspapers. Finally, households also depend on other media such as TV and radio for information on expectations (D'Acunto et al. 2019b). Therefore, the derived measures of news reporting may not fully cover all the news households are exposed to. A more comprehensive analysis of all types of reporting with a more complex text analysis approach may lead to quantitatively more relevant and more precise results.

Generally, however the results support the hypothesis that surprise tightening announcements increase inflation expectations due to the information effect. After surprise tightenings, newspapers report more about high or increasing inflation, which in turn influences household inflation expectations. They thereby provide a potential transmission channel from the central bank to the households.

2.5 Conclusion

This chapter shows that monetary policy directly affects households' inflation expectations in the US. A surprise tightening of monetary policy increases inflation expectations. The effect is quantitatively small, also due to the small size of the high-frequency surprises used to quantify monetary policy announcements. However, the direction of the effect is significant and robust to different specifications. The results support the hypothesis that agents in the economy look to central bank decisions to learn something about the current outlook for the economy. This interpretation is supported by the fact that news reported and news heard react in a similar way. After surprise tightening announcements more articles mention high or increasing inflation and households report more frequently that they heard news about high or increasing inflation.

The results have important implications for monetary policy. First, they highlight that monetary policy does have a direct effect on expectations. This is in line with other recent analyses such as Lewis et al. (2019) and Enders et al. (2019b). At the same time, it is not necessarily contradicting evidence that households are not very informed about central banks (Binder 2017b). Evidence by D'Acunto et al. (2019b), among others, shows that households do update their expectation if they are directly confronted with news about inflation or the central bank. Thus, my results indicate that after monetary policy meetings a sufficiently large amount of people reads the news about inflation and reacts accordingly. Given the effects of direct information treatments in survey experiments, this implies quantitative effects should increase when central bank communication reaches more people.

The second implication is that more and better communication has non-trivial effects on the economy. The information effect highlighted here leads to updating in the opposite direction of the one intended by the policy change. If the central bank is more worried about inflation than previously thought and accordingly raises interest rates in order to reduce inflation pressures, an increase in household inflation expectations limits the effectiveness of this policy. This is especially relevant when conventional policy is constrained due to the effective lower bound. Nakamura and Steinsson (2018) argue that in normal times providing additional information is likely welfare improving even with an offsetting information effect because the central bank can always counteract with additional policy changes. In addition, it is likely that the information would have been revealed to households anyway, only with a delay. In this case, the early revelation of the information by central banks does not alter the path of the economy much. However, at the effective lower bound negative spirals are possible. In that case, it may be useful to withhold negative information (Nakamura and Steinsson 2018).

Wiederholt (2015) models the effects of more communication by the central bank explicitly in a simple New Keynesian model with information frictions. The author shows that there is one central trade-off central banks face. Generally, more information can be considered welfare improving because it allows agents to make efficient decisions and because it reduces dispersion across households. However, incomplete information on the part of some agents provides an insurance mechanism against shocks hitting the economy, as some agents are not aware of how bad (good) the current state is. If the economy is at the effective lower bound and more people learn about this due to central bank communication, this may make the situation worse as they reduce consumption further.

In light of these potentially problematic consequences of central bank communication, additional research on the effects and transmission mechanisms is needed. In particular, the results of this study should be confirmed with more detailed survey data, especially regarding the timing of the responses. In addition, the measure of news reporting about monetary policy could be improved in terms of media covered, topics considered (beyond inflation) and the text analysis approach used. In light of the monetary policy response to the next crisis, these are highly relevant topics – especially, given the fact that the ECB is still at the effective lower bound and the Fed just moved back to it in March 2020.

Appendices to Chapter 2

B1 Additional tables and figures

Figure B.1: Cross-sectional distribution of inflation expectations

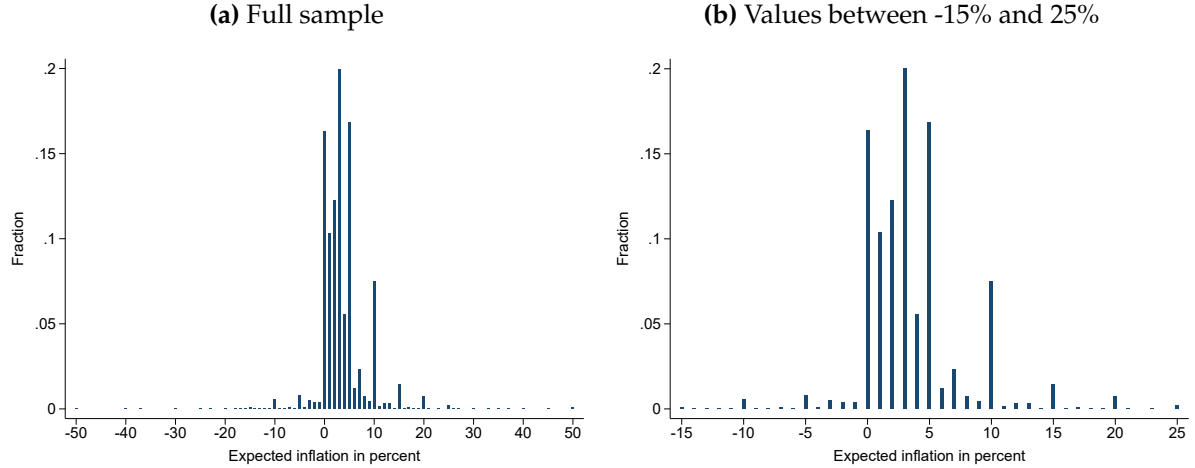

Notes: Histograms how quantitative inflation expectations for the next 12 months from the Michigan Survey of Consumers, from January 1994 to December 2016.

Table B.1: Demographic distribution in different samples

	All	With exp.	Base-line		All	With exp.	Base-line
Gender				Education			
Male	0.47	0.49	0.50	High school or less	0.34	0.32	0.30
Female	0.53	0.51	0.50	Some college, no degree	0.26	0.27	0.26
Renter/owner				College degree	0.24	0.25	0.26
Owns home	0.76	0.77	0.80	Graduate studies	0.16	0.17	0.18
Rents home	0.24	0.23	0.20	Income			
Cohort				bottom 33%	0.27	0.25	0.23
Born before 1970	0.81	0.80	0.83	middle 33%	0.34	0.35	0.35
Born 1970 and later	0.19	0.20	0.17	top 33%	0.38	0.40	0.42
Age				Investment			
18-39 years	0.30	0.30	0.28	No stocks	0.44	0.41	0.37
40-59 years	0.39	0.40	0.41	Inv. in bottom 33%	0.17	0.17	0.18
60 years and older	0.31	0.30	0.31	Inv. in middle 33%	0.18	0.19	0.21
				Inv. in top 33%	0.20	0.22	0.24

Notes: 'With exp.' refers to households who provided an answer to the question on quantitative inflation expectations for the next 12 months (Q2). 'Baseline' refers to sample of household responding twice to the survey and reporting inflation expectations both times. Specific household types are more likely not to respond to the inflation expectations question, and not to respond to the second interview. This leads to differences in the shares in both samples. Generally, women and households with a lower socio-economic status are less likely to respond twice and to respond to inflation expectations.

Table B.2: List of newspaper included and number of articles

	Newspaper	No. art.		Newspaper	No. art.
1	The New York Times	678	36	San Gabriel Valley Tribune	18
2	The Washington Post	417	37	The Hill	14
3	Pittsburgh Post-Gazette	243	38	Inland Valley Daily Bulletin	13
4	Charleston Gazette	207	39	The Detroit News	11
5	St.Louis Post-Dispatch	169	40	The New York Sun	9
6	Telegraph Herald	164	41	Fort Wayne Journal-Gazette	8
7	Deseret Morning News	155	42	San Bernardino Sun	8
8	St.Paul Pioneer Press	148	43	Lowell Sun	7
9	Bismarck Tribune	126	44	Pasadena Star- News	7
10	The Philadelphia Inquirer	118	45	Long Island Business News	6
11	San Jose Mercury News	113	46	Alameda Times-Star	5
12	The Columbian	89	47	New Orleans City Business	5
13	Chicago Daily Herald	70	48	Lincoln Journal Star	4
14	South Bend Tribune	69	49	San Mateo County Times	4
15	WashingtonPost.com	68	50	Tri-Valley Herald	4
16	The Pantagraph	66	51	Daily Journal of Commerce	3
17	Austin American-Statesman	63	52	Intelligence Journal	3
18	The Atlanta Journal and Constitution	59	53	Long Beach Press-Telegram	3
19	Monterey County Herald	51	54	Marin Independent Journal	3
20	Dayton Daily News	48	55	Santa Fe New Mexican	3
21	The Daily Record	46	56	Tampa Bay Times	3
22	The New York Post	44	57	The Wall Street Journal	3
23	The Tampa Tribune	43	58	Daily Camera	2
24	Pittsburgh Tribune Review	39	59	The Daily Record of Rochester	2
25	The Orange County Register	38	60	The Evening Sun	2
26	St.Petersburg Times	37	61	Colorado Springs Business Journal	1
27	Wisconsin State Journal	36	62	Finance & Commerce	1
28	Daily News	28	63	Mississippi Business Journal	1
29	The Daily News of Los Angeles	27	64	New York Observer	1
30	Inside Bay Area	25	65	Roll Call	1
31	Philadelphia Daily News	25	66	The Indianapolis Business Journal	1
32	Star Tribune	22	67	The Mecklenburg Times	1
33	The Oklahoman	21	68	Tribune-Review	1
34	Capital Times	20	69	Wyoming Tribune-Eagle	1
35	Creators Syndicate	19		Total	3750

Notes: List of all newspapers from which articles are used in the analysis, i.e. all articles from LexisNexis database published in US newspapers within seven days of FOMC meetings in the sample period, and containing the three keywords 'inflation', 'central bank' and 'fed'.

Table B.3: Demographics and news heard

	Dependent var.: indicator of having heard news in t			
	(1) General	(2) Economic situation	(3) Inflation	(4) Interest rates
Female	-0.03***	0.01	-0.09***	-0.18***
	(0.01)	(0.01)	(0.01)	(0.02)
Income, bottom 33%	-0.15***	-0.14***	-0.00	-0.29***
	(0.01)	(0.01)	(0.02)	(0.03)
Income, middle 33%	-0.08***	-0.06***	-0.01	-0.15***
	(0.01)	(0.01)	(0.02)	(0.02)
HS degree or less	-0.60***	-0.45***	-0.14***	-0.43***
	(0.01)	(0.01)	(0.02)	(0.02)
Some college, no degree	-0.30***	-0.21***	-0.04**	-0.25***
	(0.01)	(0.01)	(0.02)	(0.02)
College degree	-0.14***	-0.08***	0.01	-0.06***
	(0.01)	(0.01)	(0.02)	(0.02)
Rent	-0.02*	-0.05***	-0.03*	-0.08***
	(0.01)	(0.01)	(0.02)	(0.02)
No stocks	-0.34***	-0.28***	-0.11***	-0.20***
	(0.01)	(0.01)	(0.02)	(0.02)
Stocks, bottom 33%	-0.11***	-0.08***	-0.02	-0.07***
	(0.01)	(0.01)	(0.02)	(0.02)
Stocks, middle 33%	-0.11***	-0.08***	-0.04**	-0.06***
	(0.01)	(0.01)	(0.02)	(0.02)
18-39 years	-0.23***	-0.18***	-0.06***	0.15***
	(0.01)	(0.01)	(0.02)	(0.02)
40-59 years	0.04***	0.05***	0.12***	0.20***
	(0.01)	(0.01)	(0.02)	(0.02)
Constant	0.96***	0.52***	-1.41***	-1.40***
	(0.01)	(0.01)	(0.02)	(0.02)
Observations	109898	100160	110394	110394
Pseudo R^2	0.06	0.04	0.01	0.06

Notes: FF surprise: surprise change in the 3-month federal funds future around FOMC meetings. Baseline categories are male, income in top 33%, graduate degree, homeowner, stockowner with inv. amount in top 33%, and 60 years and older. Dependent variables are dummy variables, where 1 indicates household heard news referring to the topic in the column header. 0 indicates no news heard. Details on news questions can be found in Appendix B2.2. Coefficients are from probit estimations. Cluster-adjusted robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.4: Effect for additional demographic groups

	Dependent var.: inflation expectations in t				
	(1) Age	(2) Cohort	(3) Joint spec. 1	(4) Joint 2	(5) Joint 3
FF surprise, *t-1*	0.96	0.97*	0.84	2.84*	1.62
	(1.03)	(0.54)	(0.79)	(1.63)	(1.97)
18-39 years × FF surprise, *t-1*	0.60				1.25
	(1.40)				(2.56)
40-59 years × FF surprise, *t-1*	-0.35				1.77
	(1.27)				(1.66)
Born 1970 and later × FF surprise, *t-1*		0.17			0.78
		(1.50)			(2.71)
Female × FF surprise, *t-1*			-0.22	-1.39	-1.33
			(1.05)	(1.39)	(1.40)
Rent × FF surprise, *t-1*			-1.93	-1.31	-1.56
			(1.40)	(1.88)	(1.93)
Income, bottom 33% × FF surprise, *t-1*			0.39	-0.16	0.48
			(1.61)	(2.21)	(2.29)
Income, middle 33% × FF surprise, *t-1*			1.45	2.11	2.30
			(1.18)	(1.71)	(1.73)
HS degree or less × FF surprise, *t-1*				-0.64	-0.77
				(2.03)	(2.04)
Some college, no degree × FF surprise, *t-1*				-2.96	-3.06
				(2.01)	(2.01)
College degree × FF surprise, *t-1*				-0.61	-0.75
				(1.85)	(1.85)
No stocks × FF surprise, *t-1*				0.31	0.10
				(2.10)	(2.13)
Stocks, bottom 33% × FF surprise, *t-1*				-1.60	-2.07
				(2.27)	(2.34)
Stocks, middle 33% × FF surprise, *t-1*				1.33	1.12
				(1.94)	(1.99)
Gas price inflation, t	0.04***	0.04***	0.04***	0.04***	0.04***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Food price inflation, t	0.62***	0.62***	0.65***	0.87***	0.89***
	(0.06)	(0.06)	(0.07)	(0.09)	(0.09)
Observations	86542	86540	79924	53716	59982
Households	43271	43270	39962	26858	33193
Within R^2	0.01	0.01	0.01	0.01	0.01

Notes: FF surprise: surprise change in the 3-month federal funds future around FOMC meetings. Inflation expectations are from the Michigan Survey of Consumers, and refer to the next 12 months. Realized inflation is measured as the month-on-month change in the price level for food and gasoline products, respectively. Baseline categories are (1) 60 years and older, (2) born before 1970, (3)-(5) male, income in top 33%, graduate degree, homeowner, stockowner with inv. amount in top 33%, 60 years and older, and born before 1970. Cluster-adjusted robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.5: Monetary policy and self-reported news

	Dependent variable: type of news heard on					
	economy		inflation		interest rates	
	(1)	(2)	(3)	(4)	(5)	(6)
FF surprise, *t-1*	0.05	0.05	0.83**	0.74**	1.63***	1.65***
	(0.13)	(0.13)	(0.34)	(0.34)	(0.52)	(0.52)
Gas price inflation, *t*		0.00***		0.01***		-0.00
		(0.00)		(0.00)		(0.00)
Food price inflation, *t*		0.07***		-0.00		-0.03
		(0.02)		(0.07)		(0.07)
Observations	22892	22892	1026	1026	960	960
Households	11446	11446	513	513	480	480
Within R^2	0.00	0.01	0.01	0.09	0.02	0.02

Notes: FF surprise: surprise change in the 3-month federal funds future around FOMC meetings. Realized inflation is measured as the month-on-month change in the price level for food and gasoline products, respectively. Dependent variables are dummy variables, where 1 indicates 'good/better' for economy news and 'high/increasing' for inflation and interest rates. 0 indicates the opposite. Cases of contradicting news or no news are excluded. Details on news questions can be found in Appendix B2.2. Coefficient results from a fixed effects regression. Cluster-adjusted robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.6: Expectations, monetary policy and self-reported news

	Dep. var.: inflation exp. in *t*		
	(1) Economic sit.	(2) Inflation	(3) Interest rates
FF surprise, *t-1*	1.00**	0.93*	0.97*
	(0.50)	(0.50)	(0.50)
Good/better situation	0.02		
	(0.04)		
Bad/worse situation	-0.04		
	(0.04)		
Contradicting news	0.11**		
	(0.05)		
Higher prices/increasing infl.		0.69***	
		(0.08)	
Lower prices/decreasing infl.		-0.27***	
		(0.10)	
Higher interest rates			0.15*
			(0.08)
Lower interest rates			-0.14*
			(0.08)
Gas price inflation, *t*	0.04***	0.04***	0.04***
	(0.00)	(0.00)	(0.00)
Food price inflation, *t*	0.62***	0.60***	0.63***
	(0.06)	(0.06)	(0.06)
Observations	86988	86988	86988
Households	43494	43494	43494
Within R^2	0.01	0.01	0.01

Notes: FF surprise: surprise change in the 3-month federal funds future around FOMC meetings. Inflation expectations are from the Michigan Survey of Consumers, and refer to the next 12 months. Realized inflation is measured as the month-on-month change in the price level for food and gasoline products, respectively. Additional explanatory variables are dummy indicators for type of news heard on different topics: economic situation in column (1), inflation in column (2) and interest rates in column (3). Cluster-adjusted robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

B2 Questions in the Michigan Survey

B2.1 Inflation expectations

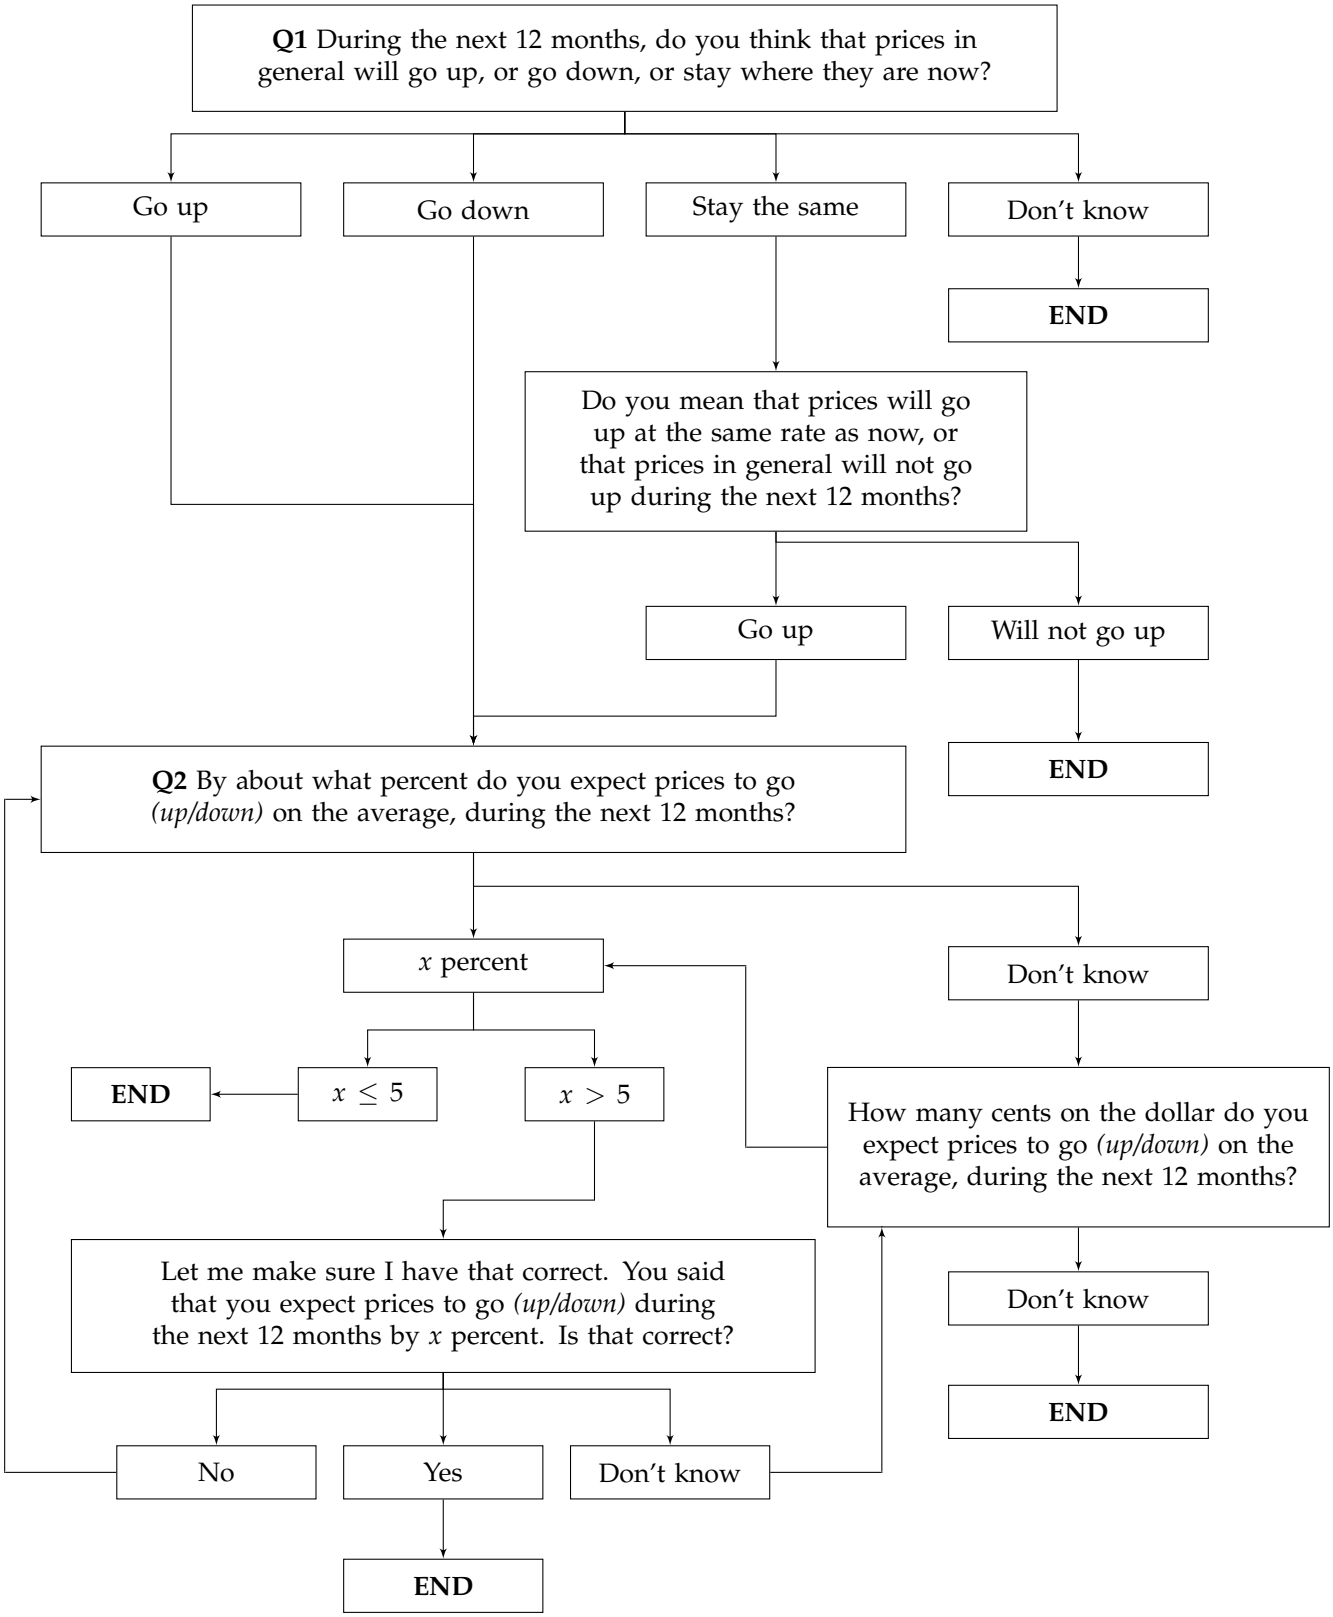

B2.2 Other questions

Q3 Homer ownership Do you (and your family living there) own your home, pay rent, or what? *(Open answer possible)*

Q4 Stock ownership The next questions are about investments in the stock market. First, do you (or any member of your family living there) have any investments in the stock market, including any publicly traded stock that is directly owned, stocks in mutual funds, stocks in any of your retirement accounts, including 401(K)s, IRAs, or Keogh accounts? *(Yes or no answer)*

Q4a Stock amount If "yes" for Q4: Considering all of your (family's) investments in the stock market, overall about how much would your investments be worth today? *(Probe: What is your best estimate?)*

Q4b Stock amount bracket If "don't know" for Q4a: Would the total be $100,000 or more?

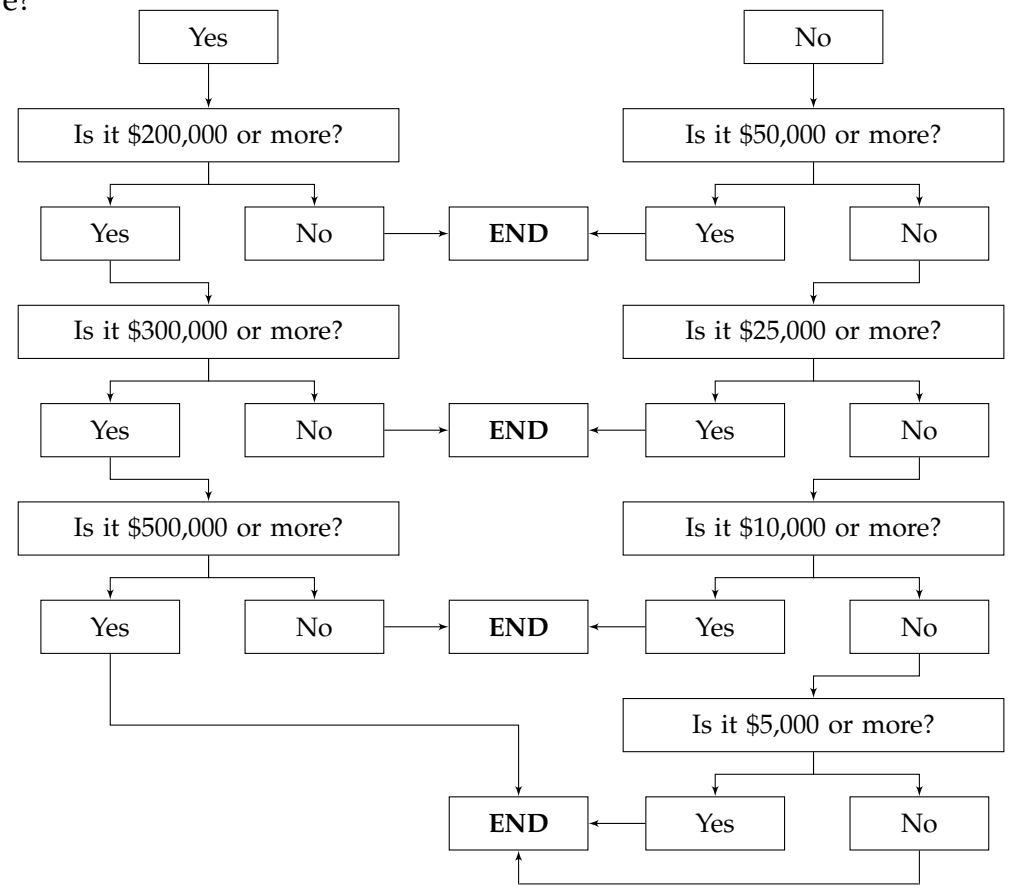

Q5 News heard During the last few months, have you heard of any favorable or unfavorable changes in business conditions? *(Yes or no answer)*

Q5a News type If "yes" for Q5: What did you hear? *(Open answer, up to two answers are reported to researchers using classification scheme detailed in Table B.7 and Table B.8.)*

B2.3 Classification of self-reported news

Table B.7: Classification of news, items categorized as favorable by respondents

MSC report	Economic Sit.	Inflation	Interest rates
Elections, admin, Congress, President			
More military spending, more war/tensions			
Less military spending, few tensions			
Gov't programs improved			
Specific gov't programs incr/cont			
Specific gov't programs decr/end			
Taxes, changes/reforms, rebates			
Other references to gov't			
Fiscal policy, budgets, deficits			
Gov't good/better business condition	good/better		
Opening of plants, factories, stores	good/better		
Consumer/auto demand high	good/better		
Purch power high, wages high	good/better		
Employ is high, plenty of jobs	good/better		
Population increase, more people to buy	good/better		
Low debts, higher savings/assets, invest up	good/better		
Other references to employ and purch power	good/better		
Production increasing, GNP is up	good/better		
Unemp has risen, good for economy			
Tight money, int rates high	good/better		increasing/high
Lower/stable prices, less inflation	bad/worse	decreasing/less	
Higher prices, inflation is good	good/better	increasing/more	
Easier money, credit easy to get, low int rates	bad/worse		decreasing/low
Profits high/rising	good/better		
Stock market, rise in price of stocks	good/better		
Other references to prices/credit			
Balance of payments, dollar devalue			
Controls (price or wage)			
Better race relations, less crime			
Union disputes settled, relations good			
Times/business is good in the coming year	good/better		
Bad times can't last, due for good times	good/better		
R sees sign of improvement already	good/better		
Improvements in specific industries	good/better		
Farm situation good, crops good	good/better		
Other good factors or favorable ref			
Economy more stable, optimism	good/better		
Energy crisis, pollution			

Notes: This classification is provided by the MSC when downloading the data from `https://data.sca.isr.umich.edu/sda-public/cgi-bin/hsda?harcsda+sca`.

Table B.8: Classification of news, items categorized as unfavorable by respondents

MSC report	Economic Sit.	Inflation	Interest rates
Election, new admin/President			
More military spending, more war/tensions			
Less military spending, few tensions			
Specific gov't spend programs changed			
Specific gov't spend programs eliminated			
Gov't progams begun/increased			
Taxes, changes/reforms, rebates			
Other references to gov't			
Fiscal policy, budgets, deficits			
Gov't not good/better business conditions	bad/worse		
Closing of plants, factories, stores	bad/worse		
Consumer/auto demand low	bad/worse		
Lack of purch power, no money to spend	bad/worse		
Drop in employ, less overtime	bad/worse		
Population increase, immigration			
High(er) debts, lower savings/assets	bad/worse		
Other references to employ/purch power			
Production decreasing, GNP down	bad/worse		
Prices falling, deflation	bad/worse	decreasing/less	
Prices high, inflation	good/better	increasing/more	
Tight money, int rates high	good/better		increasing/high
Profits low, falling	bad/worse		
Profits high, too high	good/better		
Stock market decline	bad/worse		
Other price/credit references			
Balance of payments, dollar devalue			
Controls (price or wage)			
Bad race relations; more crime			
Excessive wage demands by unions; labor unrest			
Times are bad now and won't change in next year	bad/worse		
Good times can't last, due for a fall	bad/worse		
R sees downward trends, has heard business is bad	bad/worse		
Decline in specific industries	bad/worse		
Farm situation is bad, low farm prices, drought	bad/worse		
Other unfavorable/bad factors			
Economy in general less stable, lack of confidence	bad/worse		
Energy crisis, pollution, less natural resources			
Change mentioned but NA whether favorable/unfavorable			

Notes: This classification is provided by the MSC when downloading the data from `https://data.sca.isr.umich.edu/sda-public/cgi-bin/hsda?harcsda+sca`.

B3 Controlling for other economic news

As discussed at the end of Subsection 2.3.1, other economic news releases may correlate with the monetary policy surprises. Therefore, I also control for these other news. Following Del Negro et al. (2012), I use surprises measured as the difference between a survey forecast and the actual release. In particular, I use the Reuters Poll surprises available at ThompsonReuters Datastream. Unfortunately, most series start later than 1994. I therefore, have to limit the sample for this analysis. In order to ensure comparability, I also report the baseline estimation results for the sample when Reuters data is available.

The news surprises I include are measured as the difference in the median forecast from the Reuters Poll and the actual release. I consider the following monthly variables: industrial production, capacity utilization, the unemployment rate, and the manufacturing purchasing manager index (PMI). These series are available since June 1999. Given the lack of information about the timing of the MSC interviews, both news from the previous months and the current month may be relevant. Therefore, I control for both.

Table B.9 shows the results. The first two columns focus on news surprises from the current month. In the sample from June 1999 to December 2016, the coefficient is larger than in the baseline estimation (2.14 compared to 1.00). However, including the other economic news only changes the coefficient slightly to 2.21. The same is true for the results with news from the previous month (columns (3) and (4)), and news from both months (columns (5) and (6)). In all cases, the inclusion of the news measures only changes the effect of the monetary policy surprises slightly, and actually increases it. Furthermore, most of the other economic news do not have a significant effect on inflation expectations. The difference between the coefficients in these samples and the one from the baseline sample highlights the state dependence of the quantitative effect. Depending on which years are included, the effect is smaller or larger. However, the effect is always positive, indicating the presence of information effects in all samples.

Table B.9: Role of other economic news releases

	Dependent var.: inflation expectations in *t*					
	News, month *t*		News, month *t-1*		News, both	
	(1)	(2)	(3)	(4)	(5)	(6)
FF surprise, *t-1*	2.14***	2.21***	2.73***	2.80***	2.66***	2.98***
	(0.66)	(0.67)	(0.67)	(0.68)	(0.67)	(0.69)
Industrial production, surprise, *t*		0.11				0.00
		(0.08)				(0.08)
Capacity utilization, surprise, *t*		-0.18*				-0.08
		(0.10)				(0.10)
Unemployment rate, surprise, *t*		0.40***				0.59***
		(0.13)				(0.14)
Manufacturing PMI, surprise, *t*		-0.00				0.00
		(0.01)				(0.01)
Industrial production, surprise, *t-1*				-0.18**		-0.24***
				(0.08)		(0.08)
Capacity utilization, surprise, *t-1*				0.02		0.08
				(0.09)		(0.10)
Unemployment rate, surprise, *t-1*				0.16		0.17
				(0.14)		(0.14)
Manufacturing PMI, surprise, *t-1*				-0.00		0.00
				(0.01)		(0.01)
Gas price inflation, *t*	0.04***	0.04***	0.04***	0.03***	0.04***	0.03***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Food price inflation, *t*	1.02***	1.08***	1.12***	1.12***	1.15***	1.19***
	(0.09)	(0.09)	(0.09)	(0.09)	(0.09)	(0.09)
Observations	60386	60386	59668	59668	56724	56724
Households	30193	30193	29834	29834	28362	28362
Within R^2	0.01	0.01	0.01	0.01	0.01	0.02

Notes: FF surprise: surprise change in the 3-month federal funds future around FOMC meetings. Inflation expectations are from the Michigan Survey of Consumers, and refer to the next 12 months. Realized inflation is measured as the month-on-month change in the price level for food and gasoline products, respectively. News surprises are measured as the difference between the median expectation in the Reuters poll and the actual release of the respective variable. Data available since June 1999, and taken from Datastream. Columns (1), (3) and (5) show the baseline regression for the sample in which Reuters news data is available to make results comparable. Cluster-adjusted robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

B4 Classification of newspaper articles

The method of text analysis I use is called dictionary or lexicon method. Since I am looking for very specific topic in the articles, I use my own list of keywords. In addition, I make use of the grammatical structure of the articles by only considering sentences. This is more specific than the simple bag-of-words method often used, where classification is simply based on all words which occur within a text, independent of their order or distance (see Shapiro et al. 2019). Below I provide the list of keywords I use. In addition, I describe the detailed adjustments I make based the order of words as well as the distance of words. The last section discusses some checks I carried out.

B4.1 List of keywords

Table B.10: Keywords used to classify articles

	high/increasing		low/decreasing	
all positions	accelerate	quickening	decline	falls
	accelerating	rampant	declined	fell
	elevated	rapid	declining	low
	fast	rise	decrease	lower
	faster	risen	decreased	lowered
	galloping	rises	decreases	lowest
	heightened	rising	decreasing	slow
	high	rose	down	slower
	higher	soaring	downward	slowest
	highest	spur	drop	slowing
	increase	spurring	dropped	sluggish
	increased	strong	dropping	subdued
	increases	stronger	drops	waning
	increasing	up	fall	weak
	quicken	upward	fallen	weaker
			falling	
only next to 'inflation'	more	raises	less	
	much	raising	little	
	raise	raised		

B4.2 Adjustments

After locating the word 'inflation' in each sentence, I separately check the five words before and after inflation for the keywords referring to the two categories, see Table B.10. In order to ensure the classification is correct, I adjust these raw counts in several ways:

1. Some words are only considered if they are directly next to 'inflation', see bottom part of Table B.10. This is done because these words are likely to carry a different meaning or to be linked to a different word if they are further away from 'inflation'.

2. I remove any misleading compound words and phrases including 'inflation' from the text, such as 'inflation-adjusted', 'adjusted for inflation', or 'inflation hawk'.

3. I remove misleading expressions including the keywords, such as 'down the road'.

4. I remove comparative statements, such as "higher than".

5. I exclude several cases with conjunctions and negations. For all cases where a negation or conjunction occurs within the five-word range, I checked all cases manually and chose how to treat them based on how much misclassification they introduce. Here I took a conservative approach in the sense that I valued cases introducing false positives more strongly than cases introducing false negatives. This leads to the exclusion of the following cases:

   i. negation precedes keyword

   ii. negation occurs between keyword and 'inflation'

   iii. 'and', 'or' or 'by' occur between keyword and 'inflation'

   iv. 'up' keyword followed by 'or' (this is usually a phrase like "inflation will be higher or lower", it is common not to say "lower or higher")

   v. keyword followed by 'by'+'percent'/'percentage' (this usually is a phrase like "inflation rose by 1.3%", which does not necessarily imply high or increasing inflation)

6. In a few cases keywords from both categories occurred around the same instance of 'inflation'. For these I chose the classification manually.

B4.3 Manual checks of results

In order to check how well my dictionary performs, I drew a random sample of 300 sentences from all sentences, which were classified and published on the day of the meeting or within 7 days after the meeting. Then, I manually checked the classification: 10.9% of the 176 sentences classified to refer to high or increasing inflation were incorrect, i.e. either did not refer to any direction for inflation or the opposite (although the majority referred to no clear direction). For the 135 sentences classified to refer to

low or decreasing inflation, 11.9% were incorrect. In total 11.7% of all 300 sentences were incorrectly classified.

I also considered a random sample of 300 sentences including the word inflation but not classified to refer to high/increasing or low/decreasing inflation. Of these sentences 31 or 10.3% were classified wrongly, 13 should have been classified to refer to high/increasing inflation, and 18 should have been classified to refer to low/decreasing inflation. The differences in both error rates between the two categories may also partially explain the asymmetric response to monetary policy surprises: the measure of 'down' reporting may be more imprecise than the 'up' measure.

Chapter 3

Firm Expectations and Economic Activity*†

*This chapter is based on joint work with Zeno Enders (University of Heidelberg) and Gernot Müller (University of Tübingen). An earlier version of this chapter is available as CESifo Working Paper No. 7623.

†We thank participants of various conferences and seminars for useful comments. We also gratefully acknowledge valuable suggestions by Nikolai Histrov in the early stage of the project, and the technical support by the team of the LMU-ifo Economics & Business Data Center (EBDC) in Munich and Mark Brisske. This research has received financial support by the German Science Foundation (DFG) under the Priority Program 1859.

3.1 Introduction

To what extent do firms' expectations affect current decision making? According to theory, expectations should have a first-order effect. Expectations about the business cycle take center stage in modern macroeconomic theory: firms decide on production, investment, and hiring as well as on prices in a forward-looking manner (see e.g., Kydland and Prescott 1982; Lucas 1973; Mortensen and Pissarides 2009; Woodford 2003). This, in turn, is essential for why and how cyclical impulses propagate and how policy announcements shape economic outcomes (see e.g., Del Negro et al. 2012; Eggertsson and Woodford 2003). Yet at an empirical level, the systematic exploration of how expectations affect economic decisions and hence economic outcomes is still in its infancy. Arguably, two major difficulties are to blame. First, expectations are not directly observable. Second, expectations are responsive to changes in the economic environment – identifying a causal effect of expectations on economic decisions is therefore challenging.

In this study, we take up the issue by exploiting a particular data set and a novel identification strategy. Specifically, we base our analysis on the EBDC Business Expectations Panel (BEP), maintained by the LMU-ifo Economics & Business Data Center (EBDC) in Munich. Our sample comprises monthly observations for the period 1991 to 2016. In each month, about 2000 German firms in the manufacturing sector report their expectations regarding future production in a qualitative manner: it may increase, not change, or decrease. Similarly, firms report expectations about business cycle conditions. The survey is the basis for the ifo business climate index, a widely observed leading indicator for economic activity in Germany (Becker and Wohlrabe 2008). In addition, the BEP contains a rich set of observations for each firm. These include a large range of measures that capture the economic and financial conditions under which firms operate.

We exploit these data in order to identify the causal effect of firm expectations on their behavior, notably in terms of production and price setting. For this purpose, we match firms based on fundamentals and compare price-setting and production decisions of firms that have the same fundamentals but differ in their views about the future. Formally, we estimate a probit model and match optimistic and pessimistic firms, in turn, with neutral firms using their propensity scores (Rosenbaum and Rubin 1983). Intuitively, we consider "optimism" and "pessimism" as a treatment that is randomly assigned across firms with the same fundamentals: we estimate the average treatment effect on the treated by comparing the behavior of treated and untreated firms with the same probability of being treated.

We find that expectations have a significant effect on production and prices. In the

impact period, optimistic firms are 15 percentage points more likely to raise production than neutral firms. Similarly, we also find that optimistic firms are more likely to raise prices. For pessimistic firms, we find an opposite effect of about the same magnitude: they are more likely to reduce production and prices.

These results are consistent with two distinct hypotheses on how expectations affect economic decision-making. Under the first hypothesis, expectations that are orthogonal to current fundamentals are not necessarily orthogonal to future fundamentals. Put differently, expectations represent genuine information ("news") about the future, which is not yet reflected in current fundamentals. Under this interpretation, expectations matter as a transmission channel, but not as an exogenous source of variation. A number of influential contributions suggest that news are indeed an important source of business cycle fluctuations (Barsky and Sims 2012; Beaudry and Portier 2006; Schmitt-Grohé and Uribe 2012). Yet, these studies provide only indirect evidence on the role of expectations as news. In contrast to our analysis, they do not analyze expectations data explicitly.

Under the second hypothesis, changes in expectations are fully exogenous and different labels are used to capture this notion, such as "noise," "sentiment," or "animal spirits".[1] In this spirit, a number of recent contributions have put forward modern models of the business cycle in which "noise shocks" play a key role (Angeletos and La'O 2013; Lorenzoni 2009). Again, these contributions also do not exploit expectations data directly. Instead, they show that noise helps quantitative business cycle models to account for key features of aggregate time-series data.

The unique nature of our data set allows us to test these two hypotheses directly. For not only do we observe firm expectations regarding future production and business conditions, we also observe actual production and business conditions. We are thus able to construct a measure of firms' forecast errors and identify firms whose optimism or pessimism turns out to be incorrect or "undue" from an ex-post point of view (Pigou 1927). In the second step of our analysis, we match, in turn, incorrectly optimistic and pessimistic firms with ex-ante neutral firms. We find that incorrectly optimistic firms are also relatively more likely to raise output and prices.

In a third step, we quantify the contribution of incorrect optimism and pessimism to aggregate fluctuations. For this purpose, we compute an aggregate measure of incorrect optimism and pessimism in our population of firms. Specifically, we use an ordered probit model to measure the extent of optimism and pessimism at the firm level and classify such sentiment as incorrect whenever we observe a forecast error ex-post. Finally, we aggregate across firms and project macro variables of interest on the

[1] According to Keynes, animal spirits are "a spontaneous urge to action rather than inaction", which drive economic decisions beyond considerations based "on nothing but a mathematical expectation" (Keynes 1936, pp. 161–162).

resulting time series of incorrect optimism and pessimism. We find that optimism in particular causes industrial production and prices to rise.

Our analysis relates to studies that focus on the expectation formation process. Coibion and Gorodnichenko (2012, 2015a) document the presence of information rigidities using survey data from different sources. More recently, Coibion et al. (2018) exploit a survey of New Zealand firms. There is also work on expectation formation with the data from the ifo survey. An early study by Nerlove (1983) finds evidence in support of an adaptive expectations model. Bachmann and Elstner (2015) show that at most one-third of the firms in the ifo survey make systematic forecast errors. Massenot and Pettinicchi (2018), in turn, identify various factors which account for forecasting errors of firms in the ifo sample. In Enders et al. (2019b), we show that firm expectations respond systematically to monetary policy announcements.

Few studies investigate empirically how expectations affect economic decision making. An exception is Boneva et al. (2020). They find that expectations of UK firms feature significantly in an estimated version of the New Keynesian Phillips curve. Tanaka, Bloom, David, and Koga (2020) analyze GDP forecasts of Japanese firms and show an association with employment, investment, and output growth at the firm level. Coibion et al. (2020) use a survey of Italian firms to estimate the effect of decision makers' inflation expectations on firm decisions. Bachmann and Zorn (2018) use ifo data to study the drivers of investment and find, among other things, a role for firm expectations. Gennaioli, Ma, and Shleifer (2015) analyze the Duke University quarterly survey of Chief Financial Officers and show that CFOs' expectations of earnings growth explain firm investment.

The remainder of the chapter is structured as follows. The next section provides details on our data set as well as a number of descriptive statistics. Section 3.3 describes the estimation approach and the results of the first step of our analysis. In Section 3.4, we zoom in on the transmission channels of firm expectations and distinguish between firms with and without forecast errors. Afterwards we quantify the aggregate effects of firm expectations using local projections. Section 3.6 concludes.

3.2 Data

The EBDC Business Expectations Panel (BEP) combines monthly survey data from the ifo institute and annual balance sheet data from the Amadeus and Hoppenstedt databases (EBDC-BEP 2017). The data combines different surveys covering German firms in four sectors: manufacturing, retail, construction, and services. The surveys include the same basic stock of questions for each sector, but the wording of these questions and answers may differ at times. In our analysis, we focus on the manufacturing

survey, which is the longest running and includes the largest number of firms. Also, in this case the wording of the survey questions is particularly suitable for the purpose of our investigation. One caveat to note is that the responses to the survey and the balance sheet data come at different frequencies: while the survey is conducted monthly, balance sheet data is only available annually. We will use balance sheet data to predict firm expectations. To ensure that we do not use information that is not yet available when firms report expectations, we only use the most recent balance sheet data at a given point in time.[2]

The BEP sample period starts in 1986 with the manufacturing survey. In our analysis, we use data from 1991 to 2016 because some of the variables we rely on are available only since 1991. The unit of observation in the manufacturing survey is a product. As a result, some firms respond to several questionnaires each month or different plants of one firm respond separately. However, in our sample this is the case for less than 10% of firms. We conduct our analysis at the product/plant level and do not explicitly account for whether a product/plant is part of a multi-product firm. In our analysis below, we refer to the individual observation as a "firm" in order to ease the exposition.

We report a number of basic statistics for the firms in our sample in Appendix C1. For this purpose, we distinguish between the sample of all firms ("full sample") and the sample for which balance sheet data is available. The latter sample is smaller but still includes about different 5,000 firms and more than 300,000 firm-month observations, see Table C.1. Table C.2 shows that means and standard deviations of firms' responses are generally quite similar across the samples, although firms for which balance sheet data are available tend to be somewhat larger. Firms with balance sheet data stay in the survey for 87 months (7 years) on average and provide answers in 74 months (6 years), implying that they respond in 83% of the months they are in the sample. We also stress that independently of the starting period sample attrition is moderate, as shown in Table C.3.

The BEP covers a large set of questions, but only a subset of those are asked regularly. Within firms, the questions are typically answered by the top management. In more than 80% of small and medium-sized firms and more than 60% of large firms, the CEO or owner responds. Otherwise, the response is typically provided by the head of the relevant department (for details, see Sauer and Wohlrabe 2019). In our analysis, we focus on four main questions, listed in Table 3.1. Some questions vary over time. Es-

[2]For example, if a firm publishes balance sheet data every September, we will use this data for all the following months until the next balance sheet is published. Hence, our specification is conservative as we neglect potential information known to firms in the months close to but preceding the publication of the balance sheet. In Appendix C4.1, we pursue an alternative strategy but find the results obtained for the baseline specification robust.

Table 3.1: Selected ifo survey questions

Label	Question[1]	Possible answers
Q1	Expectations for the next 3 months: Our domestic production activity regarding good XY will probably ...	increase [1] not change [0] decrease [-1]
Q2	Expectations for the next 6 months: Taking economic fluctuations into account our state of business will be ...	rather more favorable [1] not changing [0] rather less favorable [-1]
Q3	Tendencies in the previous month: Our domestic production activities with respect to product XY have ...	increased [1] not changed [0] decreased [-1]
Q4	Tendencies in the previous month: Taking changes of terms and conditions into account, our domestic sales prices (net) for product XY have ...	increased [1] not changed [0] decreased [-1]

Notes: Tables shows our translation of the most recent formulation of the question in the German questionnaire. Additional questions used are listed in Table C.6. Changes to the questions are listed in Table C.7

pecially in 2002, many changes were implemented due to a harmonization of business and consumer surveys in the European Union. The changes relevant to our analysis are documented below.

Our measure of firm optimism is based on question Q1, which refers to expectations about production activity in the next three months. The wording of this question has changed over time. Since July 1994, firms can additionally report that they have no significant domestic production. These firms are not included in our analysis. Furthermore, the question contained a note to ignore seasonal fluctuations until the end of 2001. Since these are minor changes affecting all firms in the same way, they are unlikely to matter for our results. Further details on the wording of the questions can be found in Table C.7 in the appendix.

Q2 is a broader question regarding expectations for the state of business over the next six months. Combined with a question on the current state of business it provides the basis for the ifo business climate index. In a sensitivity analysis, we consider optimism as reflected in the answer to this question and find similar results as in our baseline. Furthermore, the answers to both questions tend to be highly correlated, see Figure C.1 in the appendix. In our baseline analysis, we use Q1, though, because its wording is more specific and the time horizon in question is shorter. Q2 also used to include an additional note to ignore seasonal fluctuations. This note was dropped in 1997, see Table C.7.

Questions Q3 and Q4 refer to our outcome variables: changes in production and prices. These questions changed in 2002. Previously both questions asked about the change in production and prices in the current month *compared to* the previous month.

Figure 3.1: Distribution of production expectations, 1991 to 2016

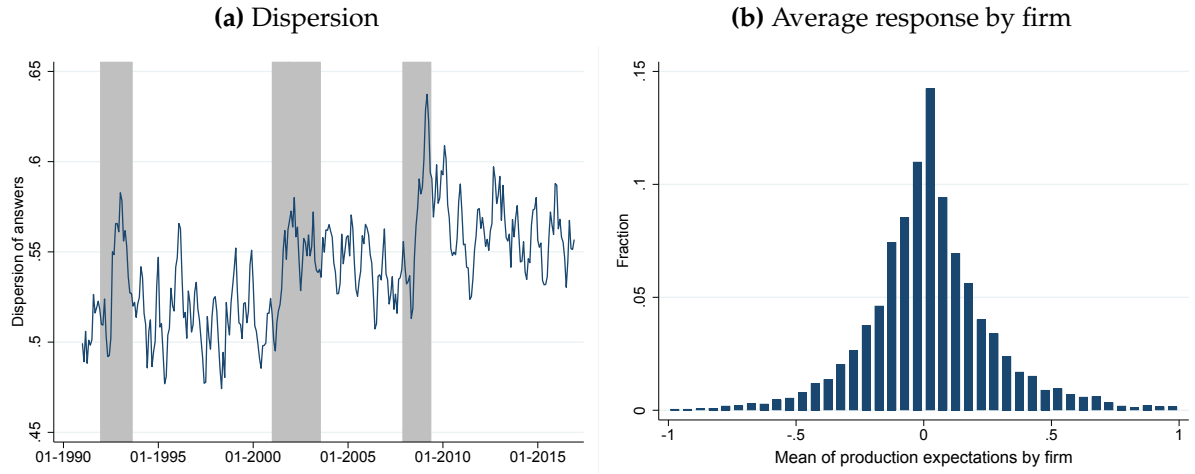

Notes: Full sample. Shaded areas in Panel (a) indicate recession periods as defined by the German Council of Economic Experts. Panel (b) only includes firms which respond at least 10 times.

Since 2002 both questions ask about the change in these variables *in* the previous month. We adjust the data to account for this change in timing. To make sure that this adjustment does not affect results, we also consider a reduced sample, which starts in 2002 in our sensitivity analysis. The results are very similar to those for the full sample (see Section 3.3).

In what follows, we compute a number of descriptive statistics for the variables of interest. For the purpose of our quantitative analysis, we assign a value of 1 to positive responses (*increase/improve*) and a value of -1 to negative responses (*decrease/worsen*) and a value of 0 otherwise. Figure 3.1 shows the distribution of responses over time and across firms. Because distributional statistics are not straightforward to compute with qualitative data, we report results from two alternative approaches. In Panel (a) of Figure 3.1, we plot a common measure of dispersion based on the shares of positive and negative responses in a given month.[3] Dispersion generally increases during crisis periods and tends to decrease afterwards. In Panel (b) of Figure 3.1, we plot the distribution of the average response to Q1 per firm. It shows that the share of notorious optimists and pessimists is limited, and that for most firms the average response is zero. The distribution is close to normal.

Figure 3.2 plots the cross-correlation function for expected production and six leads and lags of realized production and prices. The correlation between expectations and reported outcomes is positive in both cases, across all leads and lags (Panels (a) and (b)). The contemporaneous correlation of current production and prices, on the one

[3]Dispersion of expectations based on qualitative survey data is measured as $\sqrt{\text{frac}^+ + \text{frac}^- - (\text{frac}^+ - \text{frac}^-)^2}$, where frac^+ and frac^- are the fraction of positive and negative responses in each month, respectively. This measure is also used by Bachmann et al. (2013).

Figure 3.2: Correlation of expected changes in production with changes in realized production and prices in the manufacturing sector

(a) Expected production in t and reported production $t + k$, firm level

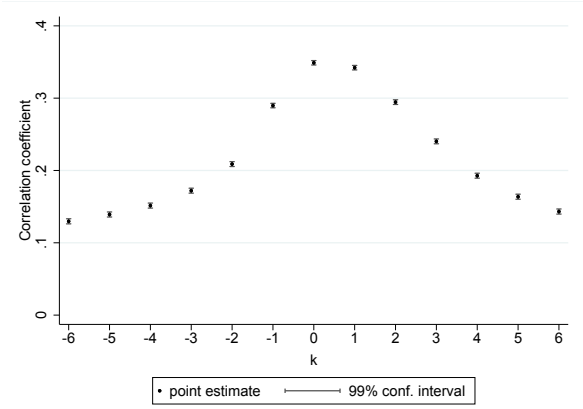

(b) Expected production in t and reported prices $t + k$, firm level

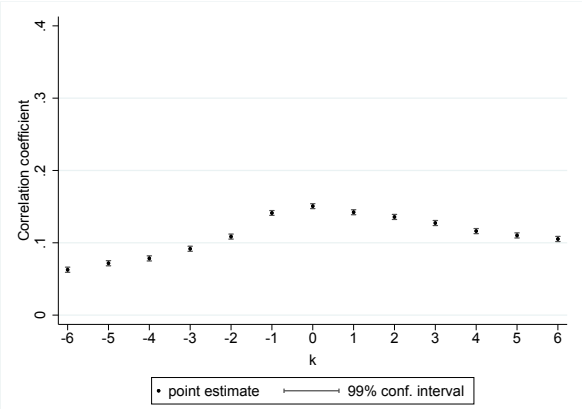

(c) Mean expected production in t and industrial production in $t + k$, monthly

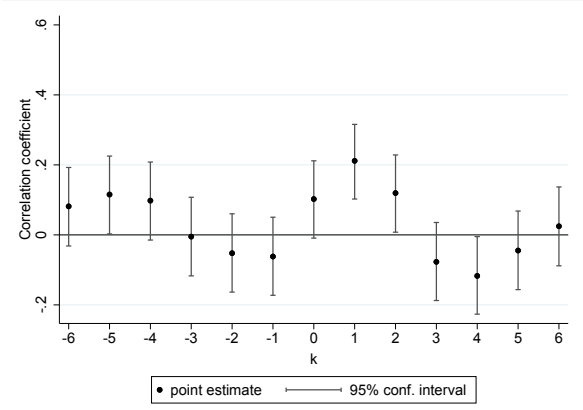

(d) Mean expected production in t and producer prices in $t + k$, monthly

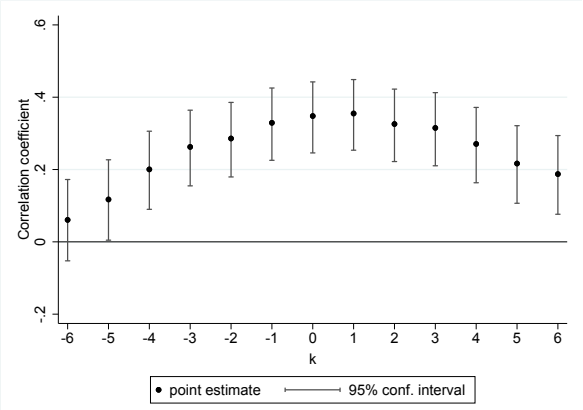

Notes: Full sample; indices of production and prices in manufacturing from the German Statistical Office and the Bundesbank, respectively. Industrial production is measured as the month-on-month change in the not seasonally adjusted index. Producer prices are measured as month-on-month changes in the calendar adjusted price index.

hand, and production expectations, on the other hand, is particularly strong. In our analysis below, we seek to establish the causal effect of production expectations on production and price-setting decisions.

We also compute the correlation of average firm expectations and aggregate production and prices. Panel (c) of Figure 3.2 shows the cross correlation function for average production expectations across firms within a month, and leads and lags of the monthly growth rate of industrial production. The two time series are strongly correlated for small leads of industrial production. This is consistent with the well-established fact that the ifo business climate index is a leading indicator of economic activity in Germany (Abberger and Wohlrabe 2006; Henzel and Rast 2013). We com-

pute the same statistic for month-on-month changes in the producer price index and average production expectations and show results in Panel (d). Here the correlation is very strong for more lags and leads.

3.3 Do firm expectations matter?

The main purpose of our analysis is to identify the effect of firm expectations on firm decisions. Specifically, in our baseline specification, we assess to what extent firms' production and price-setting decisions depend on production expectations. To this end, we compare the behavior of firms that expect an increase (decrease) of production to firms that expect production to remain unchanged. A key challenge in this regard is to identify variation in expectations that is orthogonal to current fundamentals. For only to the extent that firms are comparable in terms of fundamentals, we may think of expectations as a "treatment" into which some firms are randomly selected and others are not.

Put differently, as we compare the behavior of firms with different views about the future, we face a selection problem because firms with better fundamentals are also more likely to enjoy a more favorable outlook. In order to address this selection problem, we rely on propensity score matching (see e.g., Caliendo and Kopeinig 2008; Imbens and Rubin 2015). The idea is to mimic randomized control trials where treatment is actually assigned in a random fashion and hence orthogonal to observable characteristics. The matching approach is particularly suited for the purpose of our analysis since we are dealing with qualitative data on expectations: firms may be either optimistic, neutral, or pessimistic, i.e. expect an increase, no change, or a decrease, respectively. Hence, in our analysis, to the extent that firms receive a treatment, they are treated either with "optimism" or with "pessimism." Of course, our analysis does not require optimism or pessimism to be literally assigned in a random way. We merely assume that the assignment is orthogonal to current fundamentals. Note also that we do not require optimism/pessimism to be unrelated to future fundamentals. We take up this issue in more detail in the Section 3.4.

In general, the matching approach offers several advantages over conventional regression analysis. First, it ensures that the distribution of control variables is similar across treated units and the control group (Dehejia and Wahba 2002; Imbens and Rubin 2015). This is important because differences in the distribution of controls can lead to a significant bias when estimating treatment effects (Heckman, Ichimura, Smith, and Todd 1998). Second, the matching approach disciplines the analysis since the control group is specified prior to and independently of the estimation of the treatment effect (Imbens and Rubin 2015). Lastly, after matching, the treatment effect is estimated by

a simple mean difference, thus allowing for a non-parametric estimation (Dehejia and Wahba 1999; Heckman et al. 1998).

3.3.1 Propensity score matching

We now briefly outline our approach following Caliendo and Kopeinig (2008). Inference is based on estimating the potential outcome of a treated firm under no treatment, that is, the (unobserved) counterfactual outcome had the treated firm not been treated. Formally, the object of interest is the average treatment effect on treated (ATT) firms:

$$\theta = \mathbb{E}[Y(1) - Y(0)|D = 1] = \mathbb{E}[Y(1)|D = 1] - \mathbb{E}[Y(0)|D = 1],$$

where $D = 1$ indicates treatment, $Y(1)$ the potential outcome of a treated firm, that is, a firm which is optimistic (pessimistic), and $Y(0)$ the potential outcome in the absence of treatment. Since we do not observe the latter for treated firms, we can only estimate the following relationship:

$$\mathbb{E}[Y(1)|D = 1] - \mathbb{E}[Y(0)|D = 0] = \theta + \mathbb{E}[Y(0)|D = 1] - \mathbb{E}[Y(0)|D = 0]. \tag{3.1}$$

This is equivalent to the ATT only if

$$\mathbb{E}[Y(0)|D = 1] - \mathbb{E}[Y(0)|D = 0] = 0,$$

that is, the potential outcomes are independent of treatment assignment. In randomized control trials, this holds true due to the random assignment of treatment. In observational studies, additional assumptions are required. One approach is to assume that treatment is assigned randomly given a set of relevant covariates X:

$$Y(1), Y(0) \perp D|X.$$

Covariates are relevant if they affect both the (potential) outcome and the probability of being treated. In our case, this means that we need to include all information that matters for firms' expectation formation as well as for their production and price-setting decisions. We describe these variables below. Since we are only interested in the effect on the treated, we merely need $Y(0)$ to be independent of treatment status, see equation (3.1). In this case, the required conditional independence assumption simplifies to

$$Y(0) \perp D|X.$$

In the expressions above, we condition on the whole set of control variables. This

Table 3.2: Control variables in the propensity score model

Variable	Description	Frequency	Reference period
debt share[1]	total debt over assets	annual	*t-11* to *t*
financing coefficient[1]	liabilities minus provisions divided by equity plus provisions	annual	*t-11* to *t*
employees	no. of employees	annual[2]	October/November
state of business	answer to question on state of business (values: 1, 0, −1)	monthly	*t*
orders	answer to question on state of orders (values: 1, 0, −1)	monthly	*t*
foreign orders	answer to question on state of foreign orders (values: 1, 0, −1)	monthly	*t*
production	answer to question on change in production (values: 1, 0, −1)	monthly	*t-1*
prices	answer to question on change in prices (values: 1, 0, −1)	monthly	*t-1*
capacity utilization	utilization of existing capacity in %	quarterly[2]	*t-1*
demand	answer to question on demand in previous month (values: 1, 0, −1)	monthly	*t-1*

Notes: For all variables with monthly frequency three lags are also included. In addition various interaction terms are included (based on a log-likelihood ratio test).
[1] To ensure outliers and measurement error do not affect our results, we exclude the 99.99 percentile of observations for the debt share and the 0.02 and 99.98 percentiles for the financing coefficient.
[2] In months with no reporting, we use data from the most recent balance sheet/most recent quarter the question was asked (if available).

can be challenging when the number of observable controls is large. In our analysis, we include 4 continuous variables and 18 categorical variables with 3 outcomes each. If we were to split the sample by the categorical variables only, we would already have 3^{18} potential bins. This makes accounting for controls by creating sub-samples of identical observations infeasible even with a large data set. We therefore rely on a result established by Rosenbaum and Rubin (1983): asymptotically, it is equivalent to condition on the propensity to be treated, $p(X) \equiv Pr(D = 1|X)$, or to condition directly on X. The conditional independence assumption can thus be stated as follows:

$$Y(0) \perp D|p(X).$$

Conditioning on the propensity score requires the additional assumption of common support, that is, treatment is not fully determined:

$$0 < p(X) = Pr(D = 1|X) < 1.$$

In what follows, we estimate the ATT by comparing the outcome of each treated observation to one or several untreated units with the same (or very similar) propen-

sity score. In our analysis, there are two possible treatments: optimism and pessimism. To establish the effect of a treatment, we compare firms in each case to firms which do not expect production to change ("neutral firms"). In order to estimate the propensity score, we pursue two alternative approaches. Since we are dealing with two treatments, we first estimate an ordered probit model where optimism and pessimism are outcomes of a common model. Alternatively, we consider two distinct probit models for optimists and pessimists. In the first case, we estimate the probability of the latent variable, y_{it}^*, falling between two thresholds α_{j-1} and α_j for treatment j as

$$Pr(y_{it} = j) = Pr(a_{j-1} < y_{it}^* \leq \alpha_j) = \Phi(\alpha_j - X_{it}'\beta) - \Phi(\alpha_{j-1} - X_{it}'\beta), \tag{3.2}$$

where $j = \{-1, 0, 1\}$ corresponds to the three possible answers to Q1. We collect the control variables in the vector X_{it}. It includes time and sector fixed effects, the sector average of the reported state of business in each month, three lags of the dependent variables, and all firm specific variables listed in Table 3.2 (including three lags for each of the survey variables). More detailed information on the survey variables is provided in Table C.6 in the appendix.

The ordered probit does not directly yield the propensity score. In this case the propensity score, $p^m(X_{it})$ for treatment $m = \{optimism, pessimism\}$, equals the conditional probability of the treatment given the alternative of no treatment, that is, expecting production to remain unchanged (see again Caliendo and Kopeinig 2008):

$$p^m(X_{it}) = \frac{Pr(y_{it} = m|X_{it})}{Pr(y_{it} = m|X_{it}) + Pr(y_{it} = 0|X_{it})}.$$

The second approach involves two separate probit regressions – one for each treatment. The specification is the same as for the ordered probit model:

$$Pr(D_{it}^m = 1) = Pr(X_{it}'\beta) = \Phi(X_{it}'\beta), \tag{3.3}$$

where D_{it}^m is a dummy variable which is 1 for an observation responding *increase* in the case of the optimism treatment, or *decrease* in the case of the pessimism treatment, and 0 for an observation responding *no change* in both cases. We again collect the same control variables in vector X_{it}. Since the sample only includes the specific treatment group and the untreated, the estimated probability is a direct estimate of the propensity score:

$$p^m(X_{it}) = Pr(D_{it}^m = 1).$$

Caliendo and Kopeinig (2008) discuss the use of serial probit estimation compared to multinomial models in the case of multiple treatments. They argue that, generally,

Figure 3.3: Histogram of the density of the propensity scores

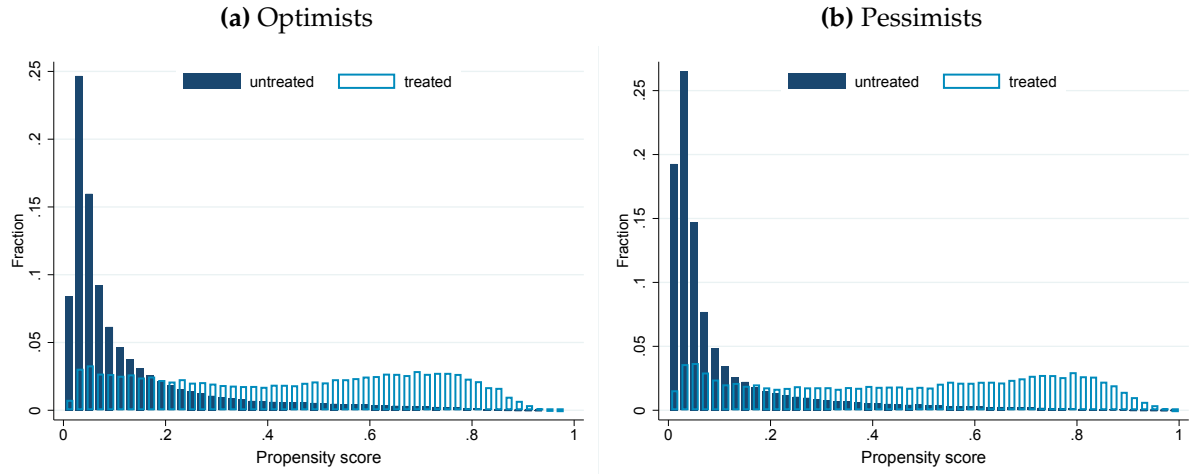

Notes: Histograms show the propensity scores for treated and untreated firms respectively, estimated as described by equation (3.3). In Panel (a) treated firms are optimistic. In Panel (b) treated firms are pessimistic.

authors found no difference or a slight advantage of using separate probit models. It turns out that also in our case the serial probit estimation has a slight advantage as it yields improved balancing statistics. We therefore use it in our baseline. However, results based on the ordered probit do not differ much from results using the two probit regressions, see Subsection 3.3.4.

As mentioned above, we include all potentially relevant variables as controls in order to capture the fundamentals of the firm, both current values as well as lags. However, we only consider realizations, which are available at the time the survey is conducted. In this regard, it is important to note that most firms respond to the survey in the first two weeks of the month. Figure C.2 in the appendix shows the distribution of participation days within the month. 50% answer within the first eight days and another 25% answer in the following week.[4]

After computing the propensity scores, we match treated and untreated observations using a variant of caliper or radius matching (Caliendo and Kopeinig 2008).[5] We match each treated observation i (optimistic or pessimistic) to all untreated observations k (neutral) *within the same month*, which satisfy

$$p(X_{it}) - 0.02 \leq p(X_{kt}) \leq p(X_{it}) + 0.02.$$

Here we allow for a radius of 0.02. This corresponds to about a tenth of the standard

[4]These statistics pertain to firms that answer the survey online. They represent more than 60% of the firms since 2004.

[5]We also test an alternative matching procedure proposed by Lechner, Miquel, and Wunsch (2011). The results are very close to our baseline results. Details can be found in Appendix C4.2.

Table 3.3: Number of matched observations

	Optimism treatment		Pessimism treatment	
	Total	Matched	Total	Matched
Panel (a): All firms				
Treated observations	26 974	25 050	23 327	20 947
Untreated observations	114 843	111 027	114 809	110 625
Panel (b): Correct firms				
Treated observations	12 366	9 995	12 123	9 493
Untreated observations	82 317	73 321	82 519	72 762
Panel (c): Incorrect firms				
Treated observations	10 634	9 671	7 641	6 614
Untreated observations	82 505	76 349	82 497	74 357

Notes: Panel (a) shows results for matching as discussed in this section. Panels (b) and (c) show results for matching based on more detailed treatments as discussed in Section 3.4.

deviation of the estimated propensity score.[6] All untreated observations to which a treated observation is matched are given equal weights: the inverse of the number of untreated observations in each match. Note that the untreated observations can be matched more than once to different treated observations.

Figure 3.3 displays the distribution of the propensity scores. The left panel contrasts the distribution for firms which receive an optimism treatment (light blue, transparent bars) with those for untreated firms (dark blue, solid bars). The right panel reports results for pessimism. In each instance, we find that there is considerable overlap of the distribution (common support), although the mass of untreated firms is more concentrated at lower propensity scores.[7] Panel (a) of Table 3.3 reports basic statistics regarding our matches.[8] We are able to find matches for about 93% of all treated optimists and for 90% of treated pessimists. This is due to the large overlap in propensity scores between treated and untreated firms.

3.3.2 Diagnostics

Before turning to the results, we report some diagnostics of the matching exercise. We compute balancing statistics in order to assess how similar the samples of treated

[6] Alternative values for the radius give similar results or, if not, fail to deliver satisfying balancing statistics (see next sections).

[7] There are also some treated observations with a larger propensity score than the largest propensity score of all untreated observations. We drop these observations in what follows. This trimming ensures that only suitable observations are matched. Specifically, we drop 15 observations for optimism but 0 for pessimism.

[8] For now only consider Panel (a). The other two panels refer to additional results from Section 3.4.

and untreated observations are. The main statistic of interest is the standardized bias between the treated and untreated sample for each control variable. Following Rosenbaum and Rubin (1983), this is computed as follows:

$$SB = 100 \frac{\bar{x}_1 - \bar{x}_0}{\sqrt{0.5(s_1^2 + s_0^2)}}, \tag{3.4}$$

where $\bar{x}_1$ is the mean of the control variable among the treated observations, $\bar{x}_0$ is the mean of the control variable for all untreated observations, s_1 is the standard deviation of the treated observations and s_0 the standard deviation of all untreated observations. Panels (a) and (b) of Figure 3.4 display the bias before and after matching for optimists and pessimists. They show that matching leads to a sizeable reduction of the standardized bias. According to a widely used rule of thumb, the matched sample is regarded as well balanced when all standardized biases are below 5% (Caliendo and Kopeinig 2008).[9] We meet this standard in all instances, see also Table C.8 in the appendix.

Rubin (2001) suggests a second measure of balancing. He argues that the variance of the part of each covariate that is orthogonal to the propensity score (the residual of a regression of the covariate on the propensity score) should be similar for treated and untreated firms. Specifically, the ratio of the variances should not be below 0.5 or above 2. Ratios between a range of 0.8 and 1.25 are considered acceptable. Panels (c) and (d) of Figure 3.4 plot the variance ratios before and after matching. Again, we find that matching firm-month observations ensures that treated and non-treated firms appear well balanced in terms of covariates. Only the ratio for the number of employees and the financing coefficient (for pessimists) falls in the "of concern" area (outside dashed lines).

3.3.3 Computation of the treatment effect

In what follows, we focus on the average treatment effect on the treated (ATT) in terms of production and price-setting decisions. For each outcome variable, we compute the ATT as the mean difference, across all matches, of treated and untreated firms.

The computation of standard errors for estimates of the ATT computed after matching is not straightforward. One can use analytical variances or bootstrapping. Since bootstrapping has sometimes been shown to be invalid (Caliendo and Kopeinig 2008), we use the methodology of Lechner (2001). He shows that in case of variants of nearest

[9] Imbens and Rubin (2015) suggest that 10% can also be considered a satisfactory value, especially when the initial bias is large.

Figure 3.4: Standardized bias and variance ration, before and after matching

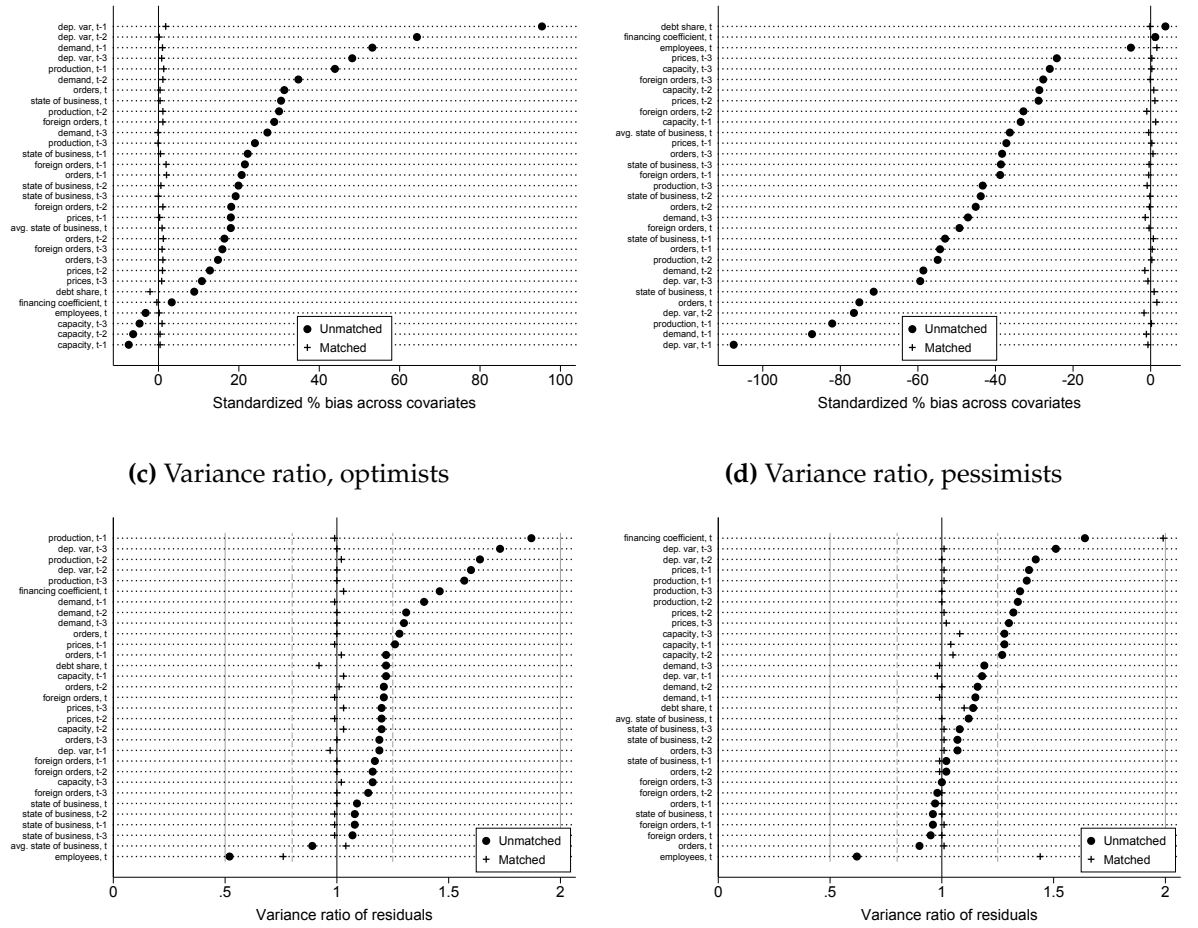

Notes: Figure shows diagnostics statistics for the matching of optimists and pessimists. The standardized bias measures the mean difference of each variable in the treated and untreated groups, as described by equation (3.4). The variance ratio measures the difference between the variances orthogonal to the propensity score. Variance ratios below 0.8 and above 1.25 (dashed lines) are considered "of concern"; ratios below 0.5 and above 2 (gray solid lines) are considered "bad", according to Rubin (2001).

neighbor matching, as in our case, the variance of the ATT, $\hat{\tau}_{ATT}$, is:

$$Var(\hat{\tau}_{ATT}) = \frac{1}{N_1} Var(Y(1)|D=1) + \frac{\sum_{j\in\{D=0\}}(w_j)^2}{(N_1)^2} Var(Y(0)|D=0),$$

where $Y(1)$ and $Y(0)$ refer to a variable of interest given the treatment indicator D equals 1 or 0. N_1 is the number of matched treated firm and w_j is the weight of untreated firm j (see above).

Table 3.4: Average treatment effect on the treated, optimistic and pessimistic firms

	(1) Baseline	(2) Radius 0.01	(3) Sample 2002-2016	(4) Sample excl. fin. crisis[1]	(5) Ordered probit	(6) Exp. state of business
Panel (a): Optimists – Production (change in current month)						
ATT	0.172***	0.170***	0.181***	0.170***	0.152***	0.182***
	(30.43)	(29.34)	(30.03)	(28.52)	(27.14)	(35.02)
Observations	129812	120335	108683	113690	128932	129706
Panel (b): Optimists – Prices (change in current month)						
ATT	0.025***	0.025***	0.027***	0.025***	0.016***	0.020***
	(5.97)	(5.80)	(5.74)	(5.52)	(3.87)	(5.18)
Observations	129858	120367	108715	113734	128977	129759
Panel (c): Pessimists – Production (change in current month)						
ATT	-0.173***	-0.170***	-0.169***	-0.172***	-0.198***	-0.204***
	(-27.77)	(-26.47)	(-24.85)	(-25.37)	(-32.53)	(-35.03)
Observations	125458	113992	104490	106764	123941	125091
Panel (d): Pessimists – Prices (change in current month)						
ATT	-0.031***	-0.033***	-0.029***	-0.035***	-0.038***	-0.029***
	(-6.13)	(-6.41)	(-5.18)	(-6.53)	(-7.80)	(-6.33)
Observations	125530	114050	104551	106821	124014	125169

Notes: Tables shows treatment effects on prices and production for different specifications. T-statistics in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
[1] Excluding the years 2008 and 2009.

3.3.4 Results

We now turn to the question that motivates our analysis: to what extent do firms' expectations affect current decision making? Table 3.4 provides a first answer. In the upper part of the table, we report the ATT of optimism regarding future production, in the lower part we report the ATT of pessimism. In each instance, we focus on production and price-setting decisions. In the columns, we consider alternative specifications.

The left-most column (1) reports results for our baseline. We find a significant positive treatment effect for production (Panel (a)). This positive effect may reflect a stronger tendency among treated firms to raise production or a reduced tendency to lower production, or both. We disentangle these effects below. For prices, we also find a significant positive effect. Although, in this case the effect is much smaller (see Panel (b)). Taken at face value, such an apparently small effect is consistent with the notion that prices are adjusted only infrequently in the short run. However, note that the effect is the outcome either of more frequent upward adjustments or of less frequent downward adjustments of prices among treated firms. Last, we note that the

Table 3.5: Average treatment effect on the treated, increases and decreases in production and prices

	(1) Prod. increase	(2) Prod. decrease	(3) Price increase	(4) Price decrease
Panel (a): Optimists				
ATT	0.149***	-0.022***	0.018***	-0.008***
	(36.93)	(-6.86)	(5.35)	(-3.20)
Observations	129812	129812	129858	129858
Panel (b): Pessimists				
ATT	-0.024***	0.149***	-0.005	0.025***
	(-7.11)	(31.35)	(-1.63)	(7.23)
Observations	125458	125458	125530	125530

Notes: Table shows treatment effects for binarized production and price indicators, i.e. separately considering increases and decreases. T-statistics in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

effect on production and prices is quite symmetric across optimism and pessimism, even though we estimate separate models for optimists and pessimists (see Panels (c) and (d) of Table 3.4).

Table 3.4 also reports results for alternative specifications. In column (2), we show results for a smaller radius in the matching procedure (0.01 instead of 0.02). In column (3), we consider a shorter sample period. It starts in 2002 rather than in 1991. In column (4), we report results for a sample which excludes observations from the financial crisis, that is, the years 2008 and 2009. Column (5) shows results for the case when we use an ordered probit model to estimate the propensity scores. Column (6) refers to the case when we use expectations for the future state of business rather than future production to define the treatments. Recall that in this case the "forecasting horizon" is 6 months rather than only 3 months. We observe that in all instances the estimate of the ATT is close to that for the baseline and always significant. Balancing statistics for the three sensitivity specifications that require changes in the matching procedure are summarized in Table C.9 and Table C.10 in the appendix.

As noted above, our results regarding the response of production and prices may reflect more upward adjustments or fewer downward adjustments, or both. In order to disentangle the overall effect, we transform the dependent variable such that we obtain two binary variables for, in turn, production and prices, a frequently used approach when dealing with qualitative survey data. We then compute the probability of treated firms to raise (lower) prices or production as the mean difference in the newly defined variable across treated and non-treated firms.

Table 3.5 shows the results for the baseline specification. Columns (1) and (2) show

that optimistic firms stand out by their increased probability of raising production. Specifically, the probability of a production increase is 14.9 percentage points higher for optimistic firms than for untreated firms. This accounts for the bulk of the overall effect discussed above. The probability of a production decrease, in turn, falls by 2.2 percentage points. Optimism also raises the probability of a price increase by 1.8 percentage points and lowers the probability of a price decrease by 0.8 percentage points, see columns (3) and (4) of Table 3.5.

Likewise, pessimism increases the probability of a cut of current production by 14.9 percentage points, while the probability of a production increase falls by 2.4 percentage points. The response of prices to pessimism is somewhat larger than the one to optimism. The probability of a price decline increases by 2.5 percentage points. The probability of a price increase is not affected significantly.

3.4 News or noise?

In the previous section, we established that firm expectations affect decisions on current production and pricing at the firm level. This raises the question of why this is the case. Two alternative hypotheses appear plausible. According to the first hypothesis, firms may have information about the future developments that is unrelated to current fundamentals. While our set of fundamentals includes forward-looking variables such as orders, one cannot rule out the possibility that firms have additional information beyond what is already reflected in current fundamentals. According to this "news" hypothesis, firms therefore have a good reason to be either optimistic or pessimistic. It is only that these reasons are not yet observable to the econometrician. Instead, according to the second hypothesis, optimism and pessimism are just "noise", that is, misperceptions about the future that are fundamentally unwarranted or "undue" (Pigou 1927). Of course, our estimate of the ATT may also reflect a mixture of news and noise.

3.4.1 Production and prices

In what follows, we seek to determine to what extent the expectations, which govern current decisions about production and prices, reflect news and noise. We do so on the basis of firms' forecast errors. Intuitively, if a firm appears particularly optimistic relative to its current fundamentals, but reports later that actual production is unchanged or declined, its view about the future appears—with the benefit of hindsight—to have been misperceived. We are thus able to classify optimism and pessimism as incorrect from an ex-post perspective. Note that we do not take a stand on whether expectations

Table 3.6: Classification of correct/incorrect firms

Expectation	Realization	Classification
expected *increase* in t	realization in t+1 to t+3 > 0	correct
expected *increase* in t	realization in t+1 to t+3 ≤ 0	incorrect
expected *no change* in t	realization in t+1 to t+3 $> \frac{1}{3}$	incorrect
expected *no change* in t	$-\frac{1}{3} \geq$ realization in t+1 to t+3 $\leq \frac{1}{3}$	correct
expected *no change* in t	realization in t+1 to t+3 $< -\frac{1}{3}$	incorrect
expected *decrease* in t	realization in t+1 to t+3 ≥ 0	incorrect
expected *decrease* in t	realization in t+1 to t+3 < 0	correct

Notes: Table shows how we classify firms into correct and incorrect based on the reported expectation and the reported realization. The latter is the simple average of the firms' responses to the question on realized production in periods t+1, t+2 and t+3 (Q3).

were rational or not from an ex-ante point of view. We simply compare them to actual outcomes. For instance, a firm may have been optimistic about some aspect of the future and correctly so. Actual production may still fall short of the expected level because of some other unforeseen development. Since our concept of optimism pertains to future production rather than to specific events or fundamentals, we classify such firms as incorrect optimists.

In order to assess whether a firm is incorrectly optimistic or pessimistic, we follow Bachmann et al. (2013) and interpret the qualitative responses to questions about expected and realized production (Q1 and Q3, respectively) as pertaining to the same latent variable. In a nutshell, we classify firms that expect a change as incorrect whenever the average realization in the three months has a different sign than expected. Firms expecting no change are considered to have made no error if they report at most one change in either direction or two changes in opposite directions. Table 3.6 provides an overview of our classification scheme.

Based on this classification scheme, we define a treatment with "correct optimism." It refers to optimistic firms (answer "increase" to Q1) that are classified as correct. The control group are all neutral firms (answer "not change" to Q1) which are classified as correct. The second treatment we consider is "incorrect optimism." Here we consider firms that were optimistic but incorrect from an ex-post point of view. The control group is the same as in the first case. The third and fourth treatments are defined analogously for pessimists. Using these four new treatment indicators, we perform the same matching procedure as described in Subsection 3.3.1.

Before turning to the results, we again consider some diagnostic statistics to ensure that the matching works reasonably well in this case as well. The statistics are the same as described in Subsection 3.3.2. Panels (b) and (c) in Table 3.3 above report the number

Table 3.7: Average treatment effect on the treated, correct and incorrect firms

	(1) Baseline	(2) Radius 0.01	(3) 2002-2016	(4) Sample ex. crisis[1]
Panel (a): Correct optimists – Production (change in current month)				
ATT	0.302***	0.298***	0.314***	0.297***
	(36.89)	(34.85)	(36.14)	(34.26)
Observations	81254	68946	68785	71391
Panel (b): Correct optimists – Prices (change in current month)				
ATT	0.035***	0.034***	0.040***	0.033***
	(5.40)	(5.18)	(5.72)	(4.90)
Observations	81254	68945	68778	71392
Panel (c): Incorrect optimists – Production (change in current month)				
ATT	0.063***	0.060***	0.072***	0.063***
	(8.58)	(7.94)	(9.20)	(8.13)
Observations	84029	74232	69715	73973
Panel (d): Incorrect optimists – Prices (change in current month)				
ATT	0.016***	0.015***	0.017***	0.011*
	(2.92)	(2.58)	(2.80)	(1.89)
Observations	84032	74232	69714	73978
Panel (e): Correct pessimist – Production (change in current month)				
ATT	-0.307***	-0.300***	-0.307***	-0.303***
	(-33.71)	(-30.52)	(-30.78)	(-32.00)
Observations	80282	66948	67112	68156
Panel (f): Correct pessimist – Prices (change in current month)				
ATT	-0.030***	-0.021**	-0.028***	-0.044***
	(-3.83)	(-2.52)	(-3.24)	(-5.66)
Observations	80285	66941	67109	68158
Panel (g): Incorrect pessimist – Production (change in current month)				
ATT	-0.086***	-0.093***	-0.074***	-0.086***
	(-9.99)	(-10.29)	(-7.92)	(-9.15)
Observations	79026	68414	65323	68835
Panel (h): Incorrect pessimist – Prices (change in current month)				
ATT	-0.003	-0.008	0.000	-0.008
	(-0.36)	(-1.07)	(0.00)	(-1.08)
Observations	79033	68420	65326	68842

Notes: Tables shows treatment effects on prices and production for different specifications. T-statistics in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
[1] Excluding the years 2008 and 2009.

of observations for which a propensity score can be computed as well as the number of observations that can be matched. Even though the number of matches is now smaller than before, there is still common support (see Figure C.3 in the appendix). We display balancing statistics in Figure C.4 and Figure C.5 in the appendix. For all four treatments balancing is achieved and no bias is above 5% (Table C.8 shows the bias in detail). In addition, the variance ratios are generally within the defined bounds with similar exceptions as before.

We report the ATTs in Table 3.7. We focus again on how firms' current production and price-setting decisions depend on optimism (upper part of the table) and on pessimism (lower part of the table). Now, however, we distinguish between correct and incorrect optimists and pessimists. Panels (a) and (b) show the results for production and prices for correct optimists, while Panels (c) and (d) display the results for incorrect optimists.

We present estimates for the baseline in column (1) and stress that results are, as before, robust across alternative specifications, see columns (2) to (4). We find that the effect of optimism on firms' current decisions is stronger for correct optimists than for incorrect optimists; likewise for pessimists. Still, also for optimism and pessimism that turns out to be incorrect in light of actual developments, we find a significant effect, except for prices in case of pessimists. As before, this effect may reflect a mixture of more upward or fewer downward adjustments compared to untreated firms. For pessimism it may reflect more downward adjustments and fewer upward adjustments.

In order to shed some light on this aspect, we rely once more on the transformation of the dependent variables into two binary variables indicating increases and decreases, respectively. Table 3.8 shows the results. We find that the probability of a production increase is 27.6 percentage points higher for correct optimists than for untreated firms and 7.2 percentage points higher for incorrect optimists. The probability of reducing production, instead, does not change much in both instances. A similar picture emerges for prices. By and large, the decisions of pessimists mirror those of optimists, i.e. there are more decreases and less increases. The results are shown in the bottom panel of Table 3.8.

In sum, we find that firm expectations matter for firm decisions. This holds not only for expectations that turn out to be correct ex-post, but also for incorrect optimism and pessimism. Hence, the role of expectations for today's decisions is not limited to a transmission channel of news. Our results show that expectations also have a noise component, that is, they cause firms to adjust prices and production even though there is no fundamental reason for firms to do so.

Table 3.8: Average treatment effect on the treated, correct and incorrect firms, increases and decreases in production and prices

	(1) Prod. increase	(2) Prod. decrease	(3) Price increase	(4) Price decrease
Panel (a): Correct optimists				
ATT	0.276***	-0.026***	0.028***	-0.007*
	(44.90)	(-5.87)	(5.45)	(-1.93)
Observations	81254	81254	81254	81254
Panel (b): Incorrect optimists				
ATT	0.072***	0.009**	0.012***	-0.004
	(13.82)	(2.11)	(2.84)	(-1.24)
Observations	84029	84029	84032	84032
Panel (c): Correct pessimists				
ATT	-0.033***	0.274***	0.000	0.030***
	(-6.62)	(39.14)	(0.01)	(5.57)
Observations	80282	80282	80285	80285
Panel (d): Incorrect pessimists				
ATT	-0.002	0.084***	0.004	0.007
	(-0.46)	(12.66)	(0.96)	(1.42)
Observations	79026	79026	79033	79033

Notes: Table shows treatment effects for binarized production and price indicators, i.e. separately considering increases and decreases. T-statistics in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.4.2 Further evidence

By now, we have established that firms respond to expectations – both to correct and incorrect ones – by adjusting current prices and production. In what follows, we turn to additional variables for which we may also expect an effect in light of this result. In each instance, we use the same framework as above, but report the ATT on variables other than production and prices.

In particular, we conjecture that optimistic (pessimistic) firms consider current inventories as too low (too high), independently of whether expectations turn out to be correct or not. In the survey, firms can evaluate the current state of their inventories as "too small" [1], "sufficient" [0], or "too large" [-1], see Table C.6 in the appendix. Panels (a) and (d) of Table 3.9 show how optimism and pessimism affect this assessment. As it turns out, both correctly and incorrectly optimistic firms evaluate their inventories as too low (Panel (a)). The opposite is true for correctly and incorrectly pessimistic firms (Panel (d)).

Furthermore, taking decisions based on expectations that turn out to be correct or incorrect should also have a bearing on future profits. Panels (b) and (c) of Table 3.9 re-

Table 3.9: Average treatment effect on the treated, correct and incorrect firms, inventories and profits

	Correct firms			Incorrect firms		
	(1) Baseline	(2) Radius 0.001	(3) Sample excl. fin. crisis[1]	(4) Baseline	(5) Radius 0.01	(6) Sample excl. fin. crisis[1]
Panel (a): Optimists – Inventories in t						
ATT	0.034***	0.030***	0.032***	0.027***	0.026***	0.027***
	(3.97)	(3.36)	(3.58)	(3.42)	(3.19)	(3.27)
Observ.	11150	9384	9543	11137	10188	9457
Panel (b): Optimists – Change in profits t+2						
ATT	0.141***	0.149***	0.136***	-0.070***	-0.074***	-0.076***
	(5.16)	(5.33)	(4.67)	(-2.91)	(-3.01)	(-2.96)
Observ.	11930	10661	10220	12149	11034	10410
Panel (c): Optimists – Change in profits t+3						
ATT	0.229***	0.240***	0.207***	-0.077***	-0.066***	-0.088***
	(8.49)	(8.79)	(7.21)	(-3.18)	(-2.63)	(-3.38)
Observ.	11403	9931	9712	11680	10562	9940
Panel (d): Pessimist – Inventories in t						
ATT	-0.067***	-0.063***	-0.072***	-0.018*	-0.026**	-0.022**
	(-5.98)	(-5.25)	(-6.34)	(-1.78)	(-2.50)	(-2.05)
Observ.	10586	9353	8768	10200	8865	8526
Panel (e): Pessimist – Change in profits t+2						
ATT	-0.122***	-0.141***	-0.148***	0.068**	0.032	0.088**
	(-3.44)	(-3.83)	(-4.16)	(2.24)	(1.01)	(2.56)
Observ.	11403	9756	9426	11267	10137	9491
Panel (f): Pessimist – Change in profits t+3						
ATT	-0.240***	-0.201***	-0.290***	0.047	0.035	-0.004
	(-6.59)	(-5.02)	(-7.57)	(1.52)	(1.09)	(-0.13)
Observ.	10356	8256	8295	10441	9087	8743

Notes: Tables shows treatment effects on inventories and demand for different specifications. T-statistics in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.
[1] Excluding the years 2008 and 2009.

port the effect of optimism on profits in the second and third month after impact.[10] We find that profits increase in both months for correct optimists. Given that those firms anticipated higher production and acted accordingly, this result appears plausible.[11]

[10] In the survey, the question about profits is asked twice a year, in May and in September. Our results are therefore based on different sets of firms. For results pertaining to profits two months after impact, we rely on expectations data from March and July. For profits three months after impact, we use responses from February and June. For details on the question see Table C.6.

[11] Profits in the first month after the price and production changes are less responsive to these changes.

Likewise, we find that profits of correct pessimists decline (Panels (e) and (f)). Profits of incorrect optimists also fall, see columns (4) to (6) in Panels (b) and (c). This decrease suggest acting based on incorrect expectations can be costly. We note, however, that according to our estimates, incorrect pessimists do not see their profits decline significantly (columns (4) to (6) in Panels (e) and (f)). This result appears to be consistent with our previous finding according to which they do not adjust their prices.

3.5 Noise and aggregate fluctuations

Up to now, we have focused on individual firms and, more specifically, we have documented that incorrect optimism and pessimism causes firms to adjust prices and production. In what follows, we investigate whether noise at the firm level matters for aggregate outcomes. Intuitively, a sufficiently large number of incorrectly optimistic firms may cause economic activity to rise at the aggregate level, and vice versa for pessimism.

Against this background, we first assess the degree to which the forecast errors of firms are correlated, both within sectors and across the entire economy. For this purpose, we now not only classify firms as correct or incorrect, but also quantify the extent to which firms are incorrect, following the approach of Bachmann et al. (2013).[12] We report descriptive statistics and the serial correlation pattern of forecast errors in Table C.4 and Table C.5 in the appendix. Table 3.10 shows that forecast errors are generally positively correlated within sectors and more strongly so than across all firms. This finding is consistent with the notion that optimism and pessimism can be contagious, assuming that firms in the same sector interact more.

3.5.1 An aggregate measure of noise

In what follows, we develop an aggregate measure of noise. Our goal is to mimic the micro-level analysis by creating a measure of noise, which considers both the current fundamentals as well as the ex-post performance. For this purpose, we rely on the ordered probit model described in equation (3.2) in Section 3.3 in order to account for fundamental determinants of expectations. However, now, rather than matching firms based on their propensity score, we compute the difference between a firm's response

[12]Specifically, the error is 0 if the firm is correct, that is, if the sign of the expectation and the average realization is the same. If the firm is incorrect, the error equals the difference between the sum of realized production in *t+1* to *t+3* and the expectation in *t*, divided by 3.

Table 3.10: Correlation of firms' forecast errors with economy and sectoral averages

Sector	Correlation with same sector	Correlation with all firms	Sector	Correlation with same sector	Correlation with all firms
All sectors	0.1967	0.1310	Rubber&plastic prod.	0.1902	0.1513
Food	0.1558	0.0383	Glass prod.	0.1889	0.1266
Beverages	0.2669	0.0186	Basic metals	0.2735	0.1977
Tobacco	0.6281	-0.0207	Fabricated metal prod.	0.1646	0.1465
Textiles	0.1985	0.1018	Computer&electronic prod.	0.1700	0.1339
Wearing apparel	0.2185	0.0397	Electrical equipment	0.1775	0.1460
Leather&related prod.	0.2965	0.0893	General-purpose machinery	0.1568	0.1333
Wood&cork products	0.2161	0.1361	Motor vehicles&trailers	0.2592	0.1966
Paper products	0.2130	0.1687	Other transport equi.	0.3299	0.1413
Printing	0.1731	0.0989	Furniture	0.2245	0.1081
Coke&refined petrol.	0.4659	0.0865	Other manufacturing	0.1969	0.1060
Chemical products	0.2126	0.1697	Repair&installation	0.3821	0.0881
Pharmaceuticals	0.3073	-0.0134			

Notes: Correlation of firms' individual forecast error with average of the forecast error in the same 2-digit WZ08 sector and the whole economy. Shown separately for each 2-digit sector. Error computed following the approach of Bachmann et al. (2013): the error is 0 if the firm is correct, that is, if the sign of the expectation and the average realization is the same. If the firm is incorrect, the error equals the difference between the sum of realized production in *t+1* to *t+3* and the expectation in *t*, divided by 3.

and the prediction of the ordered probit model.[13] Based on the ordered probit model, we classify firms as optimists whenever they expect an increase in production even though the model predicts otherwise, and as pessimists when they expect and decrease despite a different prediction by the model.[14] In a second step, to capture incorrect optimism and pessimism, we only consider firms, which are incorrect from an ex-post perspective as defined in Section 3.4 above. Finally, we compute the share of firms that fulfill these two criteria relative to all firms in a given month.

In computing the aggregate noise measure, we consider three alternative weights. First, we compute the share giving equal weight to each firm. For the second measure, we use the number of employees as weights. We drop the largest 5 percent of our observations to ensure that results are not driven by large individual firms. Finally, we

[13]We use the ordered probit, rather than distinct models for optimists and pessimists, because we seek to account for all outcomes simultaneously. Recall that the ordered probit model includes as control variables time and sector fixed effects, the sector average of the reported state of business in each month, three lags of the dependent variables, and all firm-specific variables listed in Table 3.2 (including three lags for each of the survey variables and interaction terms).

[14]The predicted response is the response to which the ordered probit model assigns the highest probability.

Figure 3.5: An aggregate measure of noise, 1991 to 2016

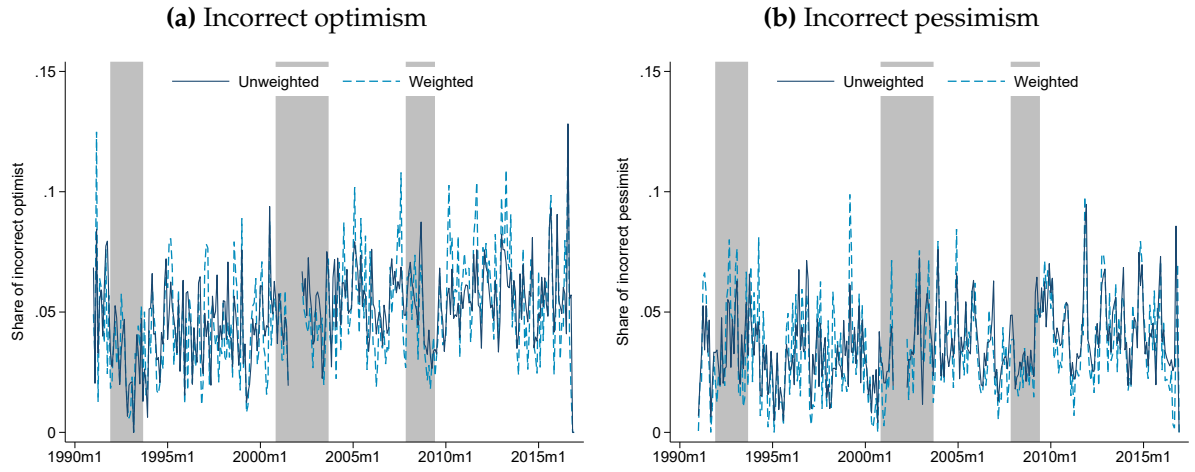

Notes: Aggregate time series for incorrect optimism and pessimism, unweighted and weighted by employees. Shaded areas indicate recession periods as defined by the German Council of Economic Experts.

weigh firms in line with the approach by the ifo institute for aggregating answers to the business climate index (Sauer and Wohlrabe 2018). This approach weights all firms within a 2-digit WZ08 sector (the German system of industry classification) using the number of employees in production as reported in the survey. Instead of using the number of employees directly, the weight is a logarithmic transformation of employment.[15] The sector averages are then aggregated using data on gross value added by sector from the German Statistical Office.

Figure 3.5 displays the unweighted and the employee-weighted time series for incorrect optimism and pessimism (using ifo weights results in a very similar time series). The computation of these shares requires firms to be in the survey for at least eight consecutive months because we need three lags for the estimation of the ordered probit model and four leads for the computation of the forecast error.[16] This leads to a gap in our time series from August 2001 to March 2002 because the ifo survey was not conducted in December 2001. In addition, it reduces the number of observations in the last five months of 2016. The main takeaway of Figure 3.5 is that there is considerable variation of incorrect optimism and pessimism over time. In addition, we note that the time series exhibit little persistence.

[15]Specifically, the weight is $w = (\log_{10}(N))^e$, with N being the number of employees, see the EBDC Questionnaire Manual. This transformation ensures that very large firms do not distort the averages.

[16]We need data on the production for the next three months. Since production is reported only for the previous month, we need four leads of the survey.

3.5.2 The aggregate effect of noise

Our noise measure is an aggregate of firm-level expectations. We compute this measure as optimism/pessimism relative to current fundamentals. In addition, we allow for time-fixed effects. As such, aggregate optimism and pessimism is unlikely to be caused by macroeconomic shocks. Moreover, even if macroeconomic shocks were to affect firms differently, it is important to note that we focus on incorrectly optimistic/pessimistic firms, that is, firms whose expectations turn out to be unwarranted from an ex-post perspective. For these reasons, we treat our aggregate measure as an explanatory variable of macroeconomic outcomes that is not itself caused by macroeconomic developments.

We may therefore rely on local projections to estimate the causal effect of noise on aggregate outcomes (Jordà 2005). Formally, using e_t^o and e_t^p to denote the time-series observations for incorrect optimism and pessimism, respectively, and x_t for the realization of a macroeconomic variable of interest, we estimate following model:

$$x_{t+h} = c^{(h)} + \sum_{j=1}^{J} \alpha_j^{(h)} x_{t-j} + \sum_{k=0}^{K-1} \beta_k^{(h)} e_{t-k}^o + \sum_{k=0}^{K-1} \gamma_k^{(h)} e_{t-k}^p + \varepsilon_{t+h}, \tag{3.5}$$

where $c^{(h)}$ is a (horizon-specific) constant.[17] In addition, we include a linear time trend. To enhance efficiency, we also include the residuals of the previous horizon when increasing the horizon by steps of one (Jordà 2005). For the estimation, we include 1 lag of the dependent variable and 12 lags of incorrect optimism and pessimism. In all specifications, we include both incorrect optimism and pessimism to account for a potential correlation between the two variables. The estimated coefficients β^h and γ^h provide a direct measure of the impulse response at horizon h, given a unit shock in period t.

We show the effect of a noise shock in Figure 3.6. It displays the response to an increase of one standard deviation in the share of incorrectly optimistic firms. The top panels show the response of industrial production in the manufacturing sector (IP), measured in percentage deviations from the trend, while the bottom panels show the response of the producer price index in the manufacturing sector (PPI), also measured in percentage deviations from the trend. The left column displays results using the unweighted measure, the middle column is based on employee-weighted shares, while the right column shows responses for ifo weights.

In each instance, time is measured in months along the horizontal axis. The blue solid line represents the point estimate, shaded areas indicate 68 and 90 percent confi-

[17] In order to account for non-stationarity, the dependent variable can be expressed relative to its pre-shock level (Stock and Watson 2018). Including lags of the dependent variable, as we do above, generally yields the same result.

Figure 3.6: Effect of a noise shock, optimism

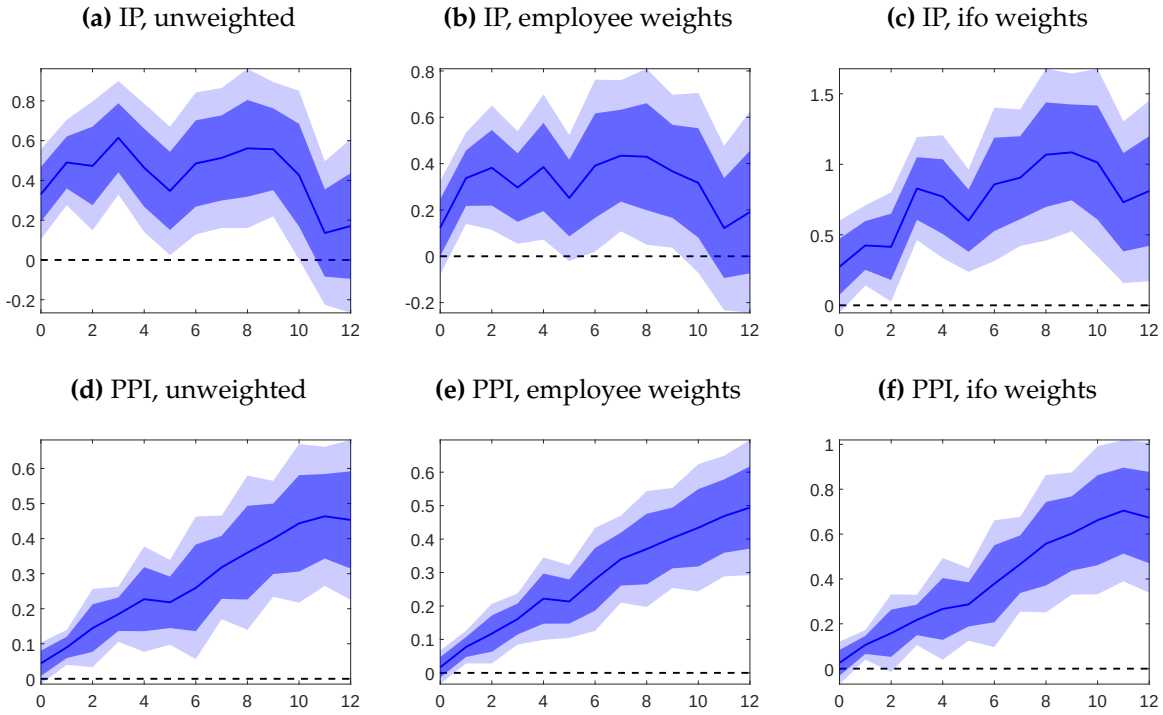

Notes: Responses of industrial production (IP) and the producer price index (PPI), both for the manufacturing sector, to incorrect optimism (one standard deviation shock). Local projections with constant, linear trend, one lag of dependent variable and 12 of the shocks. Shaded areas indicate 68 and 90 percent confidence intervals. Data for IP from the German Statistical Office. Data on the PPI from the German Bundesbank.

dence intervals. We find that industrial production responds strongly and significantly to the shock. The observed increase is temporary and becomes insignificant after approximately one year, except in case we use ifo weights. Consistent with the results at the firm level, reported in Section 3.4 above, we also find a strong and significant increase in the price level after an increase in incorrect optimism. This reaction is in line with the interpretation of noise shocks as a specific form of demand shocks (Enders, Kleemann, and Müller 2017; Lorenzoni 2009).

Figure 3.7 displays the results for incorrect pessimism. Here we find much weaker effects. Specifically, the response of industrial production is insignificant throughout and the producer price index falls (marginally significant) only in period 1 after the shock. This result is in line with the firm-level observation that incorrect pessimism does not seem to cause a significant downward adjustment of prices, see Table 3.8 above.

Finally, Table 3.11 displays a forecast error variance decomposition (FEVD) for a horizon of 12 months, using the methodology of Gorodnichenko and Lee (2017). We present results for incorrect optimism and pessimism and for all three measures of noise. We find that incorrect optimism is responsible for 10-18% of aggregate fluctua-

Figure 3.7: Effect of a noise shock, pessimism

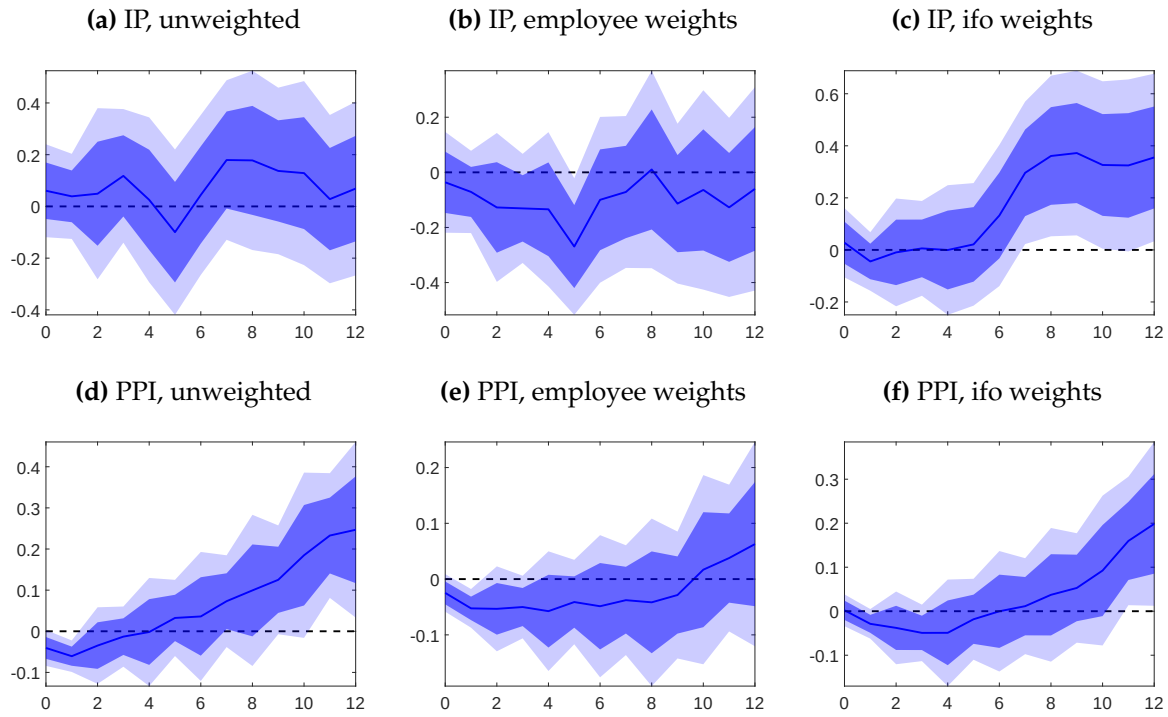

Notes: Responses of industrial production (IP) and the producer price index (PPI), both for the manufacturing sector, to incorrect pessimism (one standard deviation shock). Local projections with constant, linear trend, one lag of dependent variable and 12 of the shocks. Shaded areas indicate 68 and 90 percent confidence intervals. Data for IP from the German Statistical Office. Data on the PPI from the German Bundesbank.

tions of industrial production at a one-year horizon. At the same horizon, around 20% of the PPI is driven by incorrect optimism. This amounts to a sizeable contribution. Pessimism, on the other hand, has much a smaller effect on IP and the PPI. The specification for the PPI with unweighted observations delivers a value of around 6%, similar to that with ifo weights for IP. Regarding the small effect on prices, these results are in line with those of firm-level effects, reported in Section 3.4 above.

Table 3.11: Forecast error variance decomposition (one year horizon)

	Variable	Unweighted	Empl. weights	ifo weights
Optimism	IP	14%	10%	18%
	PPI	19%	23%	20%
Pessimism	IP	1.3%	1.7%	6.4%
	PPI	5.9%	0.9%	2.9%

3.6 Conclusion

In this chapter, we ask to what extent firm expectations matter for firm decisions. From a theoretical point of view, the answer to the question is obvious: expectations should matter a great deal. However, to date there is little direct evidence to support the theory. In this analysis, we aim at filling this gap using a particularly suited data set and a new identification strategy.

We use a large survey of firms in the German manufacturing sector. Firms report on a monthly basis whether they expect production to increase, to remain constant, or to decline. For each firm-month observation, we also observe a large number of firm characteristics, including balance-sheet information. This allows us to match firms, which differ in their expectations, but not in their fundamentals, and to identify the effect of expectations on firm decisions.

We find that optimistic firms tend to raise prices and production today. This result can be explained in two ways. According to the "news view," firms simply have additional information about future developments, which is not reflected in current fundamentals. Their optimism is thus fundamentally justified – by future fundamentals, so to speak. According to the "noise view," firms are optimistic or pessimistic for no fundamental reason. They simply have wrong ideas about the future and their current actions are driven by animal spirits.

We disentangle the effect of news and noise on firms' decision makers by considering the ex-post forecast error. Because we observe the actual developments ex-post, we can assess whether firms were right about the future or not. We match incorrectly optimistic and pessimistic firms, in turn, to neutral firms, as well as correctly optimistic/pessimistic firms. For both groups we find a positive effect of optimism on current production and prices, and a negative effect of pessimism. The effect is considerably smaller in case of incorrect expectations, however. Still, these results show that firm decisions are to some extent caused by noise.

Finally, we turn to the aggregate effects of optimism and pessimism. For this purpose, we compute an aggregate measure of noise based on firm-level data. We run local projections to estimate the effect of noise shocks on aggregate production and prices. We find, in particular, that positive noise shocks (optimism) cause industrial production and producer prices to increase. Negative shocks play only a small role in causing business cycles.

Appendices to Chapter 3

C1 Additional descriptive statistics

Table C.1: Observations and average duration in panel, 1991 to 2016

	Full sample	Sample with balance sheet data
Avg. duration in survey (months)	140.2	87.1
Avg. number of responses (months)	119.3	74.1
Response rate	82.5%	83.8%
Respondents	6625	4938
Respondents × months	620671	322839

Notes: Number of firms and duration of firms in the ifo survey. Separately for the full sample and the sample with balance sheet data.

Table C.2: Descriptive statistics for different samples, 1991 to 2016

	Full sample			Sample with balance sheet data		
	Mean	Std. Dev.	Observ.	Mean	Std. Dev.	Observ.
Employees in production	505	3024	620411	579	3665	322823
Exp. production, t	0.00	0.55	607578	0.02	0.56	313788
Exp. state of business, t	-0.01	0.61	619148	0.01	0.61	321946
Activity, t-1	-0.05	0.58	585075	-0.03	0.58	311828
Prices, t-1	0.00	0.42	597422	0.02	0.43	320617
Orders, t	-0.27	0.63	618229	-0.21	0.64	321226
Foreign orders, t	-0.28	0.57	616199	-0.22	0.58	319775
Capacity utilization, t	81.15	16.36	527754	81.10	16.13	279901
Demand, t-1	-0.02	0.64	598360	-0.00	0.65	321154
Inventories, t	-0.17	0.51	425005	-0.15	0.49	227742
State of profits, t	-0.12	0.69	51995	-0.11	0.69	42528
Change in profits, t	-0.06	0.71	51389	-0.05	0.71	42031

Notes: Descriptive statistics for the variables used in the micro-level analysis. Separately for the full sample and the sample with balance sheet data. For details on the related survey questions see Table C.6.

Table C.3: Sample attrition, 1991 to 2016

Start date	Total	Fraction of firms surviving after										
		6m	1y	2y	3y	4y	5y	6y	7y	8y	9y	10y
1991m1	1896	1	1	0.99	0.96	0.94	0.91	0.87	0.84	0.81	0.78	0.75
1992m1	2114	1	0.99	0.96	0.93	0.89	0.86	0.82	0.79	0.76	0.73	0.70
1993m1	2320	0.97	0.95	0.92	0.88	0.85	0.81	0.78	0.74	0.71	0.69	0.65
1994m1	2213	0.98	0.96	0.93	0.89	0.84	0.81	0.78	0.74	0.72	0.68	0.64
1995m1	2130	0.97	0.95	0.91	0.87	0.83	0.80	0.77	0.74	0.71	0.66	0.62
1996m1	2072	0.98	0.96	0.92	0.88	0.85	0.81	0.78	0.74	0.70	0.66	0.63
1997m1	2040	0.97	0.95	0.92	0.88	0.84	0.81	0.77	0.72	0.68	0.65	0.62
1998m1	1979	0.98	0.96	0.93	0.88	0.85	0.81	0.75	0.71	0.68	0.64	0.60
1999m1	2038	0.97	0.95	0.91	0.87	0.83	0.77	0.73	0.69	0.66	0.62	0.58
2000m1	2055	0.98	0.96	0.92	0.87	0.81	0.77	0.73	0.69	0.65	0.61	0.56
2001m1	2050	0.98	0.96	0.91	0.84	0.80	0.76	0.72	0.67	0.63	0.58	0.56
2002m1	2006	0.98	0.95	0.88	0.83	0.79	0.75	0.70	0.66	0.61	0.58	0.54
2003m1	1990	0.97	0.93	0.88	0.83	0.79	0.73	0.69	0.64	0.61	0.57	0.53
2004m1	2044	0.97	0.94	0.89	0.84	0.78	0.74	0.68	0.65	0.61	0.57	0.52
2005m1	1979	0.97	0.94	0.90	0.83	0.78	0.72	0.69	0.64	0.60	0.55	0.51
2006m1	1973	0.97	0.95	0.88	0.83	0.77	0.73	0.69	0.64	0.59	0.55	0.50
2007m1	1935	0.96	0.93	0.88	0.82	0.78	0.73	0.68	0.63	0.59	0.54	0
2008m1	1859	0.97	0.95	0.88	0.84	0.79	0.73	0.68	0.63	0.58	0	0
2009m1	1874	0.96	0.93	0.88	0.84	0.77	0.70	0.65	0.60	0	0	0
2010m1	1845	0.97	0.94	0.89	0.82	0.75	0.70	0.64	0	0	0	0
2011m1	1930	0.97	0.94	0.87	0.78	0.73	0.68	0	0	0	0	0
2012m1	2228	0.96	0.92	0.83	0.76	0.70	0	0	0	0	0	0
2013m1	2105	0.94	0.91	0.84	0.77	0	0	0	0	0	0	0
2014m1	1976	0.96	0.92	0.86	0	0	0	0	0	0	0	0
2015m1	1881	0.96	0.92	0	0	0	0	0	0	0	0	0

Notes: Table shows share of firms from initial period given by the row still in the sample after the time specified in the column, e.g. the last row shows that of the 1881 firms which were part of the sample in January 2015, 96% were still there after 6 months, 93% were still there after 1 year, and 84% after two years. The zeros in the bottom right corner are due to our sample ending in December 2016. '6m' = 6 months, '1y' = 1 year.

Table C.4: Forecast error in the full sample, 1991 to 2016

	Full sample				Excl. recession months			
	Mean	Std. Dev.	Min	Max	Mean	Std. Dev.	Min	Max
Forecast error, *t*	-0.038	0.362	-1.33	1.33	-0.022	0.358	-1.33	1.33
Exp. production, *t*	0.002	0.551	-1	1	0.028	0.545	-1	1
Realized prod., *t+1* to *t+3*	-0.152	1.297	-3	3	-0.053	1.273	-3	3

Notes: Expected production based on response to Q1 and realized production based on Q3, see Table C.6 for wording of question and coding of answers. Forecast error computed following the approach of Bachmann et al. (2013): the error is 0 if the firm is correct, that is, if the sign of the expectation and the average realization is the same. If the firm is incorrect, the error equals the difference between the sum of realized production in *t+1* to *t+3* and the expectation in *t*, divided by 3. Right panel excludes months in which Germany was in a recession according to the German Council of Economic Experts.

Table C.5: Serial correlation of the forecast error full sample, 1991 to 2016

	Correlation with forecast error in month									
	t-1	*t-2*	*t-3*	*t-4*	*t-5*	*t-6*	*t-7*	*t-8*	*t-9*	*t-10*
Forecast error, *t*	0.580	0.341	0.152	0.128	0.120	0.109	0.095	0.087	0.082	0.095

Notes: Table shows correlation of forecast errors over time at the firm level. Forecast error computed following the approach of Bachmann et al. (2013): the error is 0 if the firm is correct, that is, if the sign of the expectation and the average realization is the same. If the firm is incorrect, the error equals the difference between the sum of realized production in *t+1* to *t+3* and the expectation in *t*, divided by 3.

Figure C.1: Average expectations and industrial production, 1991 to 2016

(a) Production exp. and state of business exp.

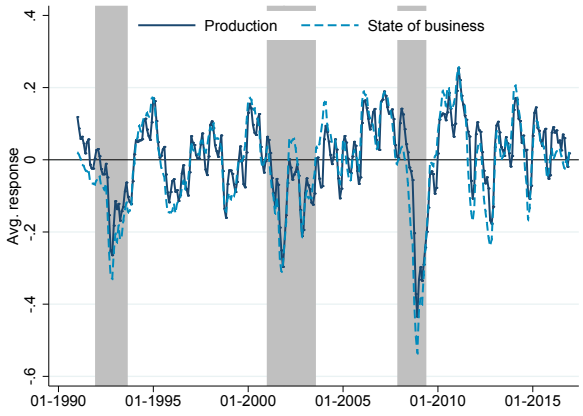

(b) Production exp. and industrial production

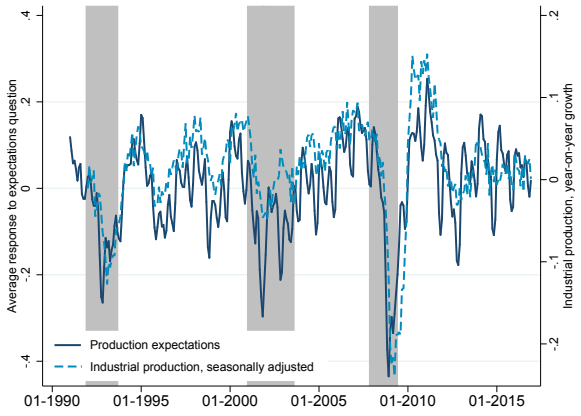

Notes: Shaded areas mark recession periods as defined by the German Council of Economic Experts. Average response computed as share of positive responses (*increase/improve*) minus share of negative responses (*decrease/worsen*). Survey data from the BEP. Industrial Production from the German Statistical Office.

Figure C.2: Distribution of days firms respond to survey within month, 2004 to 2016

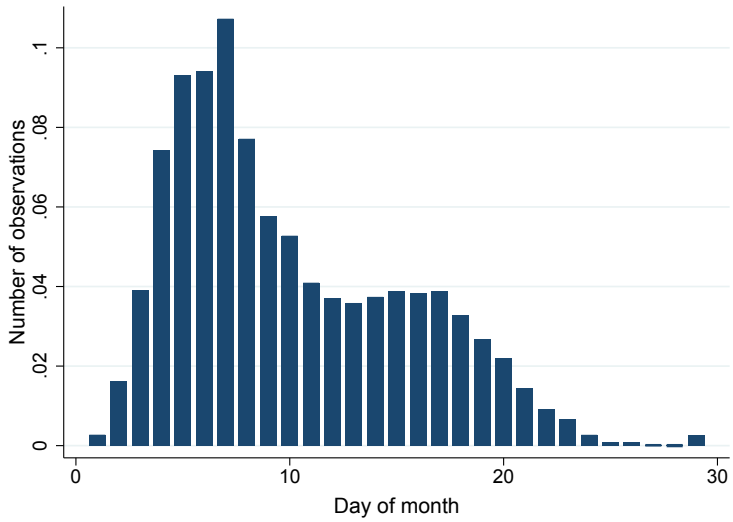

Notes: Histogram shows days of the month on which firms respond in the online sample, for which the time of handing in the questionnaire is recorded.

C2 Details on survey questions

Table C.6: Complete list of survey questions used in our analysis

Label	Name	Question	Possible answers
Q1	expected production	Expectations for the next 3 months: Our domestic production activity regarding good XY will probably ...	increase [1] not change [0] decrease [-1]
Q2	expected state of business	Expectations for the next 6 months: Taking economic fluctuations into account our state of business will be...	rather more favorable [1] not changing [0] rather less favorable [-1]
Q3	production	Tendencies in the previous month: Our domestic production activities with respect to product XY have ...	increased [1] not changed [0] decreased [-1]
Q4	prices	Tendencies in the previous month: Taking changes of terms and conditions into account, our domestic sales prices (net) for product XY have been ...	increased [1] not changed [0] decreased [-1]
Q5	employees	Number of employees: In our company (domestic enterprises only) we employ [...] persons, of which x persons are for producing product XY.	x is the number of persons employed for XY
Q6	orders	We consider our order backlog to be...	relatively high [1] sufficient [0] too small [-1] no export of XY [4]
Q7	foreign orders	We consider our order backlog for exports to be...	relatively high [1] sufficient [0] too small [-1] no export of XY [4]
Q8	capacity utilization	The current utilization of our capacities for producing XY (standard utilization = 100%) is currently x%.	x is a value between 30 and 100 divisible by 10
Q9	demand	Tendencies in the previous month: The demand situation with respect to product XY is ...	better [1] not changed [0] worse [-1]
Q10	inventories	Current situation: We evaluate our current stock of unsold finished goods of XY to be ...	too small [1] sufficient [0] too large [-1] no stock keeping [4]
Q11	change in profits	Profit situation and development: As compared to fall last year/the first quarter this year[1] the profit situation of our company measured by the operating results from customary business operations ...	improved [1] did not change [0] deteriorated [-1]

Notes: Authors' translation of the most recent formulation of the question in German according to the EBDC Questionnaire manual. We only show the answer possibilities that we consider. Specifically, we exclude "no production" or similar answers, which indicate that the question does not apply to the firm.

[1] Questions asked biannually in May and September. In May question refers to "fall last year", in September it refers to "first quarter this year".

Table C.7: Main survey questions, changes over time

Label	Time period	Question
Q1	01/1980-06/1994	Our domestic production activity regarding good XY in the next 3 months **taking economic fluctuations into account – i.e. after eliminating purely seasonal fluctuations –** will probably … increase/not change/decrease.
	07/1994-06/1997	Our domestic production activity regarding good XY in the next 3 months **taking economic fluctuations into account – i.e. after eliminating purely seasonal fluctuations –** will probably … increase/not change/decrease/*no substantial domestic production.*
	07/1997-11/2001	Our domestic production activity regarding good XY in the next 3 months **taking economic fluctuations into account** will probably … increase/not change/decrease/no substantial domestic production.
	Since 01/2002	Expectations for the next 3 months: Our domestic production activity regarding good XY will probably … increase/not change/decrease/no substantial domestic production.
Q2	01/1980-06/1997	Our state of business regarding good XY in the next 6 months taking economic fluctuations into account **– i.e. after eliminating purely seasonal fluctuations –** will be … rather more favorable/ not changing/rather less favorable.
	07/1997-11/2001	Our state of business regarding good XY in the next 6 months taking economic fluctuations into account will be … rather more favorable/ not changing/rather less favorable.
	Since 01/2002	Expectations for the next 6 months: Taking economic fluctuations into account our state of business will be … rather more favorable/ not changing/rather less favorable.
Q3	01/1980-06/1994	**In comparison to the previous month** our domestic production activities regarding good XY have … **been more lively/unchanged/weaker**.
	07/1994-11/2001	**In comparison to the previous month** our domestic production activities regarding good XY have … **been more lively/unchanged/weaker** */no substantial domestic production*.
	01/2002-02/2002	**In the last 2-3 months** our domestic production activities regarding good XY have … **been more lively/unchanged/weaker**/no substantial domestic production.
	Since 03/2002	Tendencies in the previous month: Our domestic production activities with respect to product XY have … increased/not changed/decreased/no substantial domestic production.
Q4	01/1980-11/2001	**Compared to the previous month** our domestic prices (net prices) of good XY – taking changes of terms and conditions into account – have been … increased/not changed/decreased.[1]
	01/2001-02/2002	**In the last 2-3 months** our domestic prices (net) of good XY – taking changes of terms and conditions into account – have been … increased/not changed/decreased.
	Since 03/2002	Tendencies in the previous month: Taking changes of terms and conditions into account, our domestic sales prices (net) for product XY have been … increased/not changed/decreased.

Notes: Authors' translation of the question in German according to the EBDC Questionnaire manual. **Bold font** highlights components which change from the initial formulation or drop out. *Italic font* highlights components which are added later on.

[1] In several months in 1980 the question was split into two parts, one covering regular and additional orders.

C3 Balancing statistics

Table C.8: Standardized bias, baseline specification

	Optimists			Pessimists		
Variable	All	Correct	Incorrect	All	Correct	Incorrect
Dep. var, *t-1*	1.8	1.6	1.8	-0.7	-0.7	-0.9
Dep. var, *t-2*	0.2	-0.2	-0.3	-1.7	-3.5	-2.5
Dep. var, *t-3*	0.8	0.7	0.1	-0.7	-2.8	-0.5
Production, *t-1*	1.3	0.1	-0.6	0.2	-0.3	-0.6
Production, *t-2*	1.1	0.3	-0.9	0.3	1.1	0.8
Production, *t-3*	-0.1	-0.2	-0.8	-0.9	-1.3	1.6
Prices, *t-1*	0.3	1.8	-0.0	0.3	1.5	0.1
Prices, *t-2*	1.0	-0.8	0.8	1.1	3.1	-0.1
Prices, *t-3*	0.8	1.5	0.1	0.3	2.0	-0.6
Demand, *t-1*	1.0	1.1	0.6	-1.1	-0.6	-1.5
Demand, *t-2*	1.1	2.0	-1.3	-1.5	-1.0	1.2
Demand, *t-3*	-0.1	-0.4	-0.8	-1.4	-1.7	1.6
Capacity, *t-1*	0.4	0.4	0.3	1.3	3.3	1.0
Capacity, *t-2*	0.4	0.6	0.5	0.8	3.1	1.3
Capacity, *t-3*	0.9	1.2	-0.1	0.3	3.1	0.6
Employees, *t*	0.2	-0.8	0.1	1.6	1.0	1.2
Avg. state of business, sector, *t*	0.9	1.8	-0.5	-0.5	-0.9	-0.3
State of business, *t*	0.4	-0.7	-0.5	0.9	2.0	1.2
State of business, *t-1*	0.5	-0.0	-0.7	0.7	1.0	1.0
State of business, *t-2*	0.6	0.6	-0.5	-0.2	0.9	1.8
State of business, *t-3*	0.0	0.1	-0.3	-0.4	-0.2	0.7
Orders, *t*	0.4	-1.4	-0.4	1.6	1.8	-0.4
Orders, *t-1*	2.0	2.1	-0.8	0.4	0.7	0.1
Orders, *t-2*	1.2	1.4	0.0	-0.3	-1.4	-0.6
Orders, *t-3*	1.1	0.9	0.0	0.6	-0.9	0.1
Foreign orders, *t*	1.1	-0.4	-0.4	-0.4	-2.1	-0.0
Foreign orders, *t-1*	1.9	1.0	-0.8	-0.5	-1.5	0.9
Foreign orders, *t-2*	1.1	0.6	-0.5	-1.0	-1.7	-0.6
Foreign orders, *t-3*	0.9	0.8	-0.1	-0.1	-2.1	-0.1
Debt share, *t*	-2.1	-2.8	-1.2	-0.2	-4.5	-2.6
Financing coefficient, *t*	-0.4	-0.7	-1.5	0.7	0.9	0.1

Notes: Table shows standardized bias which captures the difference in means between the treated and untreated groups. The standardized bias is the mean difference of each variable in the treated and untreated groups relative to the variances, details can be found in equation (3.4) in Section 3.3.

Table C.9: Standardized bias, specification with radius r=0.01

	Optimists			Pessimists		
Variable	All	Correct	Incorrect	All	Correct	Incorrect
Dep. var, *t-1*	1.3	2.4	1.9	-0.1	-1.1	-0.1
Dep. var, *t-2*	0.1	-0.9	-0.3	-1.8	-2.8	-2.2
Dep. var, *t-3*	0.1	0.8	0.5	-0.2	-2.6	0.6
Production, *t-1*	0.8	0.1	-1.6	0.3	0.1	-0.4
Production, *t-2*	0.8	-0.7	-0.3	0.9	2.3	2.4
Production, *t-3*	-1.0	-2.2	-0.8	-0.4	0.4	3.4
Prices, *t-1*	0.6	1.4	0.0	-0.2	3.6	-0.5
Prices, *t-2*	1.3	-0.6	2.0	0.8	4.7	1.0
Prices, *t-3*	0.7	1.2	0.1	-0.1	3.5	-0.5
Demand, *t-1*	0.9	1.6	-0.5	-1.3	-1.1	-1.6
Demand, *t-2*	0.9	0.8	-1.2	-0.3	-1.5	3.0
Demand, *t-3*	-0.4	-0.7	-0.5	-1.0	-1.6	3.5
Capacity, *t-1*	0.0	-0.7	1.2	2.8	4.3	1.4
Capacity, *t-2*	0.1	-0.6	1.5	2.3	3.7	1.1
Capacity, *t-3*	0.6	-0.6	0.8	1.7	3.8	0.5
Employees, *t*	0.7	-0.0	0.6	2.0	2.0	-0.3
Avg. state of business, sector, *t*	0.6	1.0	0.0	-0.7	-0.5	-1.0
State of business, *t*	0.8	-1.6	-0.8	1.9	3.6	1.5
State of business, *t-1*	0.5	-1.0	0.4	2.1	2.7	1.2
State of business, *t-2*	-0.0	-0.4	-0.2	1.4	1.7	1.9
State of business, *t-3*	-0.2	-1.4	0.7	0.7	0.4	1.0
Orders, *t*	0.7	-2.1	-0.8	2.0	3.6	0.5
Orders, *t-1*	1.5	0.5	-0.2	1.3	0.9	0.5
Orders, *t-2*	0.5	0.6	0.6	0.6	-0.1	-0.7
Orders, *t-3*	0.9	-0.2	1.1	1.8	0.5	-0.5
Foreign orders, *t*	1.3	-1.2	-0.9	0.5	-1.1	0.3
Foreign orders, *t-1*	1.6	-0.2	-0.4	0.3	-0.7	1.2
Foreign orders, *t-2*	0.9	-0.6	-0.9	0.0	-0.7	-1.4
Foreign orders, *t-3*	1.0	-0.4	0.1	1.1	-1.0	-0.9
Debt share, *t*	-1.7	-0.9	-1.2	0.1	-5.2	-2.2
Financing coefficient, *t*	-1.3	-0.6	-1.2	1.5	1.5	0.7

Notes: Table shows standardized bias which captures the difference in means between the treated and untreated groups. The standardized bias is the mean difference of each variable in the treated and untreated groups relative to the variances, details can be found in equation (3.4) in Section 3.3.

Table C.10: Standardized bias, specifications with ordered probit and state of business expectations

	Ordered probit		State of business exp.	
Variable	Optimists	Pessimists	Optimists	Pessimists
Dep. var, *t-1*	-0.7	-0.2	2.0	-1.0
Dep. var, *t-2*	0.6	-3.6	1.4	-1.0
Dep. var, *t-3*	2.0	-0.3	0.4	-1.1
Production, *t-1*	-7.5	-12.8	-0.1	-1.1
Production, *t-2*	-1.4	-2.7	-0.9	0.0
Production, *t-3*	-0.3	-0.4	-0.9	-0.2
Prices, *t-1*	-2.8	-2.5	0.0	-0.2
Prices, *t-2*	-1.9	-0.7	-0.5	0.1
Prices, *t-3*	-0.5	-0.3	-0.1	0.5
Demand, *t-1*	-4.3	-5.1	0.7	-1.5
Demand, *t-2*	-1.8	-2.2	-0.2	-0.4
Demand, *t-3*	-0.4	-1.1	-0.8	-0.9
Capacity, *t-1*	-6.5	-6.2	-0.8	-1.2
Capacity, *t-2*	-5.1	-3.2	-0.9	-0.8
Capacity, *t-3*	-4.6	-3.0	-0.8	-0.9
Employees, *t*	-0.2	-0.3	-0.1	1.4
Avg. state of business, sector, *t*	-0.6	-2.8	-0.4	-0.5
State of business, *t*	-6.7	-5.0	-0.7	0.2
State of business, *t-1*	-3.5	-1.2	-1.0	0.2
State of business, *t-2*	-2.6	-0.2	-0.7	-0.0
State of business, *t-3*	-2.8	0.8	-0.3	-0.0
Orders, *t*	-8.6	-6.3	-1.2	0.1
Orders, *t-1*	-4.3	-5.0	-0.4	-0.7
Orders, *t-2*	-4.3	-3.8	-0.4	-0.9
Orders, *t-3*	-3.9	-1.6	-0.2	-0.9
Foreign orders, *t*	-5.2	-2.3	-0.7	-1.2
Foreign orders, *t-1*	-2.6	-1.9	-0.5	-1.8
Foreign orders, *t-2*	-2.5	-1.6	-0.6	-1.8
Foreign orders, *t-3*	-2.5	0.2	-0.2	-1.1
Debt share, *t*	-0.5	2.5	-1.2	-0.4
Financing coefficient, *t*	-0.4	1.7	0.6	0.6

Notes: Table shows standardized bias which captures the difference in means between the treated and untreated groups. The standardized bias is the mean difference of each variable in the treated and untreated groups relative to the variances, details can be found in equation (3.4) in Section 3.3.

Figure C.3: Histogram of the density of the propensity scores

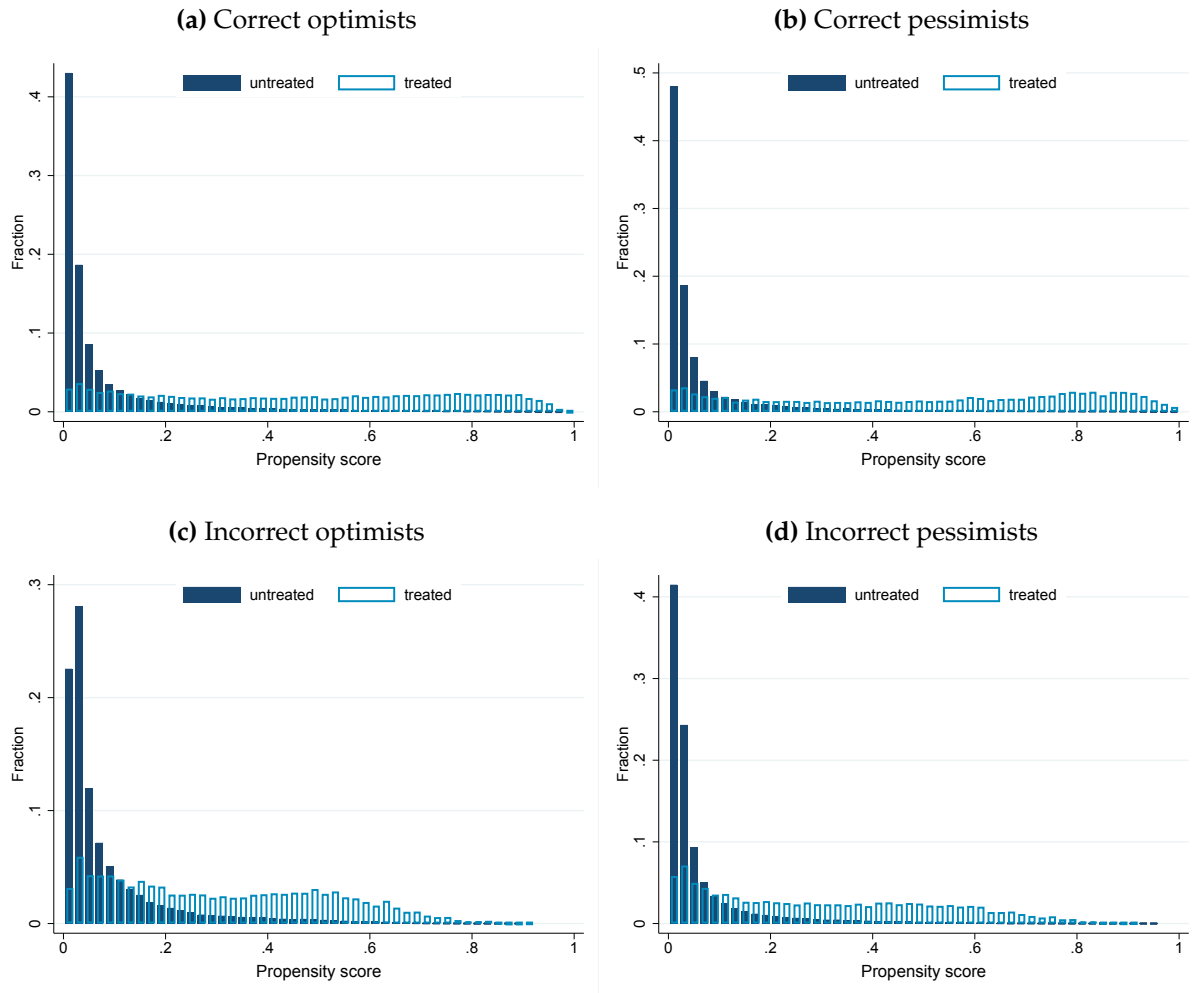

Notes: Histograms show the propensity scores for treated and untreated firms respectively, estimated as described by equation (3.3). In Panel (a) treated firms are correct optimists, Panel (b) treated firms are correct pessimists, in Panel (c) treated firms are incorrect optimists, and Panel (d) treated firms are incorrect pessimists.

Figure C.4: Standardized bias, before and after matching, correct and incorrect firms

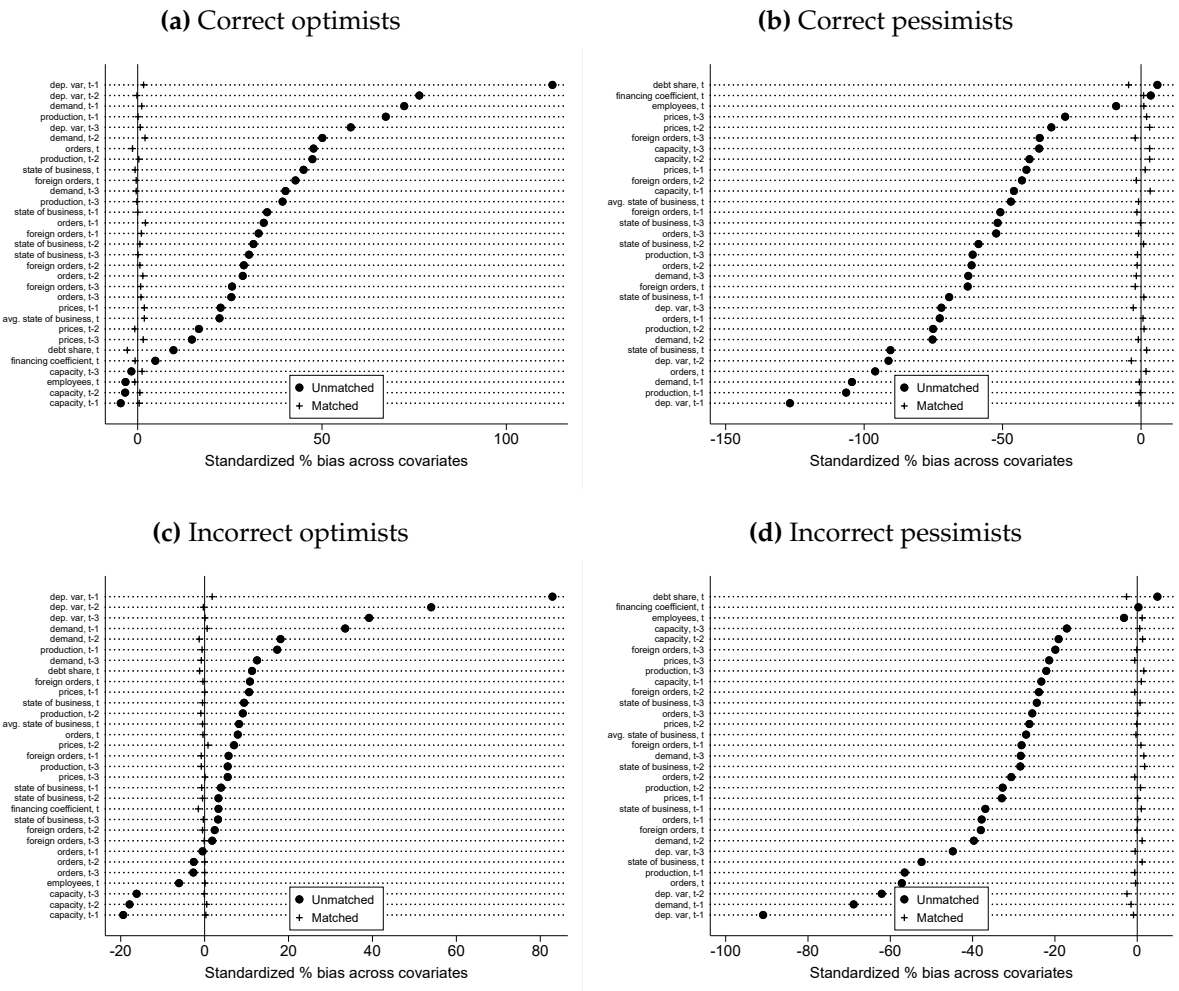

Notes: Figure shows standardized bias for the matching of correct and incorrect optimists and pessimists. The standardized bias measures the mean difference of each variable in the treated and untreated groups, as described by equation (3.4) in Subsection 3.3.2.

Figure C.5: Variance ratio of residuals, before and after matching, correct and incorrect firms

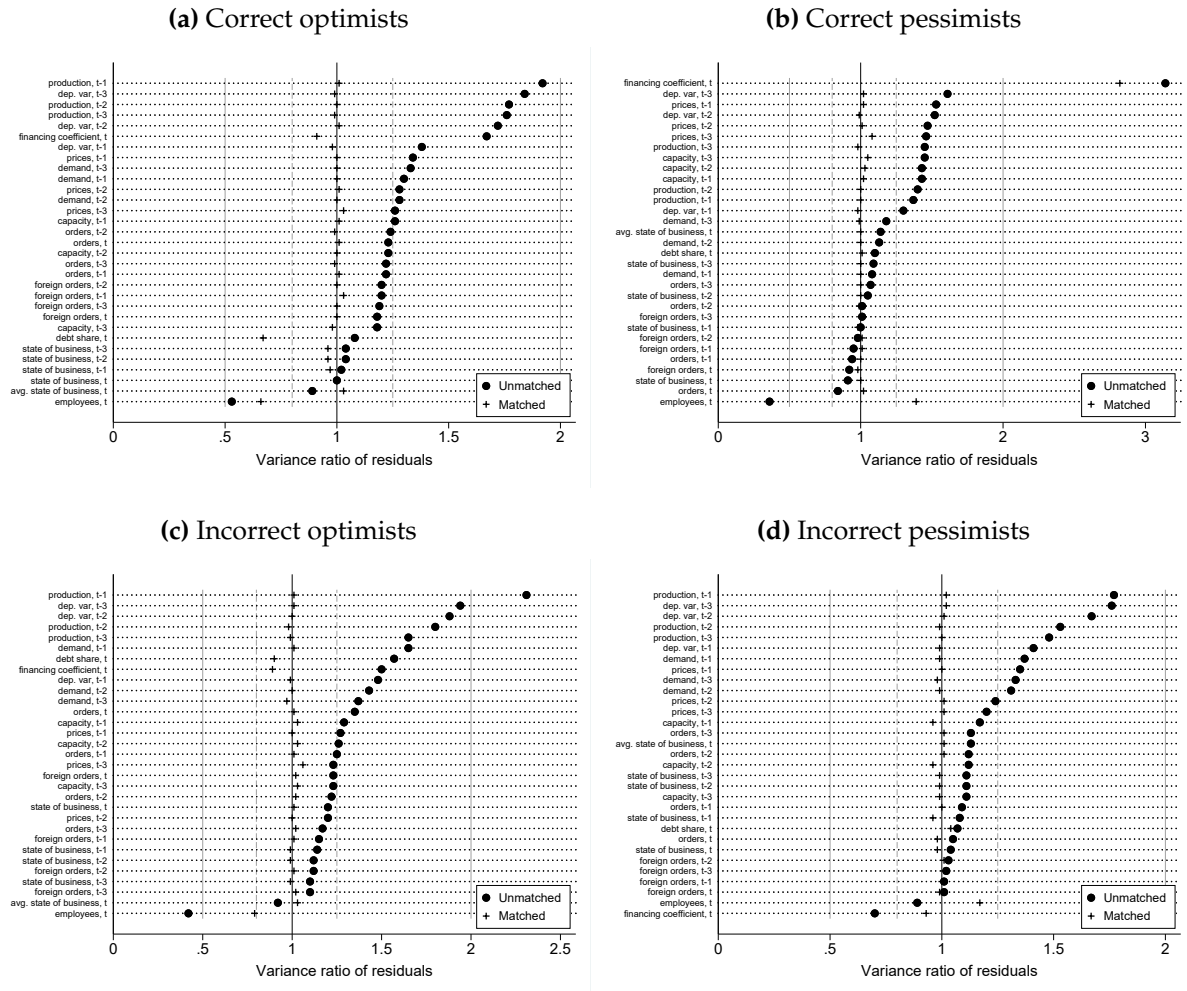

Notes: Figure shows variance ratios for the matching of correct and incorrect optimists and pessimists. The variance ratio measures the difference between the variances orthogonal to the propensity score. Details on this measure can be found in Subsection 3.3.2. Variance ratios below 0.8 and above 1.25 (dashed lines) are considered "of concern"; ratios below 0.5 and above 2 (grey solid lines) are considered "bad" according to Rubin (2001).

C4 Sensitivity analysis for Sections 3.3 and 3.4

Table C.11: Aggregate results with alternative use of balance sheet data

	Optimists		Pessimists	
	(1)	(2)	(3)	(4)
Dep. variable:	Production	Prices	Production	Prices
ATT	0.170***	0.026***	-0.175***	-0.029***
	(28.85)	(5.82)	(-27.23)	(-5.57)
Observations	120754	120802	116470	116548

Notes: T-statistics in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

C4.1 Alternative use of balance sheet data

As discussed in Section 3.2, the survey data has a different frequency than the balance sheet data. In our baseline setting, we use the most recently published balance sheet data to estimate the debt share and financing coefficient in each month. This implies that in the months before a new balance sheet is published, we use information which is almost one year old. We use this approach to avoid including any information, which is not yet available to firms at the time expectations are formed. However, one may argue that firms become aware of changing fundamentals already ahead of the publication of the new balance sheet. We therefore now propose an alternative method to link the annual balance sheet data to the monthly survey data. Specifically, for the six months following the publication of the balance sheet, we use the most recent report as before. However, for the next six months until the new balance sheet is published, we use the new data. This means that we always use the balance sheet data with the publication date closest to the respective month.

Table C.11 shows that changing the method for allocating the balance sheet data barely affects the results. Given that we only use two balance sheet variables in the probit regressions determining the propensity score this is not very surprising. Nevertheless, it is reassuring that our estimation is robust in this regard.

C4.2 Alternative matching method

In order to ensure our results are not affected by our choice of matching algorithm, we implement an alternative algorithm as described in Lechner et al. (2011). These authors propose a radius (or caliper) matching procedure, which includes weighting proportional to the distance of the match and a bias adjustment.

Table C.12: Aggregate results with alternative matching procedure

	Optimists		Pessimists	
	(1) No bias corr.	(2) Bias corr.	(3) No bias corr.	(4) Bias corr.
Panel (a): Production (change in current month)				
ATT	0.172***	0.172***	-0.174***	-0.174***
	(0.007)	(0.007)	(0.007)	(0.007)
Observations	135170	135170	131656	131656
Panel (b): Prices (change in current month)				
ATT	0.027***	0.027***	-0.037***	-0.037***
	(0.005)	(0.005)	(0.006)	(0.006)
Observations	135170	135170	131656	131656

Notes: Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Specifically, the algorithm first selects all nearest neighbors in terms of the propensity score and other variables (in the latter case using the Mahalanobis distance) without replacement. In our case, we use the propensity score from the simple probit regressions described in Subsection 3.3.1 and the month as an additional variable. The latter is done to ensure comparability to our matching procedure. In a next step the radius is computed as a function of the maximum distance within a matched pair in step one. Using this radius additional matches are selected if they are within the radius around the respective observation. This matching step is done with replacement, i.e. untreated observations can be matched to different treated observations. Weights are computed as the inverse of the distance between the untreated and treated observations in a match.

Finally, a regression bias adjustment is implemented by regressing the outcome variable on an intercept, the propensity score, the square of the propensity score, and any further variables used to define the distance. The regression is done only for the matched untreated observations using the weights obtained from matching. Using the regression coefficient, one then predicts the potential outcome under no treatment for all observations. The difference between the weighted mean of the predicted outcome in the untreated group and the mean of the predicted potential outcome in the treated group is the estimated bias. This bias is then subtracted from the estimated ATT. The variance is computed analytically.

This approach differs from our matching algorithm because the radius is determined endogenously, the weights are proportional to the distance, matches can be from different months (albeit only from close months because we include the month as an additional distance measure), and finally there is regression adjustment. We imple-

ment this procedure using the STATA code provided by Huber, Lechner, and Steinmayr (2015). For simplicity we use their default settings. The results can be found in Table C.12. Using this alternative matching procedure does not affect our results substantially. Compared to our baseline specification in column (1) of Table 3.4 results only differ at the third digit. The largest difference is observed for prices of pessimists: -0.036 compared to -0.031 in the baseline. Reassuringly the bias adjustment also does not have any effects up to three digits. This implies that using a more simple matching procedure with no bias correction is valid in our data set.

Chapter 4

Exportweltmeister: The Low Returns on Germany's Capital Exports [*†]

[*]This chapter is based on joint work with Moritz Schularick (University of Bonn) and Christoph Trebesch (Kiel Institute for the World Economy). Earlier versions are available as CEPR Discussion Paper No. DP13863 and as Kiel Working Paper No. 2133.

[†]We thank Rainer Frey, Stefan Goldbach, Ulrich Grosch, Axel Jochem, Arne Nagengast, Robert Unger and Sarah Peters from the Bundesbank for support regarding the data, in particular Ulrich Grosch for facilitating early access to detailed data on valuation changes. We thank Gerhard Schick for helpful comments, and Isabel Ahlers for proof-reading.

4.1 Introduction

In a famous scene in Michael Lewis' book "The Big Short", a senior executive at Deutsche Bank, Greg Lippmann, tours Wall Street in 2007 to short-sell securities containing US subprime mortgages. When asked who was still buying these toxic high-risk papers, "he always just said: Düsseldorf" (Lewis 2010, p. 67). As a matter of fact, Düsseldorf based IKB Bank was one of the first banks to fail in 2007. Via its "Rhinebridge" investment vehicle, IKB Bank heavily invested in the US subprime mortgage market. IKB Bank subsequently had to be rescued by the German government. These and other anecdotes have tainted the reputation of German investors in global markets, as the notion of "stupid German money" indicates.

Is there economic truth in this caricature? Are German investment returns particularly low, and if so, why? From the point of view of German savers and economic policymakers, this is a first order question. Germany today is the world's foremost exporter of savings. More than 300 billion Euros of German savings are sent abroad every year. Despite the heavy losses on American and other investments in the 2008 crisis, Germany has exported 2.7 trillion Euros in the past decade alone, equivalent to about 70% of GDP. These capital outflows came from German banks, firms, and households. Unlike in China or Japan, the public sector has played only a secondary role in the build-up of Germany's foreign asset stock, despite the Bundesbank's much-debated Target2 claims within the Eurosystem.[1]

The German public debate has generally interpreted the high current account surplus as a demonstration of the competitiveness of German industry, and not as an indication of insufficient domestic demand (see e.g., Fuest 2017; German Council of Economic Experts 2017; Schuknecht 2014; Sinn 2017). Various voices in the domestic debate view Germany's capital outflows favorably. Often-heard arguments include the potential for international risk sharing and for providing a hedge against adverse demographic trends, in line with the traditional textbook view.[2] An aging country like Germany, so the argument goes, can benefit from investing abroad and achieve better investment returns in younger, more dynamic economies than at home (see e.g., Deutsche Bank 2013). Storing wealth abroad will also improve risk sharing and consumption insurance: when a recession occurs in Germany but not in other countries, capital income transfers from abroad stabilize domestic income (see e.g., Bundesbank 2018).

Internationally, the country's large-scale capital exports have drawn more criticism. They are sometimes seen as problematic for the recovery in the euro area, and possibly

[1] In 2017, Target2 claims accounted for about 10% of Germany's total external assets.

[2] See e.g., Taylor and Williamson (1994) and Obstfeld and Rogoff (1996).

as a driver of credit and asset-price bubbles abroad (see e.g., Bernanke 2015; den Haan, Ellison, Ilzetzki, McMahon, and Reis 2016; European Commission 2016; IMF 2016; Krugman 2017). While German economists and institutions used to reject this criticism (see e.g., German Ministry of Finance 2017; Weidmann 2014), the debate has recently shifted. More voices are now questioning the size of Germany's current account surplus, which continues to exceed 7% of GDP per year (see e.g., Board of Academic Advisors 2019; Weidmann 2019). A key reason is that the economic consequences of low domestic investment have become difficult to overlook. In an ironic reference to IKB Bank's failed "Rhinebridge" investment vehicle, the deteriorating condition of actual bridges over the Rhine has become a symbol of crumbling infrastructure and growing domestic investment needs.

This study presents a comprehensive empirical assessment of Germany's investments abroad over the entire postwar period. We estimate investment returns and valuation changes across seven decades and compare Germany's returns abroad with those of 12 other advanced economies as well as the returns on domestic investment. We also assess the consumption insurance offered by foreign returns, and their role as a hedge for demographic risks. A central contribution is the international perspective, which we pursue throughout the analysis.

Since all economies have access to the same investment opportunities abroad, it is informative to compare the returns on foreign assets of one country to those of another. We ask, how large are the returns on foreign investment of each nation in international capital markets, and how does Germany compare? In other words, we focus on returns on external assets, not on the difference between returns on assets and liabilities. This is the relevant question to assess investment performance, as the liability structure of countries is a function of the investment decisions of others. Moreover, the returns on liabilities can be distorted due to tax shifting and country-specific effects that make comparisons difficult.[3]

Our analysis brings several new insights that break rank with the consensus view in Germany. First, we find that the returns on German foreign assets are considerably lower than the returns earned by other countries investing abroad. Since 1975, the average of Germany's annual foreign returns was about 5 percentage points lower than that of the US and close to 3 percentage points lower than the average returns of other European countries. Germany fared particularly bad as an equity investor where investment returns underperformed by 4 percentage points annually.

Second, we find that Germany earns significantly lower foreign returns *within* each

[3]In Germany, the difference between asset returns and liability returns has turned positive in recent years, but not because German foreign investments performed better in terms of yields or valuation gains, but rather because the yield of foreigners investing in Germany went down, partly for idiosyncratic reasons (see Appendix D6).

asset category, after controlling for risk. This suggests that Germany's weak financial performance abroad is not merely the result of a more conservative investment strategy that focuses on safer assets. The low German returns compared to other countries also cannot be explained by exchange rate movements, nor by the recent build-up of Target2 balances. Instead, valuation losses are a big part of the explanation. The valuation of Germany's external asset portfolio has stagnated or decreased, while other countries witnessed considerable capital gains. Germany's frequent investment losses are remarkable given that the world economy has witnessed a price boom across all major asset markets over the past 30 years (Jordà, Knoll, Kuvshinov, Schularick, and Taylor 2019).

Third, German returns on foreign assets were considerably lower than the returns on domestic assets. This is an important insight for the policy debate on the merits of domestic versus foreign investment. The difference was particularly pronounced in the last decade, when the average return on a domestic portfolio of German bonds, equity, and real estate was about 4 percentage points higher per year than the returns on Germany's foreign assets.[4] German capital exports surged in the past decade, but the better returns were possible home. Clearly, while it is possible that domestic returns were to some extent pushed up by capital exports, negative real interest rates on a wide range of assets provide some *prima facie* argument against a scarcity premium on domestic capital.

Fourth, we find little evidence that foreign returns have positive effects for consumption insurance. The return on Germany's external assets is highly correlated with German economic activity – even more so than domestic returns – and thus provides no hedge against domestic consumption shocks. Moreover, 70% of Germany's foreign assets are invested in other advanced economies that face similar demographic risks. In the past decade, less than 10% of capital flows went to younger, more dynamic economies outside of Europe or North America, even though emerging markets now account for more than 50% of world GDP.

Table 4.1 summarizes the main findings of the study. The table ranks countries by their average return on foreign assets, using all countries for which we have sufficiently detailed data (see Section 4.2 for details and methodology). Germany has the worst investment performance among the G7 countries. In the full country sample from 1975 to 2017, Germany ranks 12th, with only Finland performing worse. The picture looks similar if we consider the past decade (2009-2017), where Germany ranks 10th. The same is true when we use real returns, deflating each country's foreign asset returns with domestic inflation rates (see Table D.1, where Germany ranks 9th).

[4]Data on the domestic return is from Jordà et al. (2019). The return is computed as a weighted average of return on equity, housing, bonds and bills. Weights are stock market capitalization, housing wealth, and public debt (split half-half between bonds and bills).

Table 4.1: Returns on foreign assets, 1975 to 2017

	Rank	1975–2017	1999–2017	2009–2017
United States	1	10.64	8.00	9.27
United Kingdom	2	10.22	5.68	4.09
Canada	3	9.19	4.93	8.98
Sweden	4	8.99	6.45	5.38
Norway	5	8.00	5.89	7.38
Italy	6	7.96	3.31	4.39
Spain	7	7.91	2.09	2.89
France	8	7.38	4.01	4.43
Netherlands	9	6.65	4.82	6.05
Denmark	10	5.32	5.32	6.98
Portugal	11	5.01	2.71	2.20
Germany	12	4.95	3.73	3.78
Finland	13	4.43	3.90	3.39

Notes: This table shows average, nominal returns on foreign assets for various time samples. Countries are ranked by their average return. We compare nominal returns in domestic currency to abstract from different national inflation dynamics. The ranking is similar with real returns (see Appendix D2). Data for Denmark and Portugal starts in 1999 and 1993, respectively. No data is available for Japan. For details, see Section 4.2.

The cumulative effects of these bad investment returns are quantitatively large, as can be illustrated with a simple counterfactual exercise. In the decade since the 2008 financial crisis alone, Germany could have become about 2 to 3 trillion Euros richer had its returns in global markets corresponded to those earned by Norway or Canada, respectively. This implies a (hypothetical) wealth loss of 70% to 95% of German GDP (see Section 4.5 for details). On a per capita basis, this implies an amount of about 27,000 to 37,000 Euros of foregone wealth for each German citizen (compared to the performance of Norway and Canada).[5] These numbers are only an illustrative thought experiment, but they highlight the economic relevance of return differentials on foreign investments.

The large cumulative effects are also evident in Figure 4.1, which compares the total return performance of foreign asset position of Germany, the US and the UK, as well as a portfolio of domestic German assets (stocks, bonds and houses). Assume you invested 1 Euro in global capital markets in 1975 and that you reinvest any dividends or interest gains. As of 2017, you would own 40 to 60 times of that initial investment had you followed the investment strategy of the UK or the US. In comparison, the initial investment only increased by a factor of 7 using the returns on German foreign

[5]Numbers are based on the German population in 2017, which was 82.8 million according to the German Statistical Office.

Figure 4.1: Cumulated nominal returns, 1975 to 2017

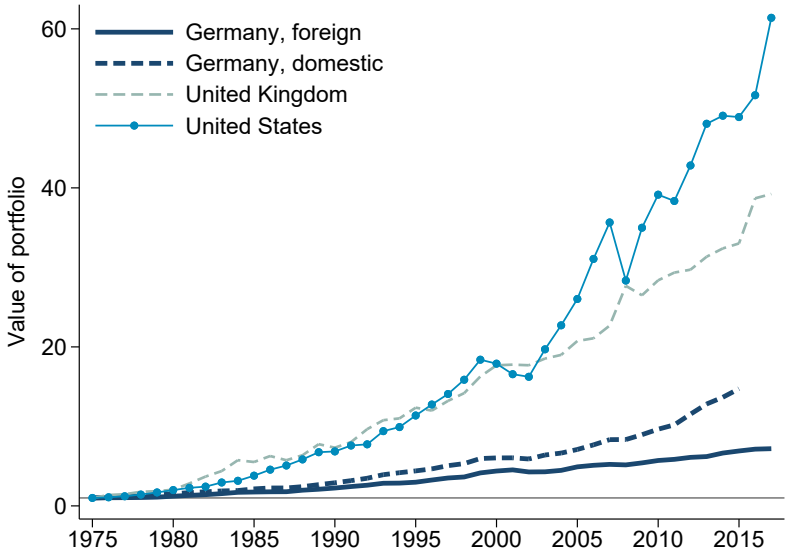

Notes: This graph shows cumulated total returns since 1975 for a portfolio with an initial value of 1. We focus on foreign nominal returns of the US, the UK and Germany, as well as the German domestic return (data until 2015 from Jordà et al. (2019)). Returns are cumulated over the years using the following formula: $\prod_{i=0}^{t}(1+r_i)$. For details, see Section 4.2.

assets (before inflation). German domestic assets increased to 14 times of the initial investment.

This chapter contributes to the existing literature in several ways. First, it compares the foreign investment returns of 13 nations in a systematic way. Much of the literature focuses on individual countries and typically compares the returns on assets abroad to those earned by foreigners on inward investments, i.e., on liabilities (for a survey see, e.g. Gourinchas and Rey 2014). A main motivation of that strand of research is to examine the size of the "exorbitant privilege", referring to the phenomenon that some countries, mainly the United States, can borrow at low yields from abroad and reinvest these into higher-yielding assets internationally (Curcuru, Dvorak, and Warnock 2008, 2013; Gourinchas and Rey 2007a; Lane and Milesi-Ferretti 2003; Meissner and Taylor 2008). Here, we take a broader perspective and benchmark returns across countries, decompose these returns, and examine their determinants in a sample spanning multiple decades. This adds to a small literature that compares international returns (Habib 2010; Lane and Milesi-Ferretti 2003, 2007a). We apply a broad range of tests across countries that so far have only been employed to analyze return differentials of individual countries. In particular, using newly gathered data and new estimates for our 13-country sample, we compute how valuation changes due to exchange rates, the asset composition, and the geographical distribution of foreign investments affect

returns.[6]

Second, this chapter is the most comprehensive analysis of the returns of Germany's external assets, going as far back as the 1950s, and by applying a similar degree of rigor as influential studies conducted for the United States. Our basic methodological approach resembles that of the Bundesbank (2014, 2018). We also get similar return estimates for the overlapping sample period from 2005 to 2017 (see Appendix D3 for a detailed comparison with these and other studies focusing on Germany). Our contribution is to put these numbers into perspective by benchmarking the investment performance to that of similar economies, by adding five decades of data, and by studying *why* German returns are so low. The results reveal how badly German investments have done in international comparison, including in the past decade. This in contrast to earlier work on German foreign investment which came to different conclusions about the profitability by only comparing German foreign investment with foreigners' investment in Germany. Furthermore, our findings on risk sharing and demographic hedging question two of the most common arguments in support of Germany's large capital outflows.

The structure of this chapter is as follows. We first guide the reader through some technical but important preliminaries about data, balance of payment arithmetic and methodology used. The next section presents long run series for returns, yields, and valuation changes on German foreign investments. We then compare returns on German foreign assets to returns on the foreign assets of other countries and to domestic German assets. The last section looks at the performance of foreign assets as a hedge for domestic consumption and demographic risks.

4.2 Data and definitions

This section gives an overview of the main data and definitions used. We start by describing the method for computing returns and the classification of investment types, then discuss the data for Germany, and move on to international data.

4.2.1 Return computation

We compute the aggregate domestic currency return as the sum of investment income, II_t, and aggregate valuation changes, VC_t, over the stock of assets at the beginning of year t:

$$\widetilde{r}_t^A = \frac{II_t^A + VC_t^A}{IIP_{t-1}^A} = \frac{II_t^A}{IIP_{t-1}^A} + \frac{VC_t^A}{IIP_{t-1}^A} = \widetilde{i}_t^A + \widetilde{vc}_t^A,$$

[6] Bénétrix, Lane, and Shambaugh (2015) also analyze a large set of countries but they mainly focus on valuation changes due to exchange rates and not on their effects on returns.

where the superscript A indicates the asset side of the economy, i.e. assets owned abroad, financial account outflows and income earned by German residents. This is the standard approach in the literature (see e.g., Habib 2010). We transform the three measures of returns to real values using consumer price inflation π_t:

$$
\begin{aligned}
i_t^A &= \frac{1+\widetilde{i}_t^A}{1+\pi_t} - 1 && \text{(real yield)} \\
vc_t^A &= \frac{1+\widetilde{vc}_t^A}{1+\pi_t} - 1 && \text{(real valuation changes)} \\
r_t^A &= \frac{1+\widetilde{i}_t^A+\widetilde{vc}_t^A}{1+\pi_t} - 1 = i_t^A + vc_t^A + \frac{\pi_t}{1+\pi_t} && \text{(real return).}
\end{aligned}
$$

Following the residual approach, we compute aggregate valuation changes as all changes in the asset position not due to transactions in the financial account:

$$VC_t^A = IIP_t^A - IIP_{t-1}^A - FA_t^A = VX_t^A + VP_t^A + VOT_t^A.$$

These aggregate valuation changes can be further split into valuation changes due to exchange rate changes, VX_t, changes due to price changes, VP_t, and changes due to other adjustments, VOT_t, which include write-offs and permanent losses but also residuals due to statistical discrepancies. The latter result mainly from changes in primary data sources and differences between in the primary data sources used for the stock and flow series. In the literature on the exorbitant privilege of the US, much attention focused on the treatment of this other adjustments term, which can be large in some cases (see e.g., Lane and Milesi-Ferretti 2009).

In the United States, the main reason for discrepancies are different revision policies for the IIP and the balance of payments, specifically revisions in asset data are more extensive than revisions in flow data (Curcuru et al. 2008). In Germany, this is not the case. The Bundesbank revises both stock and flow data up to the four previous years, and adjusts both accounts in order to ensure consistency. Therefore, we include other adjustments fully in our valuation gains following authors such as Habib (2010) and Lane and Milesi-Ferretti (2007a). However, we correct for statistical problems that are not accounted for in revisions, as discussed for example by the Bundesbank (2014) and Frey, Grosch, and Lipponer (2014). We explain our adjustments in detail in Subsection 4.2.2. Our procedure ensures that we include relevant permanent losses while, at the same time, not including changes due to purely statistical effects.

Starting in 2005, the Bundesbank provides a breakdown of the valuation changes into the three components (exchange rates, prices, other). Using this data, we show in Appendix D3.2 that different allocations of the residual term only have minor effects on the overall return. We take this as evidence that our residual approach with additional

adjustments is appropriate for the study of long-run return developments in Germany.

Following standard practice, we will focus on four broad categories of foreign investments:

(1) *Foreign direct investment (FDI)* is any kind of foreign investment associated with control or significant influence over a foreign affiliate. This category also includes any additional investment associated with the foreign direct investment relationship, including reverse investment. Furthermore, real estate investments typically fall into the FDI category.

(2) *Portfolio investment* is further split into debt and equity investment, where debt investment refers to bonds of any kind and equity refers to any direct claims not classified as FDI. Furthermore, investment fund shares are combined with the equity part of portfolio investment. In our analysis, we will usually consider debt and equity (incl. investment fund shares) separately.

(3) *Other investment* is a combination of various additional investment categories. It mainly covers financial loans, trade credit and advances, currency, and deposits. In addition, 'other investment' includes some residual 'other equity' as well as claims from pension entitlements and insurances (the latter only since the recent reform of the Balance of Payments Manual). For euro-area countries, also Target2 balances are included in the 'currency and deposits' subcategory.

(4) *Reserves* refer to any assets held by the central bank for the purpose of monetary operations.

4.2.2 German data

To compute returns, we need data from three different balance of payments accounts published by the Bundesbank: the International Investment Position (IIP), the financial account, and data on primary income from the current account.

The IIP is available since 1949 on the Bundesbank website. The flow accounts data there starts in 1971. We combine these series with data for 1949 to 1970 using a Bundesbank report published in commemoration of the 50th anniversary of Deutsche Mark, which is made available electronically by Histat/GESIS (Bundesbank 1998). The stock data and the recent flow data have been revised to match the requirements of the sixth edition of the IMF's Balance of Payments Manual (BPM).[7] The older flow data is still based on the previous edition of the BPM and it is denominated in Deutsche Mark (DM).

[7]The BPM5 was introduced in 1993 and implemented by the Bundesbank in 1995 for the current account and in 1998 for the IIP. The new BPM6 was introduced in 2009. Using the underlying data sources the Bundesbank was able to revise the old data in line with the most recent manual, which is very helpful for our purposes.

A main challenge in this context was to ensure that we have consistent time series for each asset category over the seven decades we study, for which we adapt the old (pre-1971) data to make it compatible to the newer data. The easiest part was to convert all old series into Euro using the fixed conversion rate (1.95583 $\frac{DM}{€}$). In addition, there were several changes in the BPM that relate to the classification of asset classes and their subcategories, which we deal with as follows.

First, we combine 'loans' and 'other investment' in the historical data into an aggregate category to make it consistent with the modern (BMP6) classification, which combines loans, currency, deposits and other investment activities under the label 'other investment'. The historical series, we created includes interest income from loans but not from 'other investment', since this mainly constitutes government stakes in international organizations.[8]

Second, in the old data, financial account flows of the central bank were recorded in a separate account. The data from this central bank account is similar to that in today's central bank flows in the financial account (we can compare the old and new series for a lengthy overlapping period, namely 1971-1997). Therefore, we use the data from the old Bundesbank account to measure financial account flows by the central bank in the period before 1971.

Third, in the modern data, portfolio debt investments are divided into long-term and short-term bonds, while there is no such distinction in the old data, where this category is labeled as 'fixed income assets'.[9] Since these terms refer to the same asset class, we merge the old and new series and rename them "portfolio debt". This category captures both long-term and short-term bonds in history and today.

Fourth, we face the problem that reserve assets are a distinct category in the IIP and the financial account, while reserves are combined with 'other investment' in the primary income account.[10] As a result, we merge the two series so that returns can only be computed for the sum of 'other investment and reserves'.

Fifth, portfolio equity investment includes the subcategory 'investment fund shares' but data availability and reliability of this series is limited historically. The subcategory was fully incorporated in the German IIP only in 1994, but the estimation of the liability position is noisy and imprecise until the year 2009 (Frey et al. 2014). Therefore, we only compute returns from investment fund shares from 2010 onward and exclude this series until 2009. Since investment fund shares make up no more than 6% of total

[8] See annotation in Table B6_07 in Bundesbank (1998).

[9] More specifically, in German the category in the new data is called "Schuldverschreibungen" whereas it used to be "festverzinsliche Wertpapiere".

[10] In the BPM6, countries are left the choice whether to show the income flows separately or combine them with 'other investment'. In the old manual countries could choose to include the income flows either in 'other investment' or portfolio investment. The Bundesbank chose to combine the reserve asset income with 'other investment' income, such that there is a consistent time series available.

IIP assets and less than 1% of IIP liabilities before 2009, this does not affect our results much.[11] We face similar problems for the category of FDI debt, meaning loans that are part of foreign direct investment flows. The IIP series on this subcategory only starts in 1997 while it was included in the other accounts earlier on. We therefore exclude the data on FDI loans prior to 1997 to ensure consistent return series.

Sixth, the Bundesbank added financial derivatives to the IIP in 2010, as this category had become increasingly relevant. Financial derivatives are only recorded as balances in the financial account and not included in primary income, making it difficult to compute precise returns. Therefore, we subtract financial derivatives from aggregate quantities and do not consider them in our main return calculations. However, we do show some stylized facts on the amount of outstanding derivatives in Section 4.3. More generally, whenever any further new category is added to the IIP, we subtract the increase due to this addition from the change in assets in the given year to avoid overestimating valuation gains.

Beyond the categorization of asset classes, we need to consider idiosyncratic breaks in the data series, changes in data availability, as well as mismatches between the three different accounts, as also discussed by Bundesbank (2014) and Frey et al. (2014).

First, the initial values in 1954 for IIP liability categories equity investment and debt instruments within portfolio investment are unrealistically low, compared to values in the following period, resulting in double or triple digit returns in the following year. We exclude these values as outliers.

Second, the Bundesbank changed the valuation of its reserves and other external holdings after the introduction of the Euro in 1999. Before 1999, the reserves were valued using the "lower of cost or market" concept, while they are valued at market prices afterwards. According to Bundesbank (2012), this resulted in a €26.25 billion jump in the reported value of reserves and other assets held by the Bundesbank in 1999 (of these €25.42 billion are FX reserves and €0.83 billion are other assets). To correct for this one-time change, we subtract this increase from the change in the relevant IIP asset categories in 1999.

Third, there are issues with regard to market versus book valuations. For our purpose of computing investment returns, it is crucial to value assets at market prices. This is particularly challenging for the valuation of FDI. For listed companies, FDI equity is valued at market prices by the Bundesbank since 2004, but not before. For non-listed companies, the Bundesbank uses the values reported in the parent company's balance sheet, as is standard practice in many countries and is also recommended in the BPM6

[11] The aggregate return on foreign assets is barely affected. The return on equity is lower if we include investment fund shares since their return is lower than the return on equity. Between 1994 and 2009, the inclusion of investment fund shares lowers the return from 8.71% to 6.1%. Hence, if anything, ignoring this sub-category will result in an overestimation of German returns (upward bias).

(Bundesbank 2008). Moreover, no market prices are used for other, smaller components of FDI assets, for example assets related to construction sites. Only real estate assets have always been valued at market prices (Bundesbank 2008). The lack of market values in some parts of the FDI data may lead to an underestimation of returns, which is particularly problematic if Germany uses a different valuation approach than other countries. We explore how our results may be biased due to FDI valuation issues in Appendix D8, concluding that the effects are small.

Valuation issues are less relevant for the remaining asset categories. The Bundesbank has always reported portfolio investments at market prices.[12] For reserves, the Bundesbank provides market-based values since 1999. Before 1999, it applied the lowest value accounting principle, assigning the minimum of market value and original (purchasing) costs. For loans, deposits, and currency valuation changes are secondary, except for exchange rate effects, which we consider throughout our analysis.

Taken together, these adjustments allow us to compute consistent time series of returns using primary income, financial account flows, and asset stocks for the asset categories of foreign direct investment, portfolio debt investment, portfolio equity investment, and 'other investment' (including reserves) starting in 1949.

Beyond data on German assets and liabilities, we use data on the German price level and GDP from the Macro History Database (MHD) (Jordà, Schularick, and Taylor 2017). MHD data is only available until 2016, so we append data from German Federal Statistics Office (GDP) and Eurostat (Harmonized Index of Consumer Prices (HICP)) for 2017. For GDP, we do this by using official data levels and applying the growth rates from the MHD data.

4.2.3 International data

In order to place German returns in a broader context, we compare them to other countries' returns on their foreign assets. We compute the returns of other countries as discussed above and thus require the same type of data. The main data sources are the IMF's balance of payments and international investment position statistics.

The time series generally start in 1970 but there are differences across countries. Data on assets is the most limited when drawing on national sources. Therefore, we add data on assets from the External Wealth of Nations database (EWN) of Lane and Milesi-Ferretti (2007b). We follow their recommendation regarding the starting year of when to use IMF data and when to use their estimates. This is mainly due to part of the older time series from the IMF still being book-value series. Lane and Milesi-Ferretti (2007b) on the other hand, provide estimates of market-value positions. Due

[12]The valuation used to be done by using price indices (Bundesbank 2008). Since 2006 individual securities can be tracked and valued to provide an even better estimate of foreign assets.

to the relevance of valuation changes for our returns, we rely on their estimates. This allows us to compute returns for 12 additional countries. For five countries, the return series start in 1971, for an additional five countries the series starts at the latest in 1976. We provide details on countries and time spans covered in Table D.3 in Appendix D4. To ensure comparability, we also exclude data on financial derivatives from the other countries' returns. In addition, we also check the consistency of the data sources as we do for Germany and adjust accordingly. These country-specific data issues are also described in Table D.3.

To compute real returns, we use data on inflation from the World Bank's Word Development Indicators. Furthermore, we also use nominal GDP data from the World Development Indicators database in the regression analyses.

One important influence on returns are valuation changes due to exchange rates. Unfortunately, exchange-rate specific valuation changes are only published scarcely by some countries and there is no readily available dataset across countries, especially not by asset class. Therefore, we estimate exchange rate driven valuation change for each investment category in our sample as discussed in more detail in Subsection 4.6.1. For this purpose, we need data on the currency composition of assets, but such data is also not readily available. Instead, we follow the suggestions of Bénétrix et al. (2015)[13] to approximate the currency composition. For each asset category, we use different external sources. Table D.4 in the appendix provides details on the data used for each country and asset type. We also discuss the data choices in the following.

For FDI assets, we use OECD data on bilateral FDI stocks, which starts in 1985 for most countries and covers a large set of partner countries.[14] Following Bénétrix et al. (2015), we assume that FDI in a country is always denominated in the local currency.

To estimate the currency shares of portfolio investment, we rely on data from the IMF's Coordinated Portfolio Investment Survey (CPIS). The survey collects data on the cross-border holdings of portfolio equity and debt assets starting in 2001. The holdings are broken down by country pairing and by currency. The currency breakdown includes US Dollars, British Pound, Euro, Japanese Yen, and Swiss Franc. Not all countries provide a currency breakdown since 2001. In those cases, we use information from the country breakdown to estimate currency shares. We adjust the country shares by the average ratio of country to currency shares when both are available since

[13]This is an update and extension of Lane and Shambaugh (2010), which describes the process in some more detail.

[14]Bénétrix et al. (2015) suggest using the UNCTAD database. However, we find for our sample the OECD database is more useful. The UNCTAD data only covers the years 2001-2012. Its advantage is that it covers a large set of reporting countries, and that it includes the ultimate counterparts instead of the immediate target of the investments. This is relevant especially for the investment of large multinational corporations. The former is not relevant to us since we focus only on advanced economies. The latter is unlikely to affect the currency composition in a major way.

there is no one-for-one match between the two (especially for the US Dollar). Here, we again deviate from Bénétrix et al. (2015) who only use the country data available.

'Other investment' mainly comprises of loans and deposits by banks, and Bénétrix et al. (2015) suggest using the Locational Banking Statistics (LBS) provided by the Bank for International Settlements (BIS). This data set covers cross-border lending by banks in US Dollars, British Pound, Euro, Japanese Yen, and Swiss Franc. We use these currency shares as an estimate for the currency share of 'other investment'. Data is available from 1977 onward.[15] However, for some countries data availability is limited, see Table D.4 in Appendix D4.[16]

Reserve assets also include a significant share of foreign currency assets. Therefore, we gathered balance sheet and annual report data of each of the national central banks. In some cases, only approximate shares are reported (e.g. "more than 90%"). In these cases, we resort to IMF data on reserve positions and apply the reported shares. Furthermore, we use IMF data on special drawing rights (SDRs) which are also subject to valuation changes. Again, we provide details on coverage and sources by country in Table D.4 in the appendix.

To validate our approach on exchange-rate driven valuation changes, we make use of the fact that some countries have started to publish a breakdown of IIP valuation changes in recent years, albeit usually with limited time coverage. Among these countries are Germany, the Netherlands, Portugal, Spain, the United Kingdom, and the United States. Whenever official time series are available, we use these. Otherwise, we rely on our own estimates.[17] Using these data, we can also show that our estimates are similar to the official time series, see Appendix D5 for the results. The similarity highlights that our approach works well. In the same appendix, we also compare our estimates to estimates using the currency shares of Bénétrix et al. (2015). These estimates are also very close to ours.[18]

As mentioned earlier, we also study the geographical distribution of assets, i.e., the countries where the assets are held. For this purpose, we rely on the same data sources as for the currency composition estimates. This is possible because the IMF's CPIS and

[15] Bénétrix et al. (2015) report that they have access to more detailed data directly from the BIS, potentially covering the gaps in the officially reported series

[16] A notable case are the United States for which the currency breakdown is only available from 2012 on with exception of one data point in 1998. We use this fact to linearly interpolate between 1998 and 2012 to increase data coverage. We check whether this biases the US data by comparing the resulting series on valuation changes to the valuation changes published by the Bureau of Economic Analysis (BEA) since 2002 and find only small deviations (see Appendix D5 for details).

[17] We do not include the UK data since their estimation procedure is less sophisticated and builds on less detailed data than ours, see the Appendix D5 for details on this.

[18] We provide this additional comparison because the currency data of Bénétrix et al. (2015) differs slightly from ours and in some cases is more detailed, as discussed above. However, we cannot rely on their more detailed data for our estimations because they only provide the currency shares for the aggregate asset positions. For our analysis we need the currency shares separately for each asset category.

Figure 4.2: German current account balance in % of GDP, 1872 to 2017

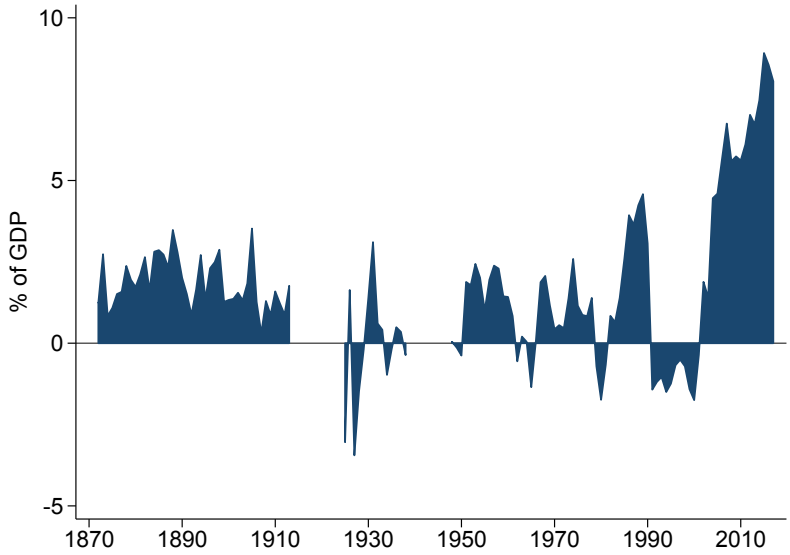

Notes: This figure shows Germany's long history of current account surpluses, which is interrupted only by few periods with deficits, in particular after Germany's reunification in 1990. The past two decades stand out, showing record surpluses both in absolute terms and as a share of GDP. No data is available for 1914–1924 and 1939–1947. Data from the Macro History Database (Jordà et al. 2017) and Bundesbank.

the BIS's LBS databases both include data on the country composition as well. The FDI data refers to countries anyways. We discuss the approach in Subsection 4.3.2. In this case, too, data is not available for all years for all countries. Table D.5 in Appendix D4 provides details.

Finally, we use data on exchange rates from the Bundesbank for Deutsche Mark (until 1998) and Euro (since 1999). All other required bilateral exchange rates are approximated using BIS data on US Dollar exchange rates.

4.3 Germany's capital exports since WW2

4.3.1 Evolution of Germany's current account and foreign assets

Germany has been running current account surpluses for a large part of its modern economic history. Notable exceptions are the late 1920s and the first ten years after the reunification. Figure 4.2 shows that the last decade is characterized by exceptionally high surpluses, even by historical standards. The recent surpluses were about three times higher relative to GDP than in gold standard times and during the so-called economic miracle in the 1950s and 1960s.

Due to its consistently high capital exports, Germany ranks among the world's top

Figure 4.3: Germany's international investment position, 1949 to 2017

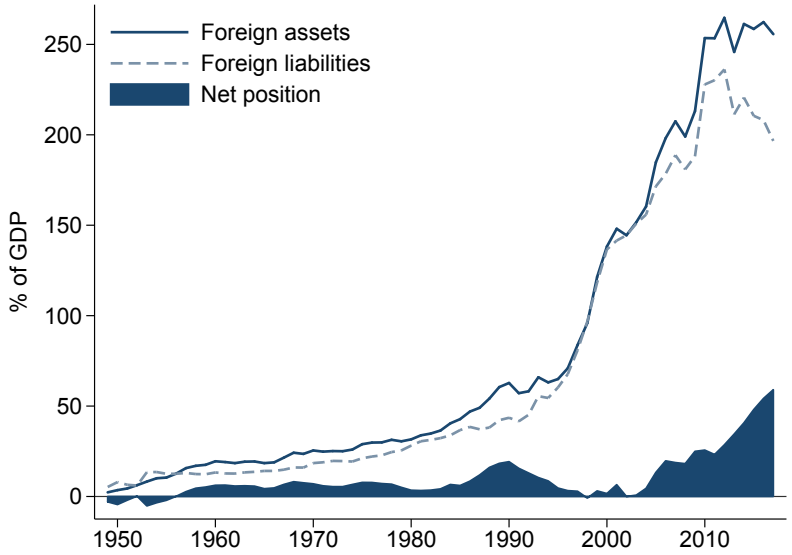

Notes: The net position is the difference between foreign assets (assets held abroad) and foreign liabilities (domestic assets held by foreigners). The graph shows the significant buildup of Germany's gross positions since the 1990s, of both assets and liabilities (financial globalization). The net position has grown most markedly over the last decade (large and sustained current account surpluses). Assets data from Bundesbank. GDP from the Macro History Database (Jordà et al. 2017) and the German Statistical Office.

external creditors, both in absolute numbers and relative to GDP.[19] Furthermore, Figure 4.3 shows that Germany not only has a large net position but also a large gross position. Both the asset and liability positions rose strongly since the mid-1990s and now amount to 256% and 197% of GDP respectively. While they initially grew in tandem leaving the net position at relatively small levels, the gap has been increasing since the mid-2000s and especially in recent years. The net position has been positive over the entire post-war period with few exceptions. It currently stands at 59% of GDP. This reflects Germany's sustained past current account surpluses.

How does Germany's international investment position (IIP) compare to accumulated capital exports? In a simple framework, one can think of Germany's external asset portfolio as a savings account. Adding up all the payments that have flowed into this account correspond to the historical book value of gross investments. The difference between historical costs and market value then reflects valuation gains on that portfolio. In other words, the larger the difference between the accumulated flow measure and the current market value of external investments, the higher the capital

[19] In absolute numbers Germany ranked second in 2017, only exceeded by Japan. In percent of GDP Germany ranked seventh, when excluding the small oil exporting countries. For details see Figure D.1 in Appendix D1.

Figure 4.4: Germany's foreign assets and cumulated financial account outflows, 1949 to 2017

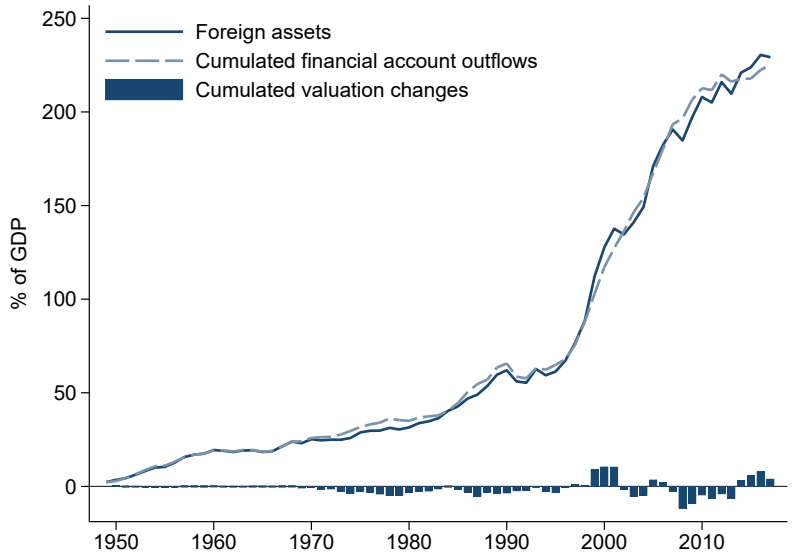

Notes: This graph shows that Germany's cumulated financial outflows (gray dashed line) closely track the stock of total foreign assets (blue solid line). This indicates small valuation gains or even losses on the gross asset position. More specifically, the difference between the two series equals cumulated valuation changes, which are negative with the exception of a few years (blue bars). Foreign assets and cumulated financial account flows are adjusted to remove statistical differences between the series, see Subsection 4.2.2 for details. Financial derivatives are excluded. Data from the Bundesbank.

gains.

Figure 4.4 demonstrates that the value of Germany's gross foreign asset position very closely tracks the cumulated financial account outflows. This implies that the valuation gains, i.e., the wedge between historical flows and current market value, have been small. The blue bars in Figure 4.4 show that this wedge has often even been negative and is generally small. In light of the multi-decade asset price boom that has characterized the world economy in the past decades, this fact is clearly noteworthy (see Jordà et al. 2019).

4.3.2 Germany's external portfolio: asset types and geography

The composition of Germany's foreign assets also changed notably over time. The balance of payments data broadly distinguishes between five different asset categories: foreign direct investment (FDI), portfolio investment, other investment, reserves, and financial derivatives. For Germany, we have data on the first four categories since 1949, but data on financial derivatives only starts in 2010 since this is an investment type that only became relevant more recently. Therefore, we will show derivatives once here and

Figure 4.5: Composition of Germany's international investment position, 1949 to 2017

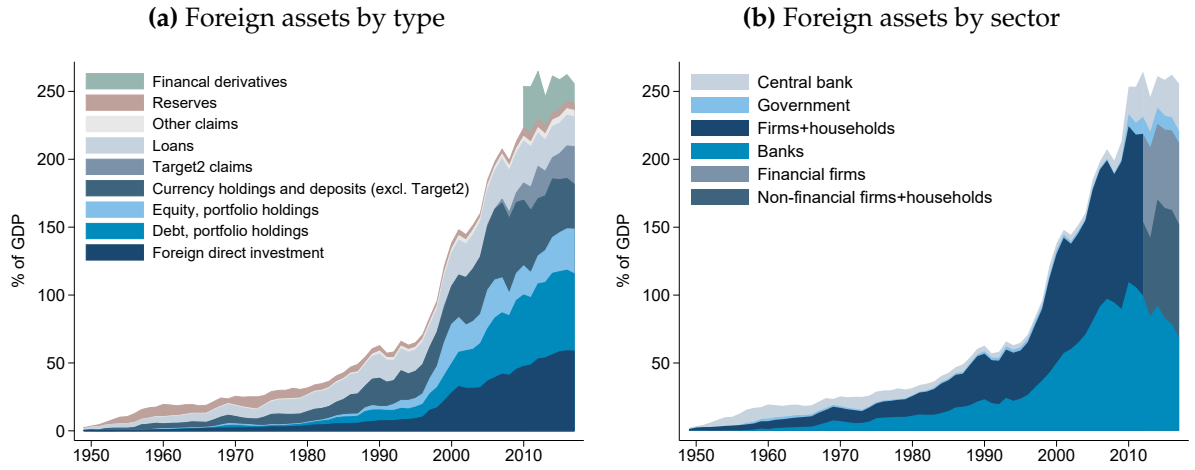

Notes: This graph shows the composition of Germany's foreign assets over time along two dimensions: by type of investment (a) and by sector (b). Asset data from the Bundesbank. The data split between firms and households is only available since 2012. GDP data from the Macro History Database (Jordà et al. 2017) and the German Statistical Office.

exclude them from the remainder of the analysis, as discussed in Section 4.2. This will facilitate the interpretation of developments over time.

Panel (a) of Figure 4.5 shows the changing composition regarding asset classes over time. The rise in the overall level in assets was largely driven by increases in foreign direct investment and portfolio investment reflecting increasing international financial integration. Reserve assets on the other hand made up 20-30% of all assets until the 1970s but have become almost irrelevant today. Target2 balances have been increasing in recent years but only represent about 10% of all assets. As Target2 balances do not generate income, they could potentially bias our estimated downwards. Throughout the analysis, we will pay close attention that our findings are unaffected by the inclusion of Target2 balances.

In addition to the composition by functional category, one can also decompose the foreign asset position by domestic sectors. Here, the balance of payments distinguishes between four broad sectors: banks, firms and households, the government, and the central bank. In more recent data, the non-bank private sector is further split into financial firms and non-financial firms plus households. Panel (b) of Figure 4.5 shows the changing composition by sector over time. The panel shows that the increase in gross position since the 1990s was mainly driven by banks increasing their exposure relative to GDP. However, since the financial crisis the banking sector reduced its exposure. This decline has been partially offset by non-financial firms.

It is equally interesting to consider the geographical distribution of assets. Unfortunately, no official data on the country of residence of the counterparties are available.

Figure 4.6: Geographical distribution of IIP assets, 1985 to 2017

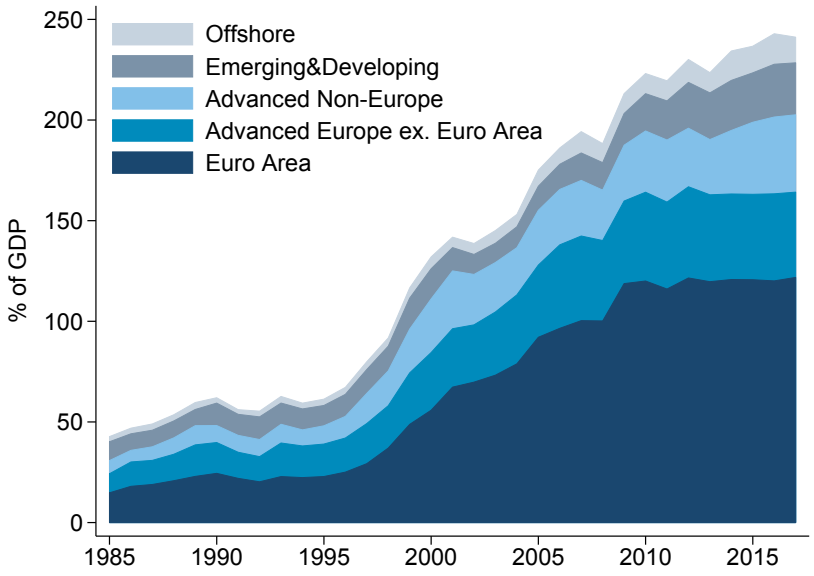

Notes: This graph shows that the majority of German foreign assets are invested in other advanced economies, especially in Europe. The geographical distribution is estimated from additional data sources, see Subsection 4.2.3. The figure excludes financial derivatives since no data on their geographical distribution available. Choice of offshore countries is based on Bundesbank list of offshore banking centers. GDP from the Macro History Database (Jordà et al. 2017) and the German Statistical Office.

Therefore, we rely on additional data sources to estimate the geographical distribution of foreign investments for Germany and other countries, as discussed at the end of Subsection 4.2.3.

Figure 4.6 shows the resulting decomposition into four regions since 1985.[20] The figure reveals that Germany mainly invests in other advanced economies, especially in fellow European countries. The introduction of the Euro in 1999 further increased the European exposure as even more investment went to other euro-area countries. Today, almost 70% of all investments are in other advanced European economies, another 15% are in non-European advanced economies (mainly the US), and only the remaining 15% are invested in other countries worldwide, including offshore destinations.

[20]The additional data sources needed do not allow for a meaningful estimation for the period before 1985.

Table 4.2: Returns on German foreign assets, 1950 to 2017

	Panel (a): Real returns			Panel (b): Nominal returns		
	1950–17	1999–17	2009–17	1950–17	1999–17	2009–17
Return, all assets	1.59	2.28	2.62	3.95	3.73	3.78
Yield, all assets	1.75	2.00	1.77	4.18	3.45	2.93
Valuation changes, all assets	-2.48	-1.12	-0.28	-0.23	0.27	0.85
Return, FDI	-0.18	3.78	4.85	2.38	5.27	6.05
Return, equity portfolio hold.	8.25	3.13	8.41	11.00	4.52	9.61
Return, debt portfolio hold.	5.66	3.18	3.40	8.39	4.65	4.57
Return, other inv.+ reserves	1.15	1.21	0.27	3.50	2.65	1.40

Notes: This table shows average real and nominal returns on German foreign assets. Returns are split by components and asset category. Returns estimated using Bundesbank data as discussed in Subsection 4.2.1. Real returns deflated using German consumer price index from Macro history Database and Eurostat.

4.4 Returns on German foreign and domestic assets

We now turn to the analysis of the profitability of German foreign investments. This section presents descriptive statistics of Germany's foreign asset returns and shows that these are lower than returns on domestic investments.

4.4.1 Germany's foreign investment returns 1950 to 2017

Table 4.2 summarizes the German return, yield, and valuation changes since 1950. For the comparison over time it is more informative to focus on real returns (deflated by national CPI), but we also show nominal returns for the main tables and figures (to save space, some were shifted to Appendix D2).

The average real annual return between 1950 and 2017 was 1.59%. The yield was 1.75% while valuation changes were negative on average at -2.48%. Recall that the real return equals the sum of the real yield and real valuation changes plus some adjustment for inflation. Average real returns increased in the more recent periods to around 2% depending on the time horizon. This is mainly due to lower valuation losses, albeit they remain negative.

The lower part of Table 4.2 reveals large differences between the asset categories.[21] Over the full sample, portfolio equity investments saw the highest returns on average, followed by portfolio debt. FDI and 'other investment' had much smaller real returns. However, since 1999 the relationships changed: equity returns fell and FDI increased so that now they are roughly similar.

[21] As discussed in the data section, we need to combine 'other investment' and reserves when computing returns since investment income data is not available separately for those two categories.

Figure 4.7: Real returns on German foreign assets, 5-year rolling means, 1950 to 2017

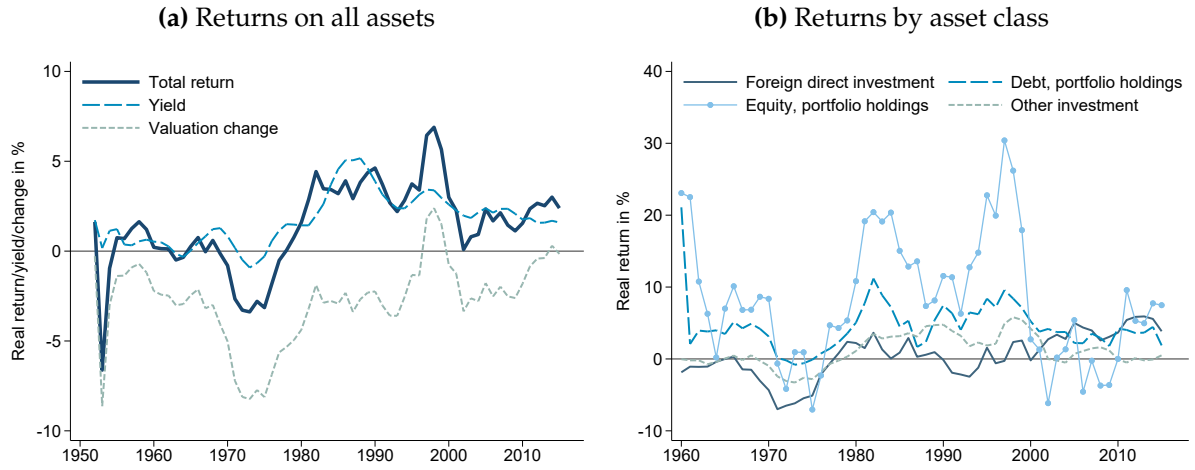

Notes: This graph shows real returns on German foreign assets as a rolling arithmetic mean computed over 5-year windows and plotted at the third year of the window. Panel (a) shows returns on all assets and the decomposition into yield and valuation changes. Panel (b) shows total return series by asset category. Returns are estimated using Bundesbank data as discussed in Subsection 4.2.1. The series are deflated using the German consumer price index from the Macro History Database (Jordà et al. 2017) and Eurostat.

To visualize developments over time, Figure 4.7 plots 5-year rolling averages of our measures of returns. Panel (a) shows the real return, yield, and valuation changes on total assets. Panel (b) plots the real returns by asset category. Several observations stand out.

First, valuation changes are more volatile than yields and drive the volatility in returns (as should be expected). Average real valuation changes were almost always negative. The improvement in average returns on German foreign assets since the 1980s was driven mainly by a significant increase in the yield, i.e., the direct income earned on investments. Average real valuation changes, however, have remained in negative territory even over much of the last decades, which is surprising as global asset markets have performed exceptionally well since the 1970s.

Second, the average return to IIP assets hides diversity across asset classes. The returns vary strongly across these asset categories, as shown in Panel (b) of Figure 4.7. Returns to foreign direct investment were low for many decades but increased in the 2000s. Portfolio investment generated larger returns than the aggregate asset position in most periods. This is mainly due to high returns on foreign equity. Finally, the 'other investment' category that, among others, includes bank loans saw returns comparable to the aggregate return.

4.4.2 Comparison to domestic returns

How do Germany's external returns compare to returns on domestic capital? There are two main options to address this question. Many studies compare the return earned abroad to the return earned on domestic capital markets. Others compare it to the return earned on the other side of international balance sheet, i.e., on inward investments by foreigners. As explained above, the comparison with liability returns is not necessarily insightful regarding the quality of foreign investment. Foreign liabilities do not cover all investment opportunities available to German investors in their own country and the returns reported by foreigners can be downward biased due to tax shifting, especially in a high-tax country like Germany. For these reasons, we focus on the first option – the comparison of foreign returns with returns on the aggregate capital stock in Germany.

Despite our emphasis on domestic portfolio returns, we also compute the return on IIP liabilities in Appendix D6. In line with earlier studies, we find that the difference between asset and liability returns in Germany was negative for a long time but decreased in the past 20 years and recently turned positive. This trend is mainly driven by decreasing yields on FDI liabilities and debt liabilities. The latter is not surprising given the flight to safety compressing German bond returns after 2008. The former may be related to tax incentives leading to the increased leverage on inward FDI, which in turn leads to relatively low reported yields (see Appendix D6 for a discussion).

The return on Germany's capital stock (held by both foreigners and Germans) is taken from the data set of Jordà et al. (2019). We make use of the return to capital, which is computed as the return to a portfolio consisting of equity, housing, bonds, and bills. To compute the return to capital, the authors compute returns for all asset categories using various data sources. The returns also include both valuation changes and direct income flows. In the case of Germany since World War II, the authors use money market rates for the bills returns, the performance index for the Bund bond market return, the German stock market index for equity, and housing returns based on the rent-price approach. Then they aggregate individual returns to arrive at an aggregate return on capital using appropriate portfolio weights. These weights are stock market capitalization for equity returns, housing wealth for housing returns, and public debt split equally between bonds and bills. Unfortunately, the return to capital series starts only in 1963.

Figure 4.8 plots five-year rolling averages of the foreign and domestic returns. It reveals that domestic returns were significantly higher than the return earned abroad for the majority of the time observed. On average, the difference was more than 3 percentage points. Only in the early 1980s and in the early 2000s until the financial crisis average returns were roughly equal. Moreover, while the average domestic return

Figure 4.8: Real foreign and domestic returns, 5-year rolling means, 1963 to 2017

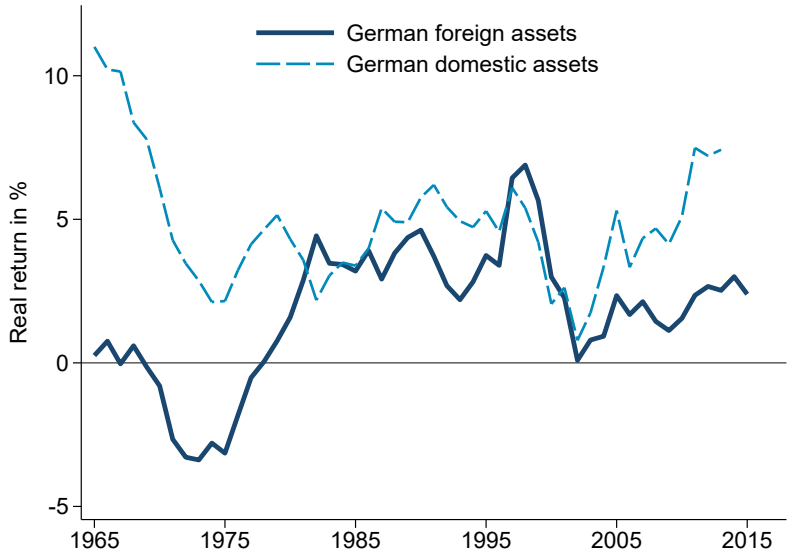

Notes: This graph shows that the return on German foreign assets (dark blue line) is lower than the domestic return on German assets at home (dotted line) for most years. The series are rolling means computed over 5-year windows and shown at the third year. The return on German domestic assets is from Jordà et al. (2019) and available 1963-2015. Both series are deflated using the German CPI.

computed by Jordà et al. (2019) rose significantly in the past decade, foreign returns did not increase.

4.5 International comparison of returns

Our main focus is comparing German foreign returns to other countries' foreign returns. This is the most plausible comparison to gauge Germany's relative performance when investing abroad since all (advanced) economies in principle have access to the same investment opportunities. For this comparison, we computed returns for a group of 12 additional advanced economies from 1975 to 2017. We started by collecting data for each of the G7 countries, although we could not find detailed, long-run data on foreign assets and their returns for Japan. We then tried to add as many OECD countries as possible. The final selection is based on data availability in particular with regard to the level of disaggregation and the years covered, since we wanted a long time horizon. We express all returns in the country's respective domestic currency.

Figure 4.9: Nominal returns in comparison, 5-year rolling means, 1975 to 2017

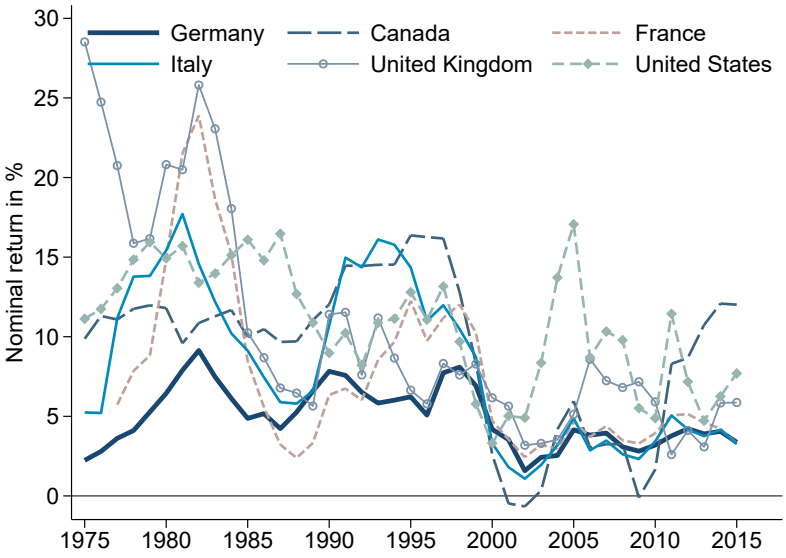

Notes: This graph shows that Germany's returns on foreign assets (dark blue line) were almost always lower than the foreign returns of other G7 members (excluding Japan due to data availability). Rolling means computed over 5-year windows and plotted at the third year of the window. We compare nominal domestic currency returns to abstract from the effects of different inflation dynamics across the countries. The overall picture is similar when plotting real returns, as shown in Appendix A1. For more details, see Section 4.2.

4.5.1 Descriptives

For now, we focus on the other G7 members for comparison to keep the graphs and tables simple (Japan is not included due to data availability). Later, we will include the additional advanced economies in the comparison group. Figure 4.9 shows that German returns abroad were consistently lower than those of other countries. Importantly, this does not apply only to the US with its "exorbitant privilege" in international finance, but also to Italy, France, Canada, and the UK.

Table 4.3 summarizes the key return statistics in comparison and over different time horizons. The table demonstrates that German returns were lower than the European average, and consistently lower than domestic returns.

4.5.2 Regression evidence

Is Germany's financial underperformance statistically significant in a broader country panel? In the next step, we test whether German returns are lower relative to a larger group of countries. As explained, the countries we consider are 12 comparable OECD economies, namely Canada, Denmark, Finland, France, Italy, the Netherlands,

Table 4.3: Comparing returns, nominal, various time horizons

	1975–17	1999–17	2009–17	1980–89	1990–99	2000–09	2010–17
Germany, foreign assets	4.95	3.73	3.78	6.68	7.12	2.77	3.65
Germany, domestic assets	7.03	6.26	8.48	6.79	8.37	4.24	8.68
Canada	9.19	4.93	8.98	10.27	15.31	1.29	8.88
France	7.38	4.01	4.43	13.56	8.60	3.41	4.13
Italy	7.96	3.31	4.39	10.23	13.17	2.28	3.72
United Kingdom	10.22	5.68	4.09	16.30	7.96	5.21	5.14
United States	10.64	8.00	9.27	14.93	10.69	7.62	7.50
Germany minus domestic	-1.94	-2.33	-4.21	-0.11	-1.25	-1.47	-4.51
Germany minus US	-5.69	-4.27	-5.49	-8.25	-3.57	-4.85	-3.85
Germany minus Europe[1]	-2.75	-0.68	-0.94	-5.02	-2.52	-0.53	-1.05

Notes: This table shows that Germany's returns on foreign assets were lower than the return on domestic assets as well as the foreign returns of other G7 members (excluding Japan due to data availability). Foreign returns are computed as discussed in Section 4.2, while the domestic German return is from Jordà et al. (2019). We compare nominal domestic currency returns to abstract from the effects of different inflation dynamics across the countries.
[1] Europe is an average of DK, ES, FI, FR, GB, IT, NL, NO, PT, SE.

Norway, Portugal, Spain, Sweden, the UK, and the US. We regress the observed annual returns for Germany and these other advanced economies on standard control variables as well as a "German dummy". We are mainly interested in the size and significance of the coefficient on the German dummy as it tells us whether German returns differ significantly compared to other countries. In the cross-country setting, we focus on nominal returns because we are interested in the direct returns that individual countries earn abroad, abstracting from inflation dynamics across the countries. Using this data, we estimate the following model using pooled OLS:

$$\tilde{r}_{it} = \alpha + \beta D_{it} + \delta Z_{it} + \gamma_t + u_{it}, \tag{4.1}$$

where D_{it} is a dummy variable which is 1 for Germany and 0 for the other countries, Z_{it} are control variables and γ_t are yearly time fixed effects. In line with the existing literature, we control for the size of the net foreign asset position as well as past financial account balances relative to GDP to capture rebalancing effects via returns in countries with large past and current external imbalances (as discussed by e.g., Gourinchas and Rey 2014).

Columns (1) to (4) of Table 4.4 present our core finding: German foreign investment returns are consistently about 2 percentage points lower than the returns of other countries. Moreover, while the exact underperformance fluctuates between 1 and 3 percentage points, the finding is robust across different periods. Another important finding is that the investment underperformance we measure is statistically significant.

Excluding the Target2 balances from foreign assets of all euro-area countries does

Table 4.4: Determinants of returns on foreign assets, 1985 to 2017

	Baseline (all assets)				Excluding Target2	
	(1) 1985–2017	(2) 1985–1998	(3) 1999–2017	(4) 2009–2017	(5) 1999–2017	(6) 2009–2017
Germany dummy	-2.27***	-3.01**	-1.48**	-2.08**	-1.36**	-1.79**
	(0.63)	(1.51)	(0.59)	(0.80)	(0.59)	(0.81)
Constant	4.20*	4.46*	14.01***	7.42***	14.04***	7.42***
	(2.47)	(2.46)	(1.31)	(1.80)	(1.32)	(1.81)
Observations	406	160	246	117	246	117
Adjusted R^2	0.38	0.27	0.41	0.18	0.41	0.19
No. countries	13	12	13	13	13	13
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows that German returns on foreign assets are significantly lower than the foreign returns of other advanced economies. This is true across different samples (columns (1) to (4)) as well as when Target2 balances are excluded (columns (5) and (6)). The dependent variable is the nominal rate of return on total foreign assets by country and year. The regressions include control variables for net foreign assets and the financial account balance (coefficients not shown). Data for Denmark and Portugal starts in 1999 and 1993, respectively. No data for Japan available. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

not alter our main finding. This is shown in columns (5) and (6) of Table 4.4, which focus on the period after 1998 when the Euro (and, thus, the Target2 system) was introduced. The coefficients confirm that Germany's returns on foreign assets are about 1.5 percentage points lower than the returns of other countries.

We also consider how the Germany dummy evolved over time to see if the German underperformance is driven by particular episodes. To test this, we interact the German country dummy with year fixed effects to estimate a time-varying effect. The regression and control variables are the same as above. Figure D.2 in Appendix D1 plots the resulting coefficients. The results show that no particular period leads to the negative Germany dummy. In addition, the Germany dummy was always smaller or equal zero, never significantly larger than zero.

4.5.3 Aggregate financial consequences – a counterfactual exercise

In this section, we aim to quantify the cumulative financial loss (or foregone gains) caused by Germany's low returns on foreign assets. For this, we need to construct a counterfactual in which German returns would have been comparable to those of other countries.

We compute Germany's counterfactual investment income assuming that Germany achieved the rates of return by the other G7 members (we also include Norway, which achieved one of the highest returns over the past decade). We then compare these counterfactual values to the realized income earned and compute the aggregate cumulative

Table 4.5: Cumulated income losses due to low returns on German foreign assets

	1975–2017		1999–2017		2009–2017	
	bn 2015 €	% of GDP	bn 2015 €	% of GDP	bn 2015 €	% of GDP
Canada	-3852.96	-105.20	-2158.89	-69.50	-3075.99	-93.77
France	-1093.66	-25.03	-332.04	-9.44	-323.80	-9.39
Italy	-1114.90	-22.57	217.41	5.02	-306.43	-8.89
Norway	-3712.11	-103.29	-2622.25	-80.10	-2245.51	-69.52
United Kingdom	-3093.65	-76.55	-2050.79	-59.02	-388.85	-13.02
United States	-5676.56	-150.60	-4331.29	-124.24	-3080.99	-92.21

Notes: This table quantifies the foregone income on Germany's foreign assets due to Germany's comparatively low investment returns abroad. Losses are computed as the difference between total income earned (yield plus valuation changes) on German assets and hypothetical income earned had Germany achieved the same return as the comparison country. Nominal losses are deflated using the CPI index with 2015=100 and then added up over time. In columns (2), (4), and (6) the nominal counterfactual losses are shown as a share of nominal German GDP of 2017.

loss or gain. To evaluate the economic size of the effects, we deflate the losses using German CPI and express them as percent of German GDP in 2017. More specifically, we compute the aggregate loss, $Loss_t$, in each year as

$$Loss_t = \tilde{r}_{t,DE} IIP^A_{t-1,DE} - \tilde{r}_{t,c} IIP^A_{t-1,DE} = (\tilde{r}_{t,DE} - \tilde{r}_{t,c}) IIP^A_{t-1,DE},$$

where $\tilde{r}_{t,\text{DE}}$ is the nominal return on German foreign assets, $IIP^A_{t-1,DE}$ is the German gross foreign asset position, and $\tilde{r}_{t,\text{c}}$ is the nominal return on the foreign assets of country c. Table 4.5 displays the resulting losses.

As can be seen, the amounts are substantial. Had Germany been as savvy an external investor as other countries, the country would be considerably richer today. The losses are largest when using US returns as counterfactual, but the numbers are also substantial when comparing to other countries. Had Germany achieved the same return on investment as France since the introduction of the Euro (in 1999), the country would be 330bn Euros richer today, according to these simple back-of-the-envelope calculations. By not achieving the returns that Italy achieved, Germany forgave wealth of about 300bn Euros in the decade since the financial crisis alone. More remarkably, since 2009, Germany would have gained an additional 2-3 trillion Euros of wealth (or 70% or 94% of its 2017 GDP) if its foreign investments had performed like those of Norway or Canada, respectively. In per capita terms, this amounts to about 27,000 and 37,000 Euros of foregone income for each German citizen in less than 10-years, a substantial

wealth loss compared to Norway and Canada, respectively.[22]

4.6 Why are German returns low?

We have established that German returns on foreign investments are considerably lower than the returns of other countries. In this section, we aim to understand the causes. More precisely, we decompose the return differential using cross-country data as far back as possible (mostly starting in 1985, when detailed data on the currency composition of foreign assets became available for most economies, see Section 4.2).

The negative German return differential may result from a range of factors linked to asset allocation (asset class selection, geography, risk profile). In addition, the differential could be due to exchange rate effects. In particular, exchange rate appreciations could have systematically lowered the "raw" returns achieved by German investors. We will start by looking at the latter.

4.6.1 Exchange rate effects

To understand the role of exchange rates, we need information on the valuation changes due to exchanges rates. The newest edition of the BPM requires countries to publish a decomposition of valuation changes into the three components exchange rates, prices and other adjustments. However, most countries either publish this breakdown only for recent years or have not started publishing it yet. Therefore, we estimate the valuation changes due to exchange rates ourselves.

In order to do this, we modify the approach of Lane and Shambaugh (2010) and Bénétrix et al. (2015). These authors show that data on the currency composition of assets is sufficient to estimate valuation changes due to exchange rate movements. To see this, first note that valuation changes due to exchange rates are the changes in the valuation of all foreign currency assets valued in the domestic currency neither due to transaction nor due to price or other changes:

$$VX_t^A = \sum_c \left(IIP_t^{A,c} E_t^c - IIP_{t-1}^{A,c} E_{t-1}^c - FA_t^{A,c} \overline{E}_t^c - (VP_t^{A,c} + VOT_t^{A,c}) \overline{E}_t^c \right), \tag{4.2}$$

where the superscript c indicates a variable, which is expressed in a different currency than the German one. E_t^c and $\overline{E}_t^c$ are the end of period and average exchange rates in year t between the Euro or DM and currency c, respectively. Here, we follow the recommendation of the IMF's balance of payments manual and use average exchange

[22] Euro values are the real values from Table 4.5. German population size was 82.8 million in 2017 (German Statistical Office). GDP shares are obtained by dividing the nominal loss by German GDP in 2017.

Figure 4.10: Valuation changes due to exchange rates, 5-year rolling means, 1985 to 2017

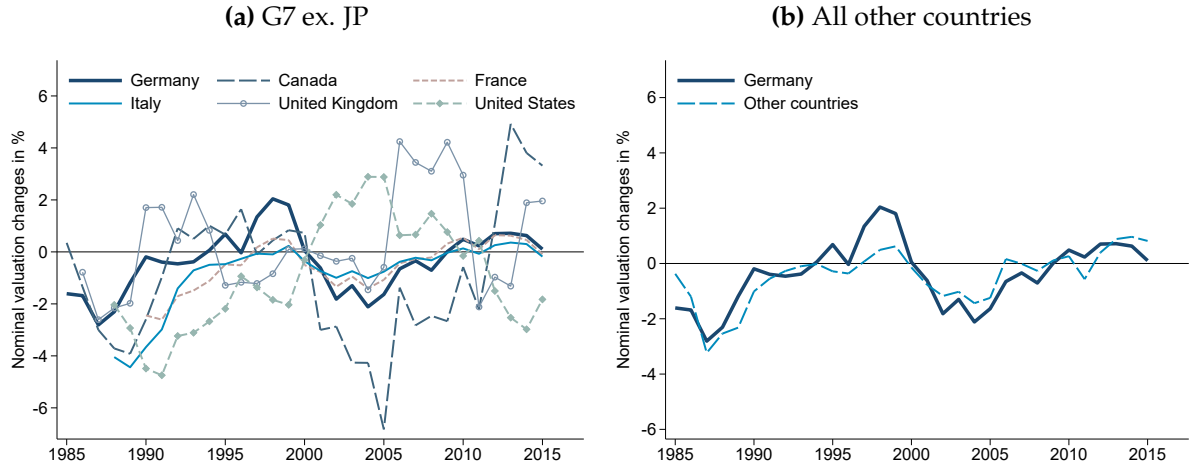

Notes: Graphs show that Germany's valuation changes due to exchange rates are not significantly different from those of other countries. The lines represent rolling means computed over 5-year windows across countries and plotted at the third year of the window. Valuation changes computed using estimated currency shares from additional data sources, see Subsection 4.6.1. "Other countries" in Panel (b) refers to CA, DK, ES, FI, FR, GB, IT, NL, NO, PT, SE and US.

rates to value transactions and valuation changes due to a lack of data on the timing of the two. Finally, note that the last part of the expression in equation (4.2) equals all changes in the value of the foreign currency assets not due to transactions when valued in the respective currency since there can be no valuation effects within the same currency, i.e.,

$$VP_t^{A,c} + VOT_t^{A,c} = IIP_t^{A,c} - IIP_{t-1}^{A,c} - FA_t^{A,c}.$$

Plugging this expression into equation (4.2) yields the following simple expression for VX_t^A after some manipulation:

$$VX_t^A = \sum_c \left(IIP_t^{A,c} E_t^c - IIP_{t-1}^{A,c} E_{t-1}^c - (IIP_t^{A,c} - IIP_{t-1}^{A,c}) \bar{E}_t^c \right).$$

Now we only need data on the currency composition of assets. As discussed in Subsection 4.2.3, this is not available directly from the countries. However, we are able to use other data sources to estimate currency shares. Despite the several approximations involved in the estimation procedure, our estimates are very close to the ones published by the Bundesbank for Germany as well as those published by other countries' statistical institutions (see Appendix D5).

Panel (a) of Figure 4.10 shows that German valuation changes due to exchange rates did not differ significantly from that of other countries. All countries experience both

gains and losses due to exchange rates.[23]

Before the introduction of the Euro, exchange rate effects tended to be relevant in Germany, with valuation effects being among the larger ones within the G7 group. In the period since then, German valuation changes due to exchange rates were rather average. We can also compare Germany to the average of all other 12 countries (as included in the regression in the previous section). Panel (b) of Figure 4.10 confirms that German valuation changes do not stand out relative to these other countries. In sum, exchange rate effects do not help explain the observed differences in returns between Germany and other countries.

4.6.2 Portfolio allocation

Lower German returns could be a result of a more conservative investment strategy that favors less risky asset classes such as bonds over equities. Figure 4.11 shows that there are indeed notable differences in the composition of foreign assets between Germany and the other countries. We show the periods 1985-1999 and 2001-2017 separately to account for the large compositional changes in the 1990s as documented in Section 4.3.

The figure reveals that Germany's large share of 'other investment' in the 1980s and 1990s was atypical compared to other countries. In addition, the shift towards more FDI and portfolio in recent decades was sizeable but other countries increased their investment in these positions even more. As a result, in the past two decades Germany invested significantly less in FDI and equities than the other countries. This provides an indication of the potential relevance of compositional differences. However, the detailed decomposition discussed in the next section reveals that differences within each asset category are even more relevant.

Other than the asset composition, the geographical allocation of foreign investments could affect performance. Unfortunately, no data is available on the returns by geographical location, so that we cannot include geography in our decomposition exercise. However, as explained, we can estimate which share of Germany's assets is located in which country (see Subsection 4.2.3). This allows us to include geography controls in our regression analysis in Subsection 4.6.4. Moreover, we know that the large bulk of Germany's investments goes to other European high-income countries, while the share of investments in the rest of the world is small and shrinking. In international comparison, Germany stands out as a country with a particularly strong

[23]The observation of negative valuation changes due to exchange rates raises a more fundamental point about international adjustment. The intertemporal approach to the current account implies that valuation changes matter for the external solvency constraint and via this constraint may be a potential channel for external adjustment (Gourinchas and Rey 2007b, 2014).

Figure 4.11: Composition of IIP assets, 1985 to 2017

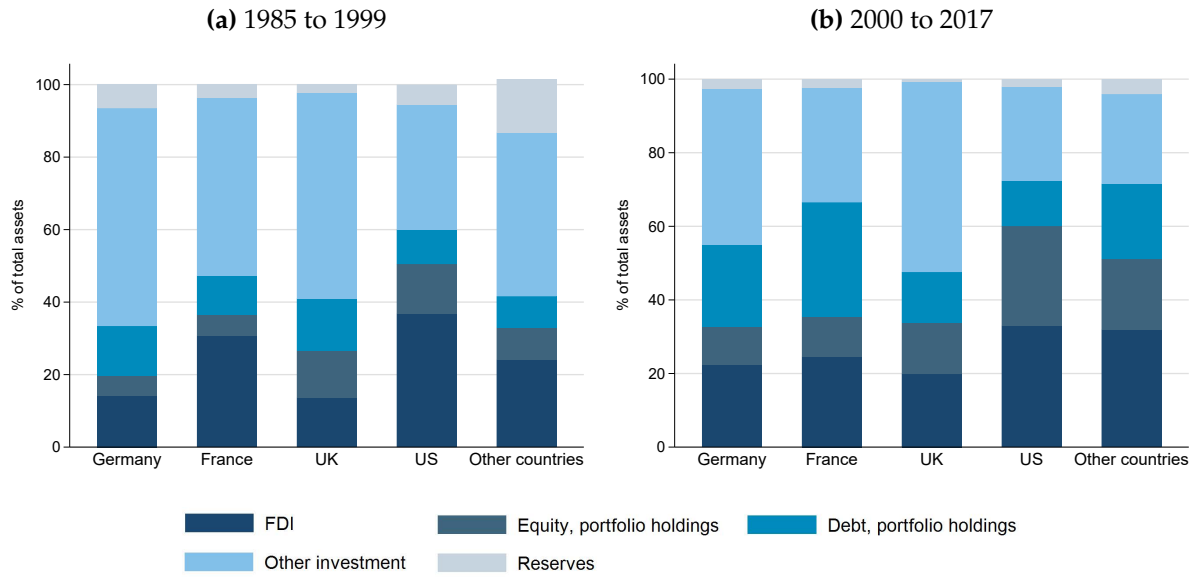

Notes: These graphs show that Germany holds a lower share of FDI and equity assets than other countries. This is true both in the earlier part of the sample (Panel (a)), and in recent years (Panel (b)). Data from Bundesbank, IMF and (Lane and Milesi-Ferretti 2007b). 'Other countries' includes CA, DK, ES, FI, IT, NL, NO, PT and SE.

"home bias" in favor of European investments.

4.6.3 Comparing exchange rate, composition and return effects

In this section, we decompose the return differential into its components to analyze the effects of exchange rates and asset composition more systematically. The exercise starts in 1990, since we need information on the returns for all asset categories separately. This is not available for most countries before the 1990s.

Our cross-country decomposition approach builds on and expands the exercise that Gourinchas and Rey (2007a) apply for the United States. Specifically, one can write the aggregate return to any portfolio p as the weighted average of the returns of the different portfolio components j using lagged weights:

$$r_t^p = \sum_{j=1}^{J} w_{j,t-1}^p r_{j,t}^p, \tag{4.3}$$

where $w_{j,t-1}^p$ is the weight of assets class j in the portfolio p, $r_{j,t}^p$ is the return to the respective asset class and J is the number of asset classes. Using equation (4.3) one can

then rewrite the difference in aggregate returns of two portfolios p and q as

$$r_t^p - r_t^q = \sum_{j=1}^{J} \frac{r_{t,j}^p + r_{t,j}^q}{2} (w_{t-1,j}^p - w_{t-1,j}^q) + \sum_{j=1}^{J} \frac{w_{t-1,j}^p + w_{t-1,j}^q}{2} (r_{t,j}^p - r_{t,j}^q). \tag{4.4}$$

The first term of equation (4.4) captures the difference in returns resulting from the different weights of each asset class in the two portfolios and is labeled the *composition effect*. The differences between weights are weighted by the average return of the respective asset class in both portfolios. The second term captures the effect of the return differential on the overall difference and is called the *return effect*.

Furthermore, we expand the exercise by subtracting valuation changes due to exchange rates from the returns before decomposing the difference. This allows us to parse out the *exchange rate effect*.

Table 4.6 shows the decomposition results, which focuses on the comparison between Germany and other G7 members. The first column shows the difference between Germany's foreign returns vis-à-vis each comparison country, averaged for the full period 1990 to 2017. In line with our findings above, German returns are lower than those of the other G7 members, so that the sign is negative in each row.

The second column shows the contribution of valuation changes due to exchange rates. The positive values indicate that most countries' returns suffered more from exchange rate-driven valuation changes (appreciation effects) than Germany's return. The only exception is the United Kingdom, where exchange rate effects can help to explain about half of the return differential with Germany (-0.8 of -1.6 percentage points overall). Taken together, however, the numbers in column (2) are small, so that exchange rate movements do not help much to explain the observed gap between German and other countries' returns.

The third column indicates that the asset composition is also not a major driver of the observed return differentials between Germany and the other countries. Only in the comparison with Canada and the US, the asset composition plays a non-trivial role, accounting for up to 25% of the return differences (-0.8 and -1.3 percentage points respectively, see column (3)).

The dominant part of the explanation are differences in returns within each asset class. This can be seen in column (4) which shows large negative numbers. The return effect explains more than three quarters of the differences in returns between Germany and other countries.

In sum, the main reason why German foreign investments produce lower returns is

Table 4.6: Decomposition of return differences: Germany vs other countries, 1990 to 2017

Comparison country	Difference in foreign returns (pp.)	Difference due to		
		exchange rates	composition (asset class)	returns within asset class
	(1)	(2)	(3)	(4)
Canada	-3.892	0.316	-0.820	-3.389
France	-0.894	0.374	0.192	-1.463
Italy	-2.007	0.340	-0.357	-1.993
UK	-1.598	-0.794	-0.200	-0.604
US	-4.106	0.441	-1.254	-3.294

Notes: Decomposition splits the difference between return on German foreign assets and other country's foreign assets into three parts: (1) difference in valuation changes due to exchange rates, (2) different composition of asset position in the four broad asset categories, and (3) difference in returns within each asset class (details in Subsection 4.6.3). Returns estimated as discussed in Subsection 4.2.1. For countries, we compare nominal returns to abstract from the effects of different inflation dynamics across the countries.

not the type of assets Germany holds (debt vs. equity vs. FDI) nor frequent exchange rate appreciations. Instead, Germany's foreign assets are less profitable within the same asset class, after controlling for exchange rate and composition effects. We now turn to understanding why this is the case.

4.6.4 Returns within asset classes

In this section, we want to understand what is driving the modest investment performance on the level of individual asset classes. For this purpose, we compare German returns to other countries' return for individual asset classes. We return to our regression model from equation (4.1) and include the geographical distribution of each country's assets as well as additional asset characteristics.

For each asset class j – portfolio equity and debt, foreign direct investment, and 'other investment' – we estimate the following regression:

$$\tilde{r}^j_{it} = \alpha_j + \beta_{1j} D_{it} + \sum_r \beta_{rj} S^{A,j,r}_{it} + \beta_{2j} \sigma^j_{it} + \delta Z_{it} + \gamma_t + u_{ijt},$$

where $S^{A,j,r}_{it}$ is the share of assets from region r in total assets of category j owned by country i in year t. Z_{it} are additional control variables and γ_t are yearly time fixed effects. In addition to the net foreign asset position and the financial account balance, Z_{it} now also includes the exchange rate effects estimated before, as these are an important driver of returns. We also include a measure of risk, σ^j_{it}. Specifically, following stan-

Table 4.7: Determinants of returns by asset class, equity and debt, 2002 to 2017

	Equity returns				Debt returns			
	(1) Baseline	(2) Val. FX	(3) Risk	(4) Geo.	(5) Baseline	(6) Val. FX	(7) Risk	(8) Geo.
Germany dummy	-4.17**	-4.16**	-4.75***	-3.04	-0.70	-0.83	-0.44	-0.87
	(1.76)	(1.72)	(1.82)	(2.21)	(0.91)	(0.90)	(0.88)	(1.07)
Valuation ch. due to ex. rates, equity		1.04***	0.97***	0.92***				
		(0.13)	(0.14)	(0.15)				
3-year rolling std. dev., equity			0.38**	0.38**				
			(0.17)	(0.17)				
Valuation ch. due to ex. rates, debt						1.01***	0.99***	1.03***
						(0.15)	(0.15)	(0.15)
3-year rolling std. dev., debt							0.23	0.11
							(0.20)	(0.21)
Advanced Europe				0.12				-0.14**
				(0.15)				(0.07)
Advanced Non-Europe				0.17				-0.08
				(0.15)				(0.08)
Emerging &Developing				0.20				-0.16**
				(0.18)				(0.07)
Constant	-20.68***	-18.10***	-25.04***	-39.40**	8.95***	10.06***	8.52***	21.44***
	(1.89)	(1.76)	(4.15)	(17.23)	(2.99)	(2.23)	(2.04)	(7.00)
Observations	175	175	164	164	175	175	164	164
Adjusted R^2	0.76	0.82	0.84	0.84	0.09	0.39	0.41	0.45
No. countries	11	11	11	11	11	11	11	11
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Germany's returns on equity are significantly lower than other countries' returns even after controlling for exchange rate effects and risk. Debt returns are comparable, albeit the coefficients of the Germany dummy are negative as well. Sample restricted to 2002-2017 due to a lack of data on geographic allocation and exchange rate valuation effects by asset class (see Section 4.2). Spain and Norway dropped entirely due to a lack of data on the geographical distribution by asset class. Net foreign assets and financial account balance included in the regressions but not shown. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

dard practice, we use the standard deviation of the respective return series, computed over 3-year rolling windows and centered around t.

Note that in this exercise we set the bar intentionally high. The foreign exchange exposure, geographic portfolio allocation, and risk are part of the investment decision of German savers or intermediaries. Investors can freely choose where and what to invest in. The regressions therefore test an even stricter version of the German returns puzzle. We ask, conditional on foreign exchange movements and other controls, did German investors receive worse returns *within* individual asset classes than other countries?

Table 4.7 and Table 4.8 present the results in four columns for each asset category. The regression in the first column only includes the controls from the baseline specifica-

Table 4.8: Determinants of returns by asset class, FDI and 'other investment', 1985 to 2017

	FDI returns				'Other investment' returns			
	(1) Baseline	(2) Val. FX	(3) Risk	(4) Geo.	(5) Baseline	(6) Val. FX	(7) Risk	(8) Geo.
Germany dummy	-3.34***	-3.61***	-3.23**	-2.90**	-1.25	-1.42*	-1.75**	-1.07
	(1.25)	(1.27)	(1.30)	(1.36)	(1.02)	(0.75)	(0.68)	(0.82)
Valuation ch. due to ex. rates, FDI		0.38***	0.36***	0.37***				
		(0.10)	(0.11)	(0.11)				
3-year rolling std. dev., FDI			0.15	0.13				
			(0.20)	(0.21)				
Valuation ch. due to ex. rates, 'other'						1.00***	1.07***	1.09***
						(0.08)	(0.08)	(0.09)
3-year rolling std. dev., 'other'							-0.16	-0.14
							(0.11)	(0.14)
Advanced Europe				-0.12				-0.09*
				(0.15)				(0.05)
Advanced Non-Europe				-0.10				-0.07
				(0.18)				(0.06)
Emerging &Developing				-0.10				-0.10
				(0.15)				(0.10)
Constant	9.83	11.69**	10.56**	21.33	3.63	12.68***	12.74***	20.72***
	(5.98)	(5.18)	(5.09)	(17.06)	(3.37)	(1.40)	(0.89)	(4.77)
Observations	339	339	328	328	227	227	212	189
Adjusted R^2	0.14	0.18	0.19	0.18	0.33	0.63	0.67	0.67
No. countries	12	12	12	12	10	10	10	10
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Germany's returns on FDI are significantly lower than those of other countries even after controlling for exchange rate effects, risk and geographical allocation. Returns on 'other investment' are lower but the effect is not always significant. Geographical distribution for the respective asset class. Net foreign assets and financial account balance included in the regressions but not shown. The results exclude Spain due to a lack of data on the geographical distribution by asset class. The results for 'other investment' further exclude Norway and Portugal for the same reason. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

tion in Section 4.4. The next three columns control for exchange rate effects, geography, and risk, respectively.

We lack data on the geographic allocation of assets and on exchange rate valuation effects for the regressions on debt and equity, so that the sample is restricted to 2002-2017 in Table 4.7. For similar reasons, also the country sample varies across different panels.[24] In particular, for the 'other investment' category, data on the geographical distribution is relatively scarce (also see Table D.5). Therefore, the number of observa-

[24]All results exclude Spain due to a lack of data on the geographical distribution by asset class. The results for debt and equity further exclude Norway due to a lack of data on geographic allocation for these assets. The results for 'other investment' exclude both Norway and Portugal for the same reason.

tions and countries changes with the inclusion of geographical composition.

The tables deliver a clear finding. Germany's returns are consistently lower across asset classes. This means that returns are lower even when we zoom in to individual asset classes and control for risk characteristics. The effects are particularly pronounced for portfolio equity investment and FDI. Both stand out economically and statistically as markets in which German returns were substantially lower, even after accounting for the effects of other allocation choices (which can also be seen as part of investment performance). For debt and 'other investment' the mean effect remains negative at about 1 percentage point but is not significant.

4.7 Benefits for consumption insurance and demographic risks

So far, we showed that German returns are systematically lower than other countries' returns even after controlling for various compositional aspects. One reason why Germans may accept these low returns is an insurance motive. If returns are countercyclical with respect to domestic consumption, they provide a hedge against volatility in consumption. Additionally, Germans may strategically invest in regions with better demographic prospects. These might not yield high returns yet but potentially will in the future. In this section, we briefly touch upon these two potential channels.

4.7.1 Consumption insurance – income smoothing from abroad?

There is large literature on the potential of international investment to reduce consumption risks. The basic idea is that foreign investments can help to buffer shocks to household consumption. Suppose Germany witnesses an economic downturn while foreign countries do not. In this situation, German households that have invested into foreign assets will benefit from their (high) capital income from abroad to counterbalance their (lower) domestic income. The foreign capital income will thus help households to smooth their consumption over time, making them better off. So far, however, the literature found only very limited effects of investment income flows on consumption smoothing (see e.g., Lane 2001; Sørensen and Yosha 1998).

Here, we test to what extent German foreign investments provide consumption insurance for German households. We again focus on total returns on the foreign assets, i.e., we combine yields and valuation changes. We base our empirical approach on the consumption capital-asset pricing model (CCAPM). The CCAPM assumes that the return of a risky asset is proportional to the consumption beta. Intuitively, this implies

Table 4.9: Correlation of real consumption and real returns

	1985–2017	1999–2017	2009–2017
Germany, foreign assets	0.298	0.504	0.444
Germany, domestic assets	0.268	0.319	0.134
Canada	0.099	0.353	0.579
Finland	-0.152	0.400	0.837
France	0.127	0.397	0.420
Italy	0.072	0.528	0.563
Netherlands	0.319	0.255	0.667
Norway	0.133	0.089	-0.274
Spain	0.053	0.381	0.675
Sweden	-0.027	0.424	0.649
United Kingdom	-0.126	-0.213	-0.516
United States	0.379	0.366	0.553

Notes: This table shows that Germany's real returns on foreign assets are positively correlated with real domestic consumption growth per capita. The correlation coefficient between returns and consumption is higher for foreign assets than for domestic German assets. The coefficient is also higher than those for most other countries. The returns and consumption growth series are deflated using each countries consumer price index. Denmark and Portugal are omitted because data only starts in 1999 and 1993, respectively.

that assets with high payoffs in bad states of the world (when consumption is low) are more desirable. This can be formalized as follows:

$$\mathbb{E}[R_{it}] - \gamma_0 = \gamma_1 \beta_{c,i}, \tag{4.5}$$

where R_{it} is the return to a risky asset, γ_0 is the return to a portfolio not correlated with consumption growth (the zero consumption beta portfolio), γ_1 is the price of risk, and $\beta_{ci} = \text{cov}(R_{it}, c_t)/\text{var}(c_t)$ with c_t as the growth rate of aggregate consumption per capita is the measure of risk (Breeden, Gibbons, and Litzenberger 1989). When the CCAPM holds, the expected return to an asset or portfolio is linear in its consumption beta. Therefore, a lower return on German foreign assets could be justified in terms of the CCAMP by a low consumption beta.

For this purpose, we compute the covariance between consumption growth and the investment returns discussed above. The CCAPM in equation (4.5) refers to spot consumption, yet empirically only period average consumption can be measured. In this study, we follow convention and choose the interpretation of consumption data as measuring consumption at the beginning of the period. Therefore, we compute consumption growth by dividing next period's consumption by current period consumption (Campbell 1999). The correlation between consumption growth and returns

Figure 4.12: Correlation of real consumption and real return, Germany, 10-year rolling windows, 1971 to 2017

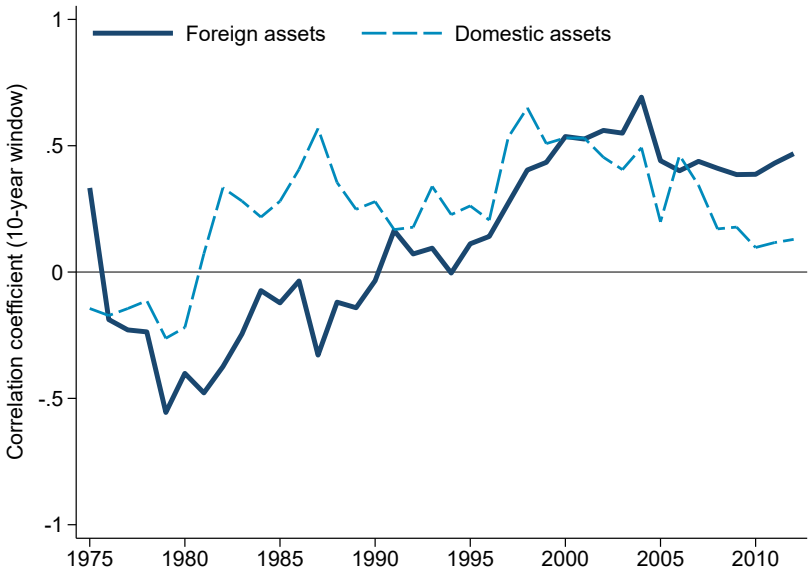

Notes: The correlation between Germany's returns on foreign assets and real consumption growth is positive in most years, especially since the 1990s. In recent years, the correlation is higher for foreign returns compared to domestic returns. Correlation coefficients plotted at the 5th year of the window. Correlation coefficient computed for consumption growth and real returns, both deflated using the German consumer price index.

is calculated using the return in period t and consumption growth between period $t+1$ and t. Data on consumption growth is taken from the Macro History Database by Jordà et al. (2017), which includes data until 2016.[25]

Table 4.9 provides two crucial insights. First, the returns on German foreign assets are more strongly correlated with domestic consumption growth than a bundle of domestic German assets. In other words, they provide *less* consumption insurance than a domestic German portfolio and their low returns are not justified by the consumption insurance that the asset provides. Second, also in comparison with other countries, the correlation of the German portfolio with German consumption appears high. The key upshot is that low German returns compared to other countries are not justified by their consumption insurance properties.

Moreover, if we take a closer look at the time path of the correlation between foreign asset returns and German domestic consumption, we find that the surge in German capital exports in recent years has gone hand in hand with a loss of consumption insurance (Figure 4.12). The correlation of foreign returns with domestic consumption growth has increased in recent years, not decreased. German foreign assets do not only

[25]The data for 2017 are from the World Development Indicators Database of the World Bank.

Figure 4.13: Geographical distribution of foreign assets by category

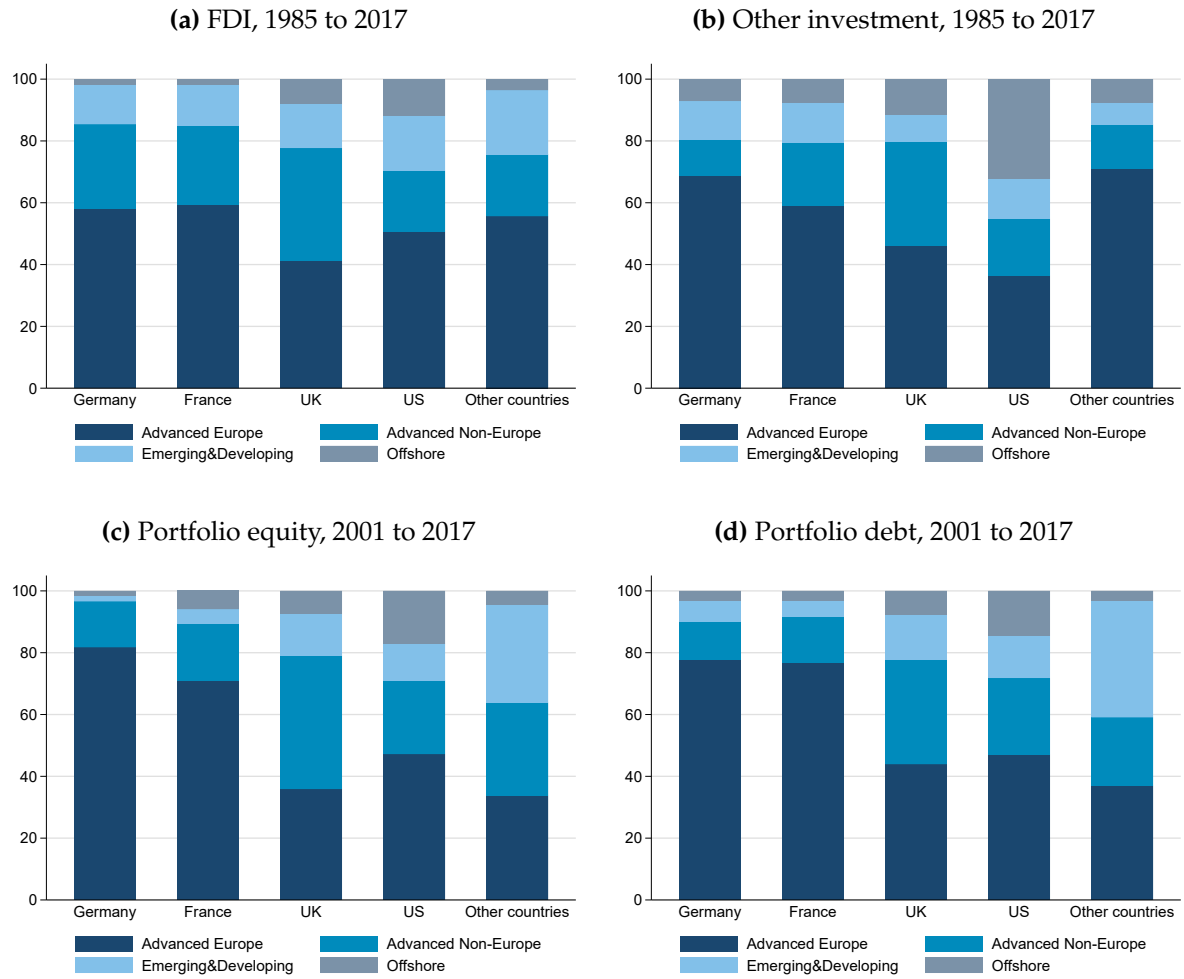

Notes: These graphs show that Germany's geographical allocation of foreign assets differs from that of other countries. Germany invests more in advanced European countries and less in developing countries. Germany's bias towards European investments is especially large for portfolio holdings (Panels (c) and (d)). Offshore countries are classified following the Bundesbank's list of offshore banking centers.

have low payoffs overall, they are also not helping to smooth consumption, and these trends have become worse in the past decade of record capital exports.

4.7.2 A hedge against demographic risks?

Germany faces increasing demographic risk from an aging population. Investing in countries with younger populations may help to hedge this risk and facilitate intertemporal income smoothing. However, the data show that German assets are predominantly invested in other advanced countries with aging populations, especially into other European countries, and increasingly so.

It is particularly remarkable that the share of German investments to younger and more dynamic developing countries and emerging markets has decreased rather than

Figure 4.14: Geographical distribution of German foreign assets, 1985 to 2017

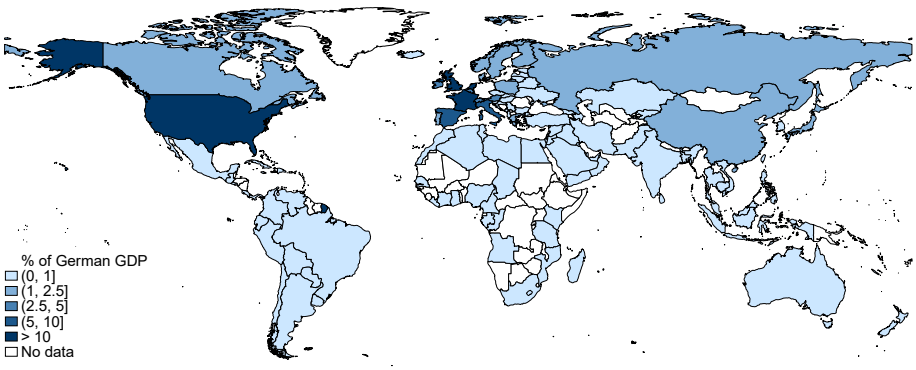

Notes: This map shows the geographic distribution of German foreign assets. The stock is expressed in % of German GDP held in each country and averaged between 1985 and 2017.

Figure 4.15: Old age population (% above 64 years), average, 1985 to 2015

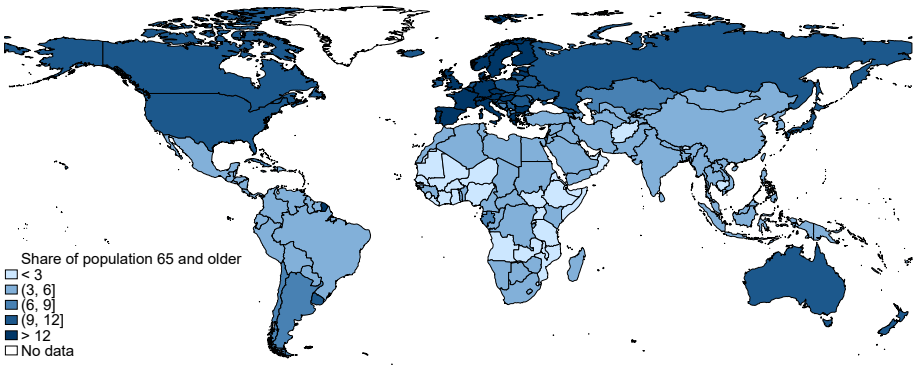

Notes: This map shows the share of each country's population aged 65-years and older, averaged between 1985 and 2015 (more recent data not available). Data from the UN Population Division.

increased, from 15-20% in the 1980s to about 10% in 2017. This drop has occurred despite the fact that developing countries such as China or India have seen record growth rates and that the emerging world now accounts for more than 50% of world GDP.[26] In other words, the "home-bias" of German investments in favor of European investments has intensified and the potential for demographic risk hedging has decreased accordingly.

The preference for investing into aging economies is more pronounced in Germany than in other countries. This can be seen in Figure 4.13 which compares the geographical allocation of foreign assets. While Germany's preference for European investments is observable for all asset categories, the focus on the euro area is especially strong for

[26]Note that before 2001 no data for portfolio assets is available and that both portfolio asset types have a low exposure to emerging and developing markets. However, this does not affect the overall dynamics because the share of portfolio investment in the total position was relatively small before 2000. When assuming a 1% exposure of equity assets and a 5% exposure for debt assets (based on the averages in the early 2000s), the exposure of total assets is still 15%. For FDI exposure to emerging markets actually increased over time but this effect is outweighed by the increase in portfolio investment in the aggregate.

portfolio equity and debt (Panels (c) and (d)). Specifically, the majority of Germany's equity assets are located in Northern euro-area countries, while the majority of debt securities were in Southern euro-area countries until recently. Now debt securities are distributed roughly equally across all member states.

Another way to explore the relevance of demography is to plot world maps. Figure 4.14 shows the geographic distribution of Germany's foreign assets across the world between 1985 and 2017. It is estimated as discussed in Section 4.2. Assets are scaled by German GDP.

The map reveals once more that Germany's foreign investments flow predominantly to Europe and other advanced countries. These are also the countries with a population structure most similar to Germany. Figure 4.15 shows the average share of the people aged 65 and older in total population between 1985 and 2015. It is clearly visible that the two maps overlap, as Germany has almost no investments in countries with a younger population structure. This illustrates that Germany's large stock of foreign assets does little for hedging against demographic risks.

4.8 Conclusion

Germany is world champion when it comes to exporting savings. In this chapter, we study the financial returns on German foreign investment. We find that the reputation of German households, firms, and banks for being bad foreign investors is mostly justified. German returns are substantially lower than those of other countries across asset classes. Moreover, foreign returns are consistently lower than domestic returns and the geographic distribution does not support the argument that the country's foreign investments hedge against demographic trends. The overwhelming share of German foreign investments is located in other industrial countries with similar demographic profiles.

We find that the underperformance of German foreign investment is particularly pronounced for equity and foreign direct investments. Importantly, the lower returns are not explained by a different risk profile of German investments. The correlation of foreign returns with domestic consumption is higher, not lower, than in other countries and higher than for domestic returns. In other words, foreign assets provide very little consumption insurance to German households. Overall, the low returns raise serious doubts with regard to the capacity of German households and the German financial sector to allocate Germany's substantial savings exports in a beneficial way.

Appendices to Chapter 4

D1 Additional tables and figures

Figure D.1: The world's largest net creditors, 2017

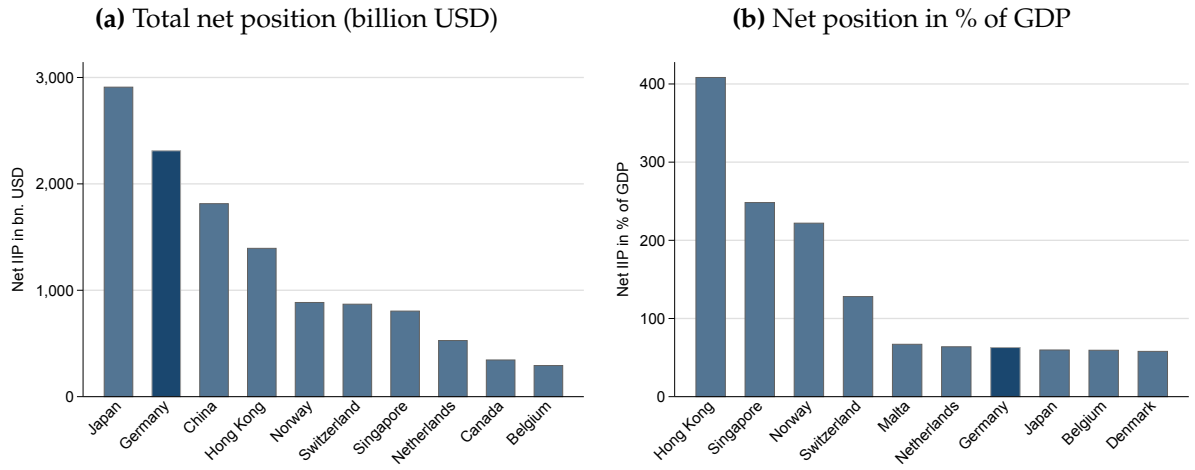

Notes: The net position is the difference between foreign assets (assets held abroad) and foreign liabilities (domestic assets held by foreigners). In USD terms Germany's net position is only exceeded by Japan. The graph excludes small oil exporting countries with large net positions in % of GDP. Data on asset positions are from the Bundesbank and the IMF, data on GDP from the World Bank.

Figure D.2: Germany fixed effect over time (country dummy)

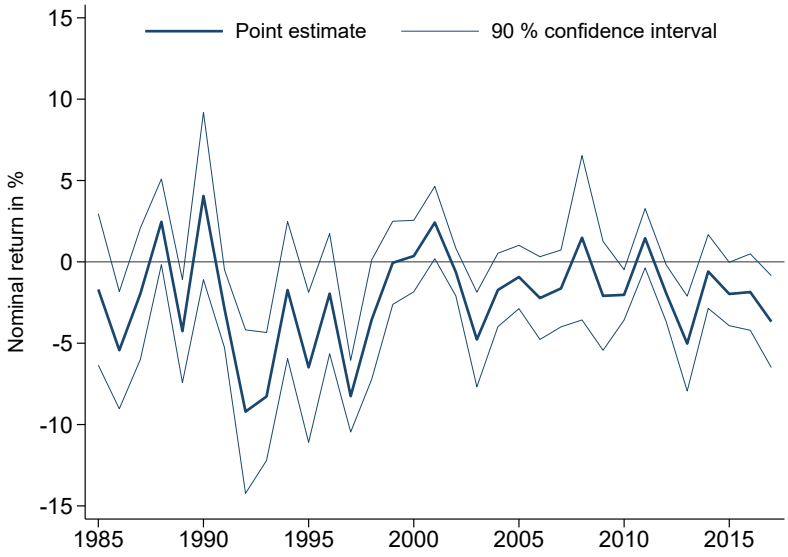

Notes: This graph plots the coefficient for the Germany dummy in each year since 1985 (dark blue line) as well as its 90% confidence interval (light blue lines). The regression builds on column (1) in Table 4.4 and controls for net foreign assets, the financial account balance and year effects. The main take away is that the underperformance of German foreign returns has been relatively persistent over time.

D2 Additional results for nominal and real returns

D2.1 Nominal returns on German foreign assets and domestic assets

When analyzing German returns on foreign assets and comparing them to domestic assets in Section 4.4, we focused on real returns. Here, we also show nominal returns over time.

Figure D.3 plots the nominal return, yield, and valuation changes (in Panel (a)) as well as returns by asset categories (in Panel (b)). Panel (a) confirms that Germany saw many years of valuation losses not just in real terms, as shown in the man text, but also in nominal terms. Panel (b) reveals that average nominal returns on FDI were positive but close to zero in contrast to the negative real returns before the 2000s. The other patterns in returns by category are similar for nominal and real returns.

Figure D.4 shows the comparison of the return on German foreign assets to the return on German domestic assets as measured by Jordà et al. (2019) for the nominal case. Compared to Figure 4.8 in the main text, the gap between nominal returns on foreign vs. domestic assets is larger than that with real returns in the 1960s and 1970s, mainly due to higher inflation. The overall picture, however, is similar for the real and nominal series.

Figure D.3: Nominal returns on German foreign assets, 5-year rolling means, 1950 to 2017

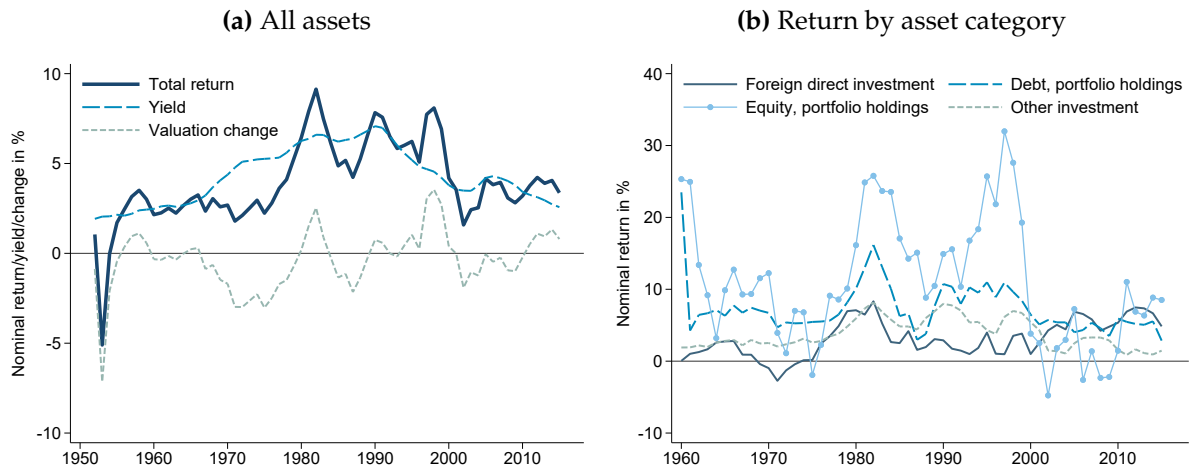

Notes: Figure shows same results as Figure 4.7 for nominal returns. In addition, in nominal terms Germany saw absolute losses in many periods. Rolling arithmetic averages computed over 5-year windows and plotted at the third year of the window. Returns estimated as discussed in Subsection 4.2.1.

Figure D.4: Nominal returns in comparison, 5-year rolling means, 1975 to 2017

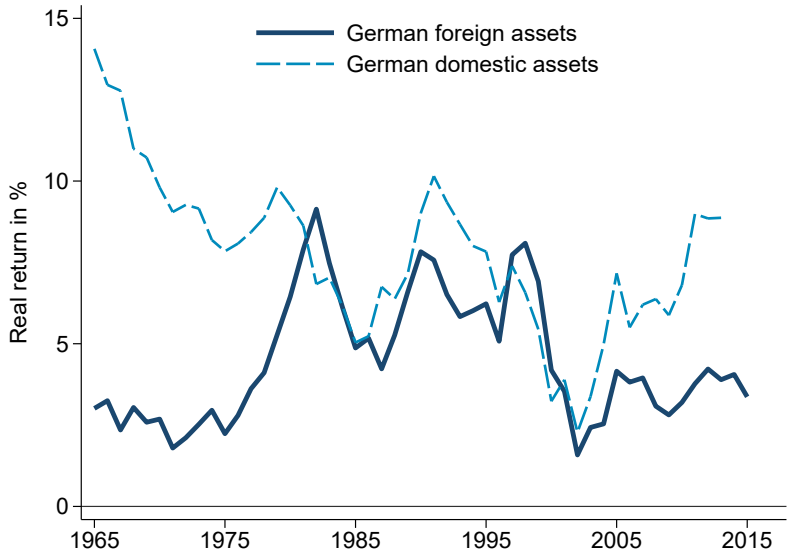

Notes: Figure shows the same comparison as Figure 4.8 for nominal returns. Given that both series are deflated using the same price index, the same results emerge. Rolling means computed over 5-year windows and plotted at the third year of the window. Return on foreign assets estimated as discussed in Subsection 4.2.1. Return on German domestic assets from Jordà et al. (2019).

D2.2 Real returns on other countries' foreign assets

When comparing German returns to other countries' returns in Section 4.5, we focus on nominal returns to abstract from different inflation dynamics across countries. Our focus on nominal returns is motivated by the idea that we want to compare investment performance on the same global level playing field – before country-specific factors (such as inflation) are taken into account. Using nominal returns, Germany ranked 12th among the 13 countries we consider. For completeness, we now also show results for real returns.

Table D.1 shows that Germany ranks 9th when considering real returns. The higher ranking is due to countries like Italy or Spain experiencing much higher rates of inflation especially in the earlier part of the sample. Therefore, in terms of their own price level, foreign returns for Italy and Spain appear lower than Germany's real returns abroad. Figure D.5 confirms that Germany compares more favorably when looking at real returns, owing to the country's relatively low domestic inflation rates.

Table D.1: Real returns on foreign assets, 1975 to 2017

	Rank	1975–2017	1999–2017	2009–2017
United States	1	6.61	5.75	7.74
Canada	2	5.21	3.01	7.40
United Kingdom	3	4.88	3.63	1.85
Sweden	4	4.63	5.22	4.60
Netherlands	5	3.77	2.95	4.55
Norway	6	3.57	3.73	5.24
Denmark	7	3.48	3.48	5.58
France	8	3.34	2.62	3.48
Germany	9	2.56	2.28	2.62
Portugal	10	2.44	0.69	1.17
Italy	11	1.88	1.50	3.18
Spain	12	1.29	-0.07	1.82
Finland	13	0.31	2.36	2.19

Notes: This table shows the same results as Table 4.1 for real returns instead of nominal returns. Nominal domestic currency returns are deflated using each country's own consumer price inflation (from Macro History Database (Jordà et al. 2017) and World Bank.) Data for Denmark starts in 1999 and data for Portugal starts in 1993.

Figure D.5: Real returns in comparison, 5-year rolling means, 1975 to 2017

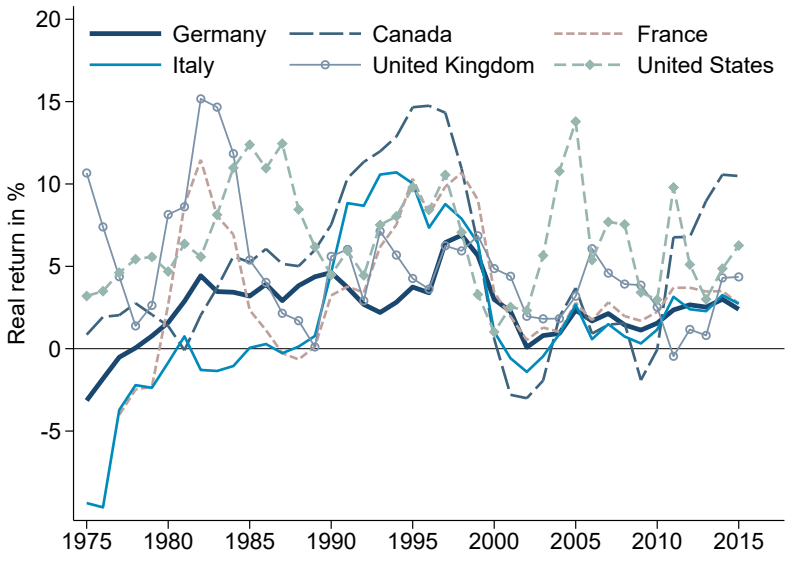

Notes: This graph shows the same comparison as Figure 4.9 for real returns instead of nominal returns. Nominal domestic currency returns are deflated using each country's own consumer price inflation (from Macro History Database (Jordà et al. 2017) and World Bank.) German returns are low compared to the other G7 members (excluding Japan due to data availability) also in real terms. Rolling means computed over 5-year windows and plotted at the third year of the window. For more details see Section 4.2.

D3 Comparison to earlier work on German foreign returns

D3.1 Overview of the literature

Several earlier studies have computed returns on German foreign assets. In this appendix, we provide an overview of this work and compare the results to ours.

The overall take away is that all studies, including ours, use similar data sources and methodology. Moreover, all earlier papers use a more limited time sample compared to our study and no previous paper conducts decomposition exercises or international comparisons as we do. Most existing papers also focus on the return differential, i.e., the difference between the return on assets and liabilities.

Table D.2 provides a concise summary of earlier estimates. For transparency, we compare the result of each study to our own using the exact same sample and variables

Table D.2: Results of other studies and comparison with our results

Authors	Data	Other valuation changes included?	Return measures and sample used	Results	Our results with same sample
Lane and Milesi-Ferretti (2007a)	External Wealth of Nations for assets, BOP	Yes	Real return on assets, 1995–2000	5.4%	6.1%
Habib (2010)	BOP and IIP before most recent revision in data	Yes	Real return differential, 1981–2007	-1.07	-0.99
Bundesbank (2014), Frey et al. (2014)[1]	BOP and IIP before most recent revision in data	No	Nominal return on assets, 2005–2013	4.0%	3.8%
Baldi and Bremer (2015) (DIW Berlin)	BOP and IIP before most recent revision in data	Yes	Nominal return differential, 1993–2012	-1.5	-0.5
Bundesbank (2018)	BOP and IIP based on BPM6	No	Nominal return on assets, 2008–2017	3.7%	3.3%
Fiedler et al. (2018) (Kiel Institute)	BOP and IIP based on BPM6	No	Nominal return on assets, 2005–2016	3.9%	4.0%

Notes: This table provides an overview of studies that have produced estimates of Germany's return on foreign assets. It describes the data and methodology used and compares the other authors' results to our results. One reason for differing results is the treatment of "Other valuation changes" as discussed in the following Appendix D3.2.

[1] Frey et al. (2014) is a more extensive version of Bundesbank (2014).

used in the respective earlier study.

The table reveals that our results are very similar overall. There are two main reasons for any remaining discrepancies. The first is differences in the data, since we use the most updated and cleaned series provided by the Bundesbank. Papers published before 2015 use data based on the old balance of payments manual, meaning prior to the many revisions that came with the introduction of the new manual (BPM6). The second reason is the treatment of valuation changes due to other adjustments. In particular, the more recent studies (like Bundesbank 2018) exclude these changes, while we include them, since there is no strong argument not to do so and since we lack a data breakdown to exclude these changes before 2005. We discuss the impact of this choice in more detail in the following section and conclude that the inclusion or exclusion of "other valuation changes" does not affect the overall results much, especially not for our international comparisons.

D3.2 Treatment of valuation changes due to other adjustments

As discussed in Subsection 4.2.1, one open issue in the computation of foreign returns regards the treatment of valuation changes due to other adjustments, meaning residual valuation changes, which can neither be attributed to exchange rate nor price movements. There are basically two options. First, one can simply exclude all residual valuation changes (see e.g., Bundesbank 2014, 2018; Frey et al. 2014). Second, one can include all or parts of these changes. Lane and Milesi-Ferretti (2009), for example, suggest including valuation changes due to other adjustments for FDI but not for portfolio investment. For 'other investment', in turn, they suggest adding it to the initial asset position.

Given that these options are pretty much arbitrary, and since the Bundesbank does not provide a data breakdown for the period before 2005, we generally include valuation changes due to other adjustments in our baseline results for Germany and other countries alike but directly adjust for known statistical breaks and discrepancies (see Section 4.2). We choose this option because the valuation changes due to other adjustments only pose a problem as far as they are due to statistical issues instead of actual changes in the position. However, we now check what impact the complete or partial exclusion would have on our results. Specifically, using Bundesbank data on the split of valuation changes after 2005, we can illustrate the effect of including or partly including these other changes compared to other method.

Panel (a) of Figure D.6 shows that excluding valuation changes due to other adjustments for the period after 2005 does not make much of a difference when appropriately adjusting for known issues. The size of the change depends on the time period cov-

Figure D.6: Robustness of estimated nominal returns, 2005 to 2017

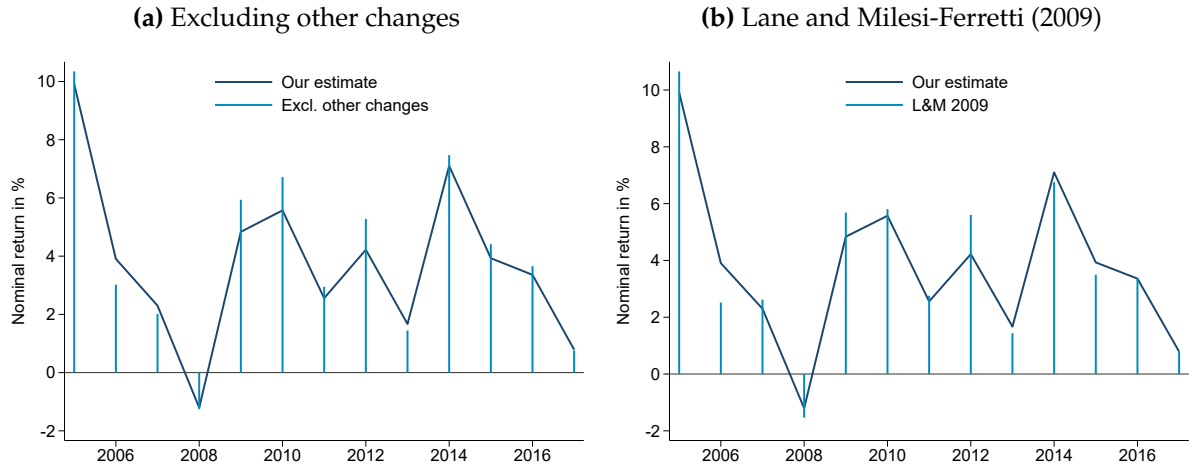

Figure D.7: Valuation changes due to other adjustments, 2005 to 2017

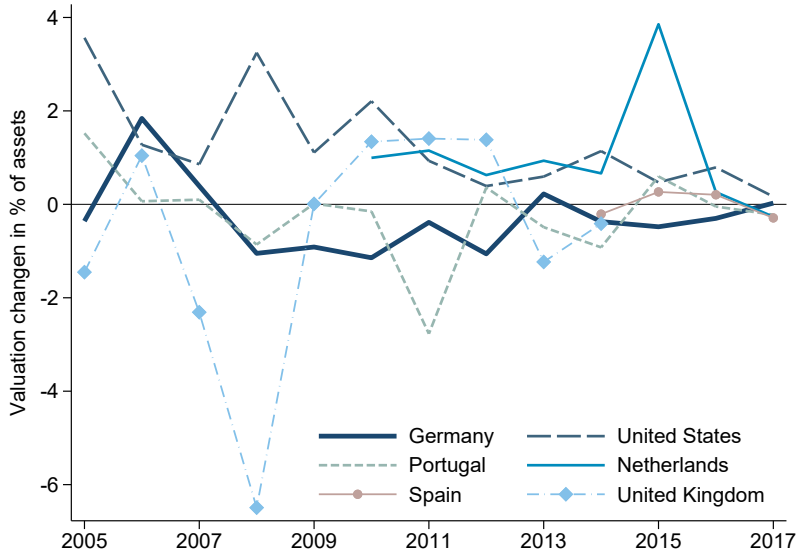

Notes: This graph shows valuation changes due to other adjustments for Germany and other countries. Data retrieved from the statistical institutions responsible for the IIP (Bundesbank (DE), Bureau of Economic Analysis (US), Banco de Portugal (PT), Dutch Central Bank (NL), Banco de Espana (ES), Office of National Statistics (UK)). The UK estimate includes financial derivatives.

ered. In the pre-crisis years 2005-2007, for example, you get *lower* German returns if you exclude 'other adjustments', while the returns are *higher* in the crisis years 2008-2010 without those changes. Panel (b) of Figure D.6 follows the recommendation of Lane and Milesi-Ferretti (2009) of excluding 'other adjustments' only for some asset categories. The takeaway is the same. In some years, the estimates are higher, while in others, they are lower. In aggregate, for all years for which there is comparison data

(2005-2017), Germany's nominal returns are 0.29 percentage points lower when excluding other valuation changes (4.06% with valuation changes due to other adjustments versus 3.77% without).

Does the choice make a difference for the international comparisons? To assess this, we retrieved data on valuation changes due to other adjustments for several countries and check how much their exclusion affects the results in comparison to Germany.

Figure D.7 illustrates the impact of valuation changes due to other adjustments in percent of assets since 2005 for Germany, the US, the UK, the Netherlands and Spain. On average, German valuation changes were slightly below zero, at -0.28%. The averages for the US and the Netherlands are positive while the average for the United Kingdom is negative. The Spanish average over the short time span available is roughly equal to zero. Thus, taken together, the discrepancy is not large enough to explain the observed differences between Germany and the other countries.

D4 Data sources and classifications

This appendix lists the data sources used for (i) estimating returns across countries (Table D.3), to compute (ii) valuation changes due to exchange rates (Table D.4), as well as on (iii) the geographical distribution of assets (Table D.5). The tables also provide information on data availability and country specific data issues. Finally, Table D.6 shows how our country groups are defined.

Table D.3: Data sources for return computation of other countries

Country	Coverage	IIP sources	Notes
Canada	1970–2016	EWN until 1989; IMF	
Denmark	1999–2016	IMF	Earlier data available but with gaps.
Finland	1975–2016	EWN until 1989; IMF	
France	1975–2016	EWN until 1993; IMF	PF not part of return before 1988 (no capital income data); capital income on PF debt in 1993 reported in Franc instead of Euro; jumps due to category changes in FDI in 1999 removed from change in assets.
Italy	1972–2016	EWN until 2003; IMF	FDI and PF not part of return before 1980 (no capital income data).
Netherlands	1970–2016	EWN until 1981; IMF	No PF assets data in EWN. No return in 2003 due to break in data.
Norway	1975–2016	EWN until 2006; IMF	PF not part of return before 1992 (no capital income data).
Portugal	1993–2016	EWN until 1995; IMF	Earlier data available but with gaps.
Spain	1975–2016	EWN until 2000; IMF	PF not part of return before 1990 (no capital income data).
Sweden	1970–2016	EWN until 1981; IMF	PF not part of return before 1997 (no capital income data).
United Kingdom	1970–2016	EWN until 1981; IMF	PF not part of return before 1984 (no capital income data).
United States	1970–2016	EWN until 1981; IMF	PF not part of return before 1986 (no capital income data).

Notes: EWN = External Wealth of Nations database by Lane and Milesi-Ferretti (2007b).

Table D.4: Data sources on the currency composition of assets

Country	OECD FDI	CPIS (currency[1])	LBS	Reserves
Canada	1985–2017	2014–2017	1977–2017	1999–2017, Bank of Canada, USD
Denmark	1991–2017	2001–2017	1977–2017	2005–2017, Danmarks Nationalbank
Germany	1985–2017	2007–2017	1977–2017	1949–2017, Bundesbank, only USD until 1999
Finland	1992–2012, 2016–2017	2013–2017	1983–2017	2002–2017, Bank of Finland
France	1987–2017	2001–2017	2007–2017	2001–2017, IMF, (based on Banque de France reports assume all in USD)
Italy	1985–2017	2001–2017	2011–2017	2005–2017, Banca d'Italia
Netherlands	1985–2017	2009–2017	2014–2017	1997–2017, Dutch central bank
Norway	1988–2017	only countries	2014–2017	1997–2017, Norges Bank
Portugal	1995–2017	2001–2017	2009–2017	2001–2017, Banco de Portugal (2011–2017 data, 2001–2010 assume 90% of reserves is in USD)
Spain	2000–2017	2007–2017	2014–2017	1999–2017, Banco de Espana
Sweden	1986–2017	2003–2014	1977–2017	1999–2017, Swedish Riksbank
United Kingdom	1987–2017	only countries	1977–2017[2]	1997–2017, Bank of England
United States	1985–2017	2003–2017	1998–2017	1999–2017, except 2001, US Department of the Treasury (US International Reserve Position report)

Notes: CPIS = Coordinated Portfolio Investment Survey LBS = Locational Banking Statistics.
[1] Country breakdown always available starting in 2001.
[2] Data for 1982 is missing.

Table D.5: Data sources on the geographical distribution of assets

Country	OECD FDI	CPIS	LBS
Canada	1985–2017	2001–2017	2007–2017
Denmark	1991,1994,1998,1999–2017	2001–2017	1977–2017
Germany	1985–2017	2001–2017	1977–2017
Finland	1992–2012, 2016–2017	2001–2017	1983–2017
France	1987–2017	2001–2017	1977–2017
Italy	1985–2017	2001–2017	2014–2017
Netherlands	1985–2017	2001–2017	1977–2017
Norway	1988–2017	2001–2017	no data
Portugal	1995–2017, 2006 missing	2001–2017	no data
Spain	2003–2017	2001–2017	2014–2017
Sweden	1986–2017	2001–2017	1977–2017
United Kingdom	1987–2017	2001–2017	1977–2017, 1982 missing
United States	1985–2017	2001–2017	1977–2017

Notes: CPIS = Coordinated Portfolio Investment Survey, LBS = Locational Banking Statistics.

Table D.6: Country groups used for geographical distribution

Group	Countries			
Advanced Europe	Austria	Belgium	Cyprus	Czech Rep.
	Denmark	Estonia	Finland	France
	Germany	Greece	Iceland	Ireland
	Italy	Latvia	Lithuania	Luxembourg
	Malta	Netherlands	Norway	Portugal
	San Marino	Slovak Rep.	Slovenia	Spain
	Sweden	Switzerland	UK	
Advanced Non-Europe	Australia	Canada	Japan	New Zealand
	Rep. of Korea	Macao	Taiwan	United States
Offshore	Antigua&Barbuda	Bahrain	Barbados	Belize
	Dominica	Grenada	Hong Kong	Lebanon
	Liberia	Marshall Isl.	Mauritius	Panama
	Philippines	Samoa	Seychelles	Singapore
	St. Kitts&Nevis	St. Lucia	Bahamas	Vanuatu
	St. Vincent&the Grenadines			
Emerging&Developing	Remaining countries			

Notes: Choice of offshore countries based on Bundesbank list of offshore banking centers.

D5 Estimation of valuation changes due to exchange rates

In this section, we compare our estimates of valuation changes due to exchange rate movements (as discussed in Subsection 4.2.1) with official data for several countries. In addition, we also compare our estimates to an alternative set of estimates that is based on the currency composition data provided for total assets (not asset categories) by Bénétrix et al. (2015). The comparisons focus on the period after 2002 when the CIPS portfolio data become available, so that all asset classes can be included. The main take away from this exercise is that the estimates are similar across sources.

D5.1 Germany

Figure D.8 shows the comparison of valuation changes due to exchange rates for all asset classes for Germany. As explained, our aggregate comparison starts in 2002 (Panel (a)). Panels ((b)-(f)) start in 2005, which is when Bundesbank estimates by asset class become available. Overall the results are very similar, except for the reserves category in 2008 and 2009. This result suggests that the publicly available data on the currency composition of reserves is not complete or that the Bundesbank uses detailed non-public data, which allows for much more precise estimate. In any case, reserves only make up a small part of total assets, so the effect of this mismatch is not large.

D5.2 Netherlands

The Dutch Central Bank (DNB) publishes a breakdown of valuation changes starting in 2010. Panel (a) of Figure D.9 shows that also for the Netherlands our approach delivers fairly similar estimates despite the lack of CPIS currency data before 2009 and the lack of any LBS data before 2014.

D5.3 Portugal

The Banco de Portugal publishes a breakdown of valuation changes since 1999. Panel (b) of Figure D.9 shows that our approach delivers fairly similar estimates despite the lack of any LBS data before 2009. However, in the first three years when no CPIS data is available the estimates differ substantially since our aggregate measure includes only FDI in this period.

D5.4 Spain

The Spanish Banco de Espana (BdE) publishes time series starting only in 2014. For the four overlapping years our estimate is only slightly lower than the published series,

Figure D.8: Valuation changes due to exchange rates, our estimate vs. official statistics, Germany

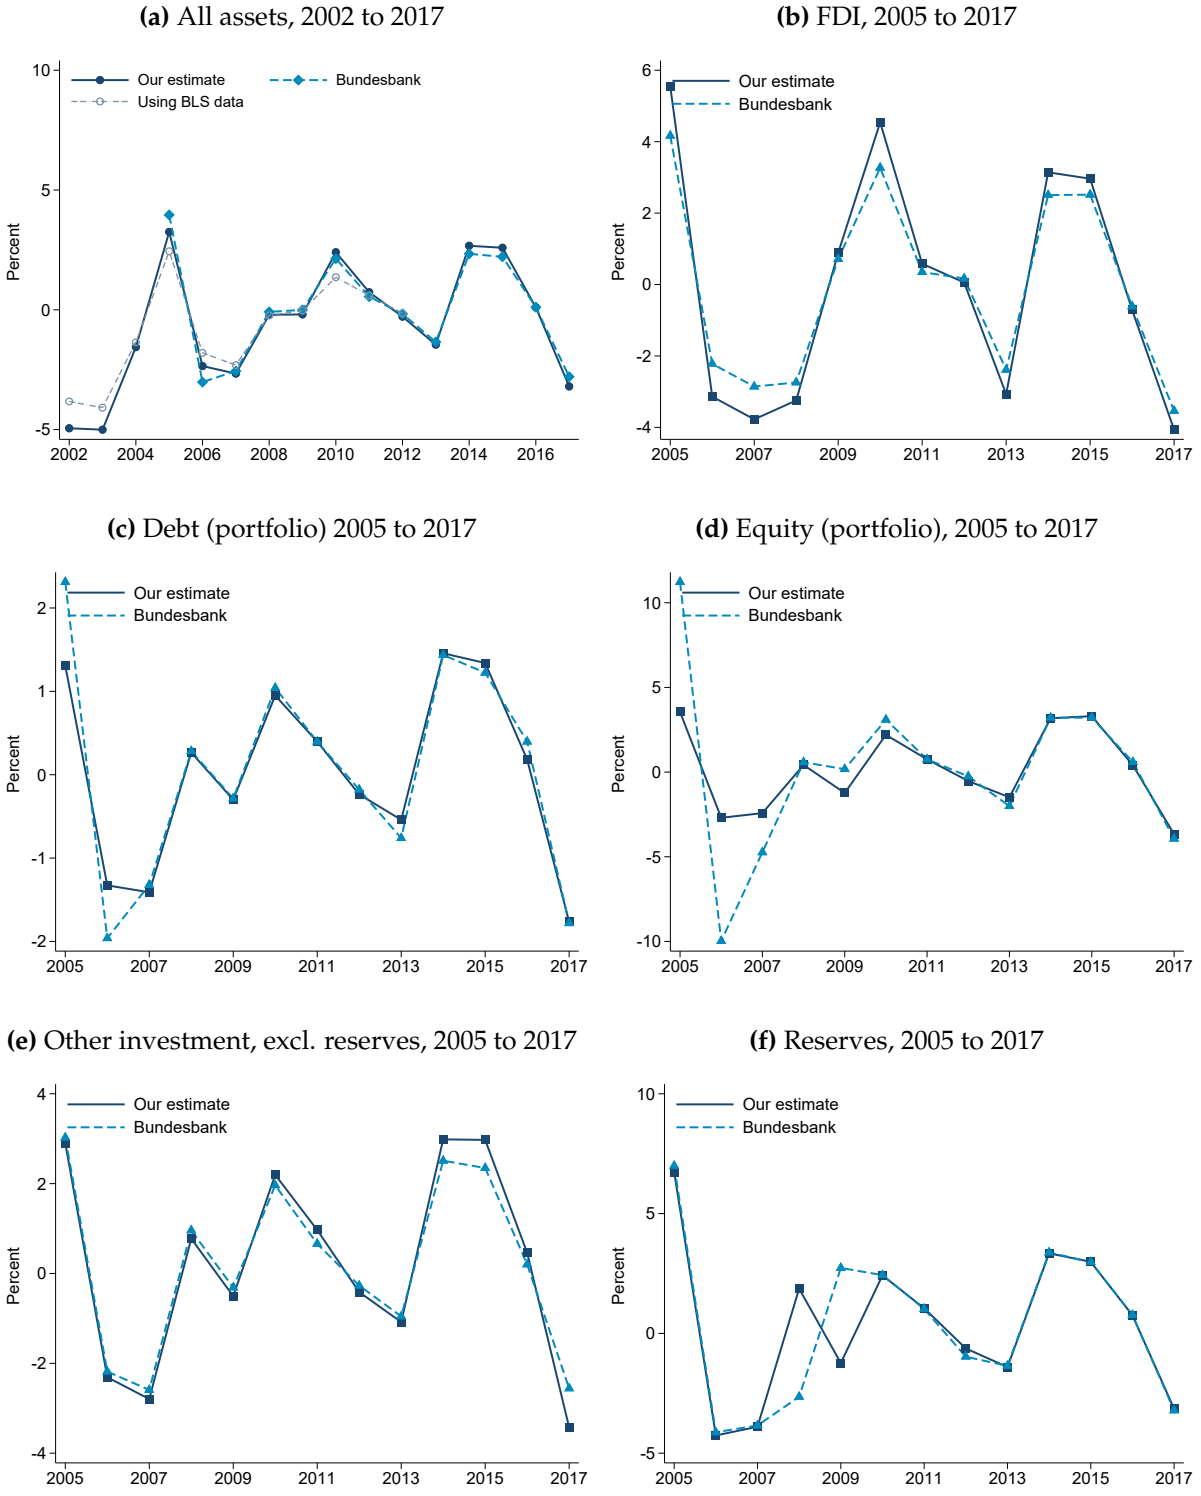

Notes: This graph compares our estimates of valuation changes due to exchange rates to those published by the Bundesbank. We find few discrepancies. The estimates using the aggregate data on total assets from Bénétrix et al. (2015) are also fairly similar. BLS refers to Bénétrix et al. (2015).

see Panel (c) of Figure D.9. The result using BLS data is quite similar.

D5.5 United Kingdom

In an article titled "Analysis of the UK's international investment position: 2016"[27] the Office of National Statistics (ONS) published estimates of the valuation changes due to exchange rate effects for 2000 to 2014. The estimation is based only on US Dollar, Euro and Japanese Yen exchange rates, uses country shares to approximate currency shares and includes financial derivatives. Panel (d) of Figure D.9 shows that this yields similar results as our more detailed approach as well as the estimates using BLS data.

D5.6 United States

For the US, the Bureau of Economic Analysis (BEA) publishes data on the valuation changes due to exchange rate movements since 2002. Panel (e) of Figure D.9 shows that also the estimates for the US are very close to the published series. This is reassuring since we need to interpolate the US currency composition of 'other investment' between 1998 and 2012.

[27] The article is available on the ONS website: `https://www.ons.gov.uk/economy/nationalaccounts/balanceofpayments/articles/analysisoftheuksinternationalinvestmentposition/2016`.

Figure D.9: Valuation changes due to exchange rates, our estimate vs. official statistics, other countries

(a) Netherlands, all assets, 2002 to 2016

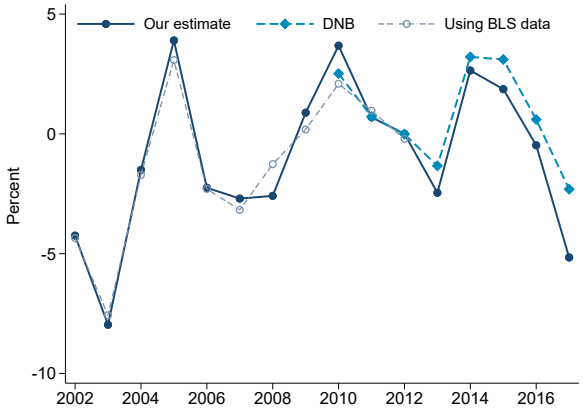

(b) Portugal, all assets, 1999 to 2017

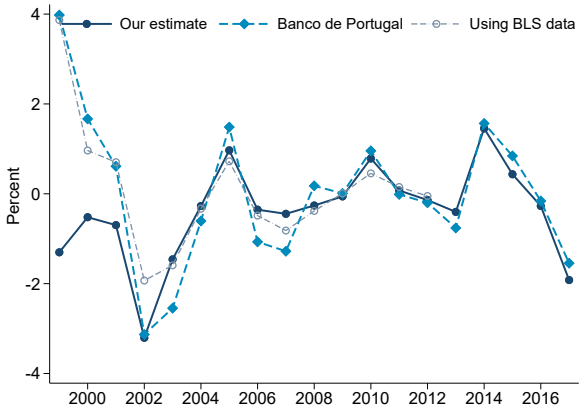

(c) Spain, all assets, 2002 to 2017

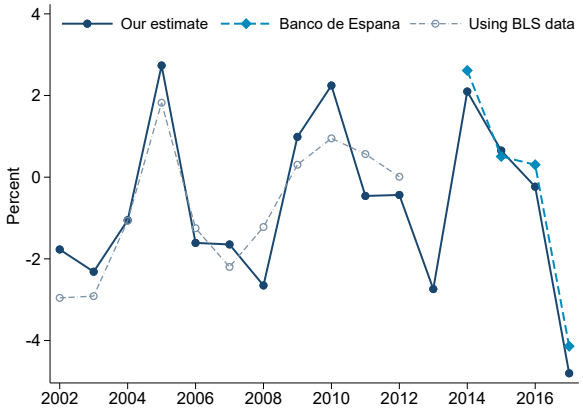

(d) United Kingdom, all assets, 2000 to 2014

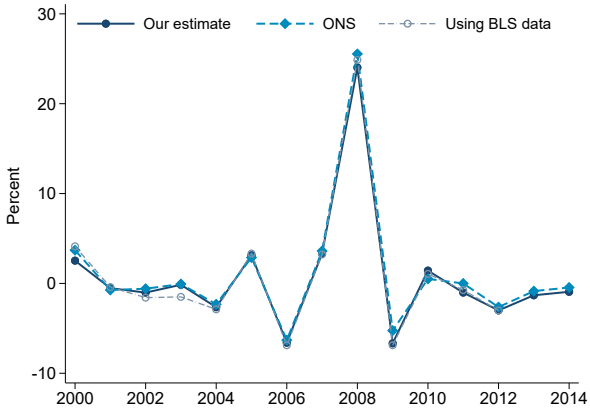

(e) United States, all assets, 2002 to 2017

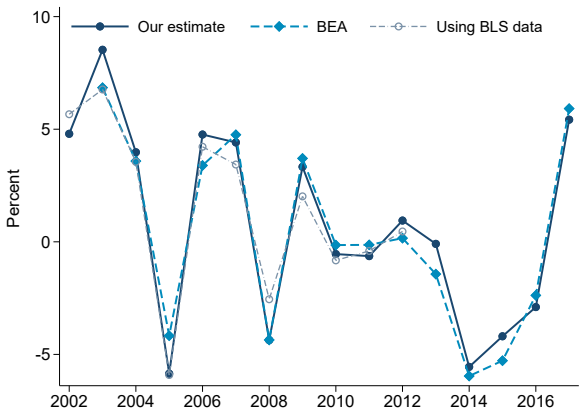

Notes: This graph compares our estimates of valuation changes due to exchange rates to the official estimates published by the respective central bank or related country authority. We conduct this comparison for all countries for which this data breakdown is made publicly available. Overall, we find few discrepancies. The estimates using data from Bénétrix et al. (2015) also produces similar results for aggregate assets. BLS refers to Bénétrix et al. (2015). DNB=Dutch Central Bank, ONS=Office of National Statistics, and BEA=Bureau of Economic Analysis.

Figure D.10: Composition of Germany's international investment position, foreign assets vs. liabilities, 1949 to 2017

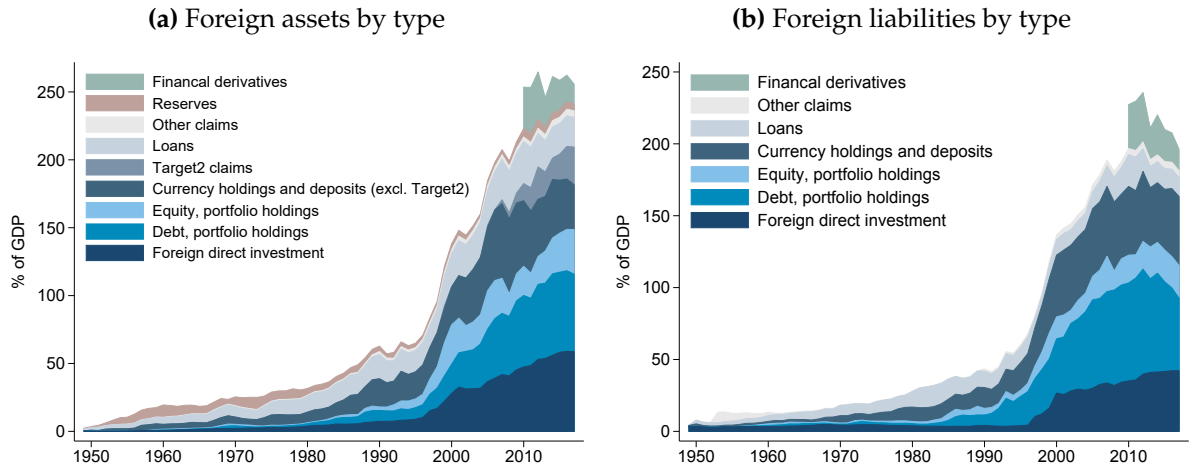

Notes: Data from Bundesbank. Detailed split of firms and households only available since 2012.

D6 Return differential between German assets and liabilities

To gauge the profitability of foreign investments, many earlier studies compare the returns on foreign assets to those earned by foreigners in the respective country (on foreign liabilities). While we prefer the comparison to other countries' returns and a broader measure of domestic returns, we also want to provide estimates of liability returns here. We will see that the difference between asset and liability return decreased in recent years. Indeed, it has turned positive over the past decade. This has led some commentators to argue that German investment performance abroad improved in recent years (Bundesbank 2014, 2018). Here we show that changes in the German differential is not the result of better foreign returns, but mainly driven by a reduction returns on liability FDI and liability portfolio debt. Both categories saw specific developments in the past decade, which complicate the comparison further.

To set the stage, Figure D.10 shows the composition of German foreign assets and liabilities. As can be seen, the asset composition of liabilities (Panel (b)) is rather similar to that of Germany's foreign assets (Panel (a)).

Panel (a) of Figure D.11 compares the returns on Germany's assets abroad to those on foreign liabilities since the 1950s. As can be seen, the return on assets was lower throughout the entire post-WW2 period, and until the 2000s. The past few years are the first time in which the difference turns (slightly) positive, on average.

For a more detailed comparison, Panels (b) and (c) of Figure D.11 show breakdowns into yields vs. valuation changes. The graphs reveal that the most recent shift is driven

Figure D.11: Real return, yield and valuation changes on German foreign assets vs. liabilities, 1950 to 2017

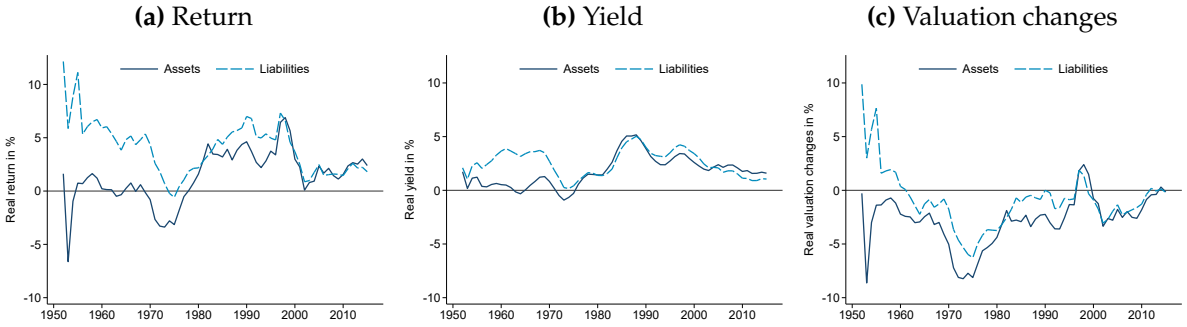

Notes: Rolling means computed over 5-year windows and plotted at the third year of the window. Return on foreign assets and liabilities estimated as discussed in Subsection 4.2.1 and deflated using the consumer price index.

by yields and not by valuation changes. Generally, there are only very few years in which the valuation changes on assets were larger than valuation changes on liabilities.

Finally, Figure D.12 displays a breakdown by asset class. To facilitate interpretation, we show the difference between asset and liability returns instead of both returns separately for each asset class. The figure shows that before the 2000s, the return differential was especially negative for FDI, but this differential has improved notably over the past two decades, with German FDI abroad showing higher returns that foreign FDI in Germany. In addition, the relative performance of portfolio debt has shifted since the early 2000s, with German portfolio assets yielding higher returns compared to German debt securities held by foreigners. In contrast, the differential remains negative for equity investments, often strongly so.

To understand the drivers of these differences, we now perform the same decomposition exercise as we did for foreign returns across countries (Subsection 4.6.3). Table D.7 shows that the differences within each asset class explain the overall difference, both historically and today. Table D.8 also illustrates the decomposition over time. It highlights that the difference due to composition did not decrease a lot since the 1970s while the within-asset class differences decreased substantially and turned positive since the 2000s. Finally, Figure D.13 shows the contribution of each type of assets, again differentiating by differences due to asset composition and differences due to returns within asset classes. The table shows that the negative values in the early sample are mainly driven by FDI and 'other investment'. The recent switch to positive differences was largely driven by FDI and portfolio debt.

To summarize, we find that the recent decrease in the gap between returns on foreign assets and foreign liabilities is mainly due to the relative changes in yields on FDI and debt portfolio holdings, especially because the return of foreigners investing

Figure D.12: Differences between real returns on foreign assets and liabilities by asset class, 1960 to 2017

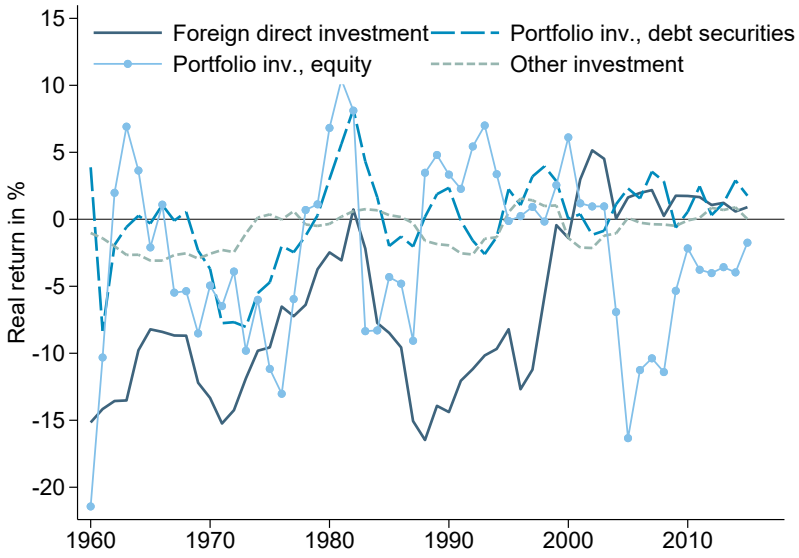

Notes: This figure shows differences between real returns (or yields or valuation changes) between German foreign assets and foreign liabilities. Rolling arithmetic means computed over 5-year windows and plotted at the third year of the window. Returns estimated as discussed in Subsection 4.2.1.

in Germany went down. This finding is consistent with the observation that the outflows of debt and FDI investments from Germany (both gross and net) have increased notably over the past 20 years, possibly due to a search-for-yield effect.

At the same time, there are important idiosyncratic effects that complicate the comparison of FDI and debt returns on assets versus liabilities. First, it is well known that, during and after the global financial crisis and the eurozone crisis, foreigners have purchased record amounts of German Bunds and other highly rated German debt securities, despite the fact that these had almost zero yields (Bundesbank 2017). This safe-haven effect improved the differential for debt securities. Second, inward FDI in Germany has become highly leveraged over the past two decades, mostly for tax shifting reasons. Today, many foreign companies load their German subsidiaries with debt to reduce the (high) effective tax rate on profits after interest. In comparison, German outward FDI is significantly less leveraged (Graf and Grimme 2017). This mismatch biases down the aggregate yields and returns reported by foreign-owned firms in Germany, even if the profitability of German firms has stayed the same (Ramb and Weichenrieder 2005). The increased within-firm tax shifting activities makes it hard to interpret the recent increase in relative yields on German FDI as improved investment performance. More research is needed to examine these effects in detail, ideally using micro-data on FDI and portfolio debt flows.

Table D.7: Decomposition of real return differential between foreign assets and liabilities, 1960 to 2017

		Difference due to	
	Difference in returns (pp.)	composition (asset class)	returns within asset class
1960-2017	-1.802	-0.463	-1.338
1985-2017	-0.859	-0.259	-0.600
1999-2017	0.185	-0.081	0.267
2009-2017	0.411	-0.268	0.679
1960-1969	-4.823	-1.186	-3.634
1970-1979	-3.342	-0.513	-2.826
1980-1989	-0.764	-0.331	-0.433
1990-1999	-1.642	-0.314	-1.328
2000-2009	-0.095	-0.249	0.153
2010-2017	0.266	-0.119	0.385

Notes: This decomposition splits the difference between returns on German foreign assets and German foreign liabilities into two parts: (1) asset class composition (using the same four broad asset categories used above), and (2) difference in returns within each asset classes. More details in Subsection 4.6.3.

Figure D.13: Decomposition of real return differential between foreign assets vs. liabilities, 1960 to 2017

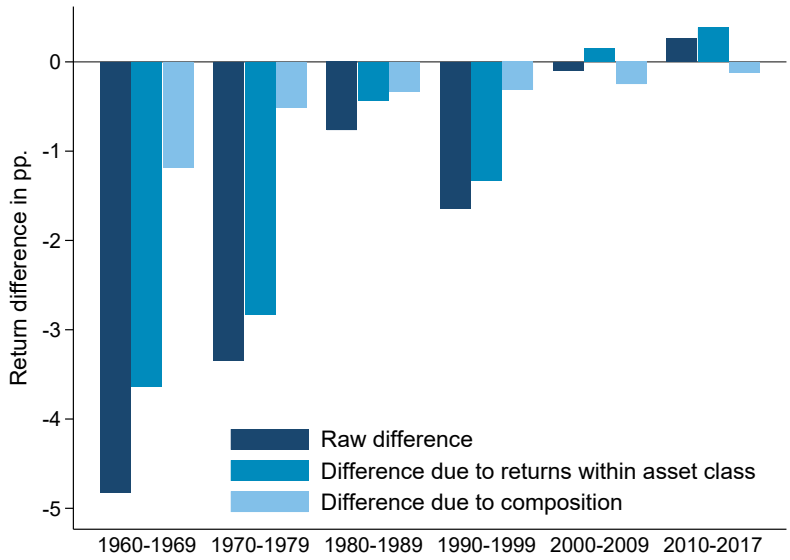

Notes: Decomposition splits difference between return on German foreign assets and German foreign liabilities into two parts: (1) different composition of asset position in the four broad asset categories, and (2) difference in returns within each asset classes (details in Subsection 4.6.3).

Table D.8: Contributions of different assets to differential between foreign assets and domestic assets (IIP), 1960 to 2017

	Total effect	FDI	Equity	Debt	Other inv.
Panel (a): Differences due to composition					
1960-2017	-0.463	-0.108	-0.194	-0.307	0.146
1985-2017	-0.259	0.156	-0.169	-0.475	0.229
1999-2017	-0.081	0.088	-0.002	-0.303	0.135
2009-2017	-0.268	0.130	-0.108	-0.307	0.018
Panel (b): Differences due to returns within class					
1960-2017	-1.338	-0.845	-0.094	0.128	-0.527
1985-2017	-0.600	-0.396	-0.167	0.244	-0.282
1999-2017	0.267	0.393	-0.290	0.393	-0.229
2009-2017	0.679	0.296	-0.238	0.516	0.105

Notes: This table shows contributions of individual asset classes to each of the components of the return differential. Components are (1) different composition of asset position in the four broad asset categories, and (2) difference in returns within each asset classes (details on decomposition in Subsection 4.6.3).

Figure D.14: Foreign asset returns and Target2 balances: Germany

(a) Share of Target2 in total assets, 1999 to 2017

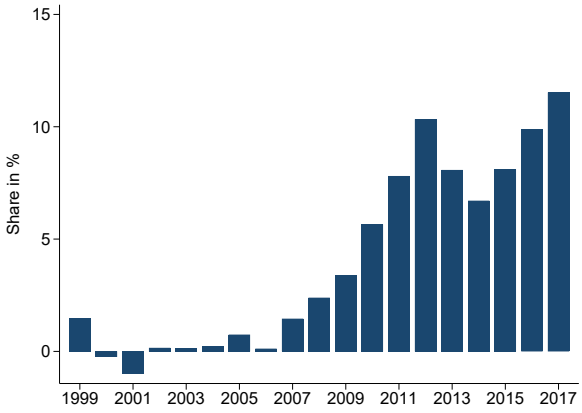

(b) Impact of Target2 on returns, 2007 to 2017

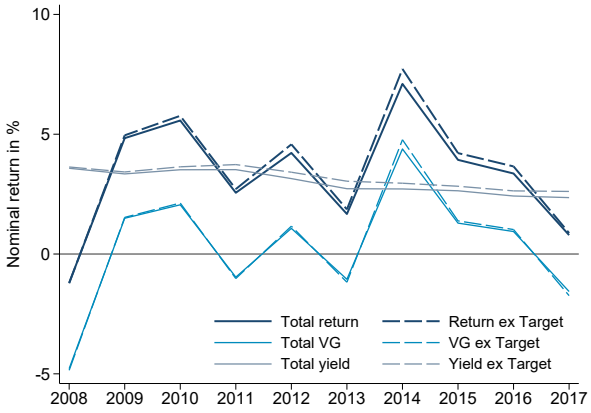

Notes: Panel (a) shows the increasing share of Target2 balances in total foreign assets excluding financial derivatives, i.e. the asset measure used for return computation. Panel (b) shows that the effect of Target2 balances for aggregate returns is not very large.

D7 Impact of Target2 balances on German foreign returns

Germany's Target2 balances have been growing fast since 2007. The crisis in the euro area led to large inflows of deposits into Germany and, thus, to higher Target2 claims. In recent years, the asset purchases by the European Central Bank (ECB) are also believed to have contributed to the growth of Target2 balances. Many of the purchased securities are sold by banks outside of the euro area, which tend to have their Target2 accounts registered with the German Bundesbank. If these investors sell their securities to the ECB, Target2 assets in Germany increase (see e.g., Bundesbank 2017).

Target2 balances are remunerated at a benchmark interest rate (the "main refinancing rate") that currently stands at zero. Target2 claims enter both German foreign assets as well as the financial account. Therefore, they do not lead to any valuation changes or yields. However, they do lower the returns (in absolute terms) by increasing the denominator.

Figure D.14 plots the share of Target2 balances in total assets, and total returns, yields and valuation changes including and excluding Target2 balances. We plot returns since 2008 because 2007 is the first year Target2 balances amounted to more than 1% of total assets. This share increased to 11.5% of total assets (excluding financial derivative) in 2017 with some fluctuations, see Panel (a) of Figure D.14.

Taken together, including Target2 lowers the yearly nominal return on total German foreign assets by 0.23 percentage points, on average, since 2008. The return effect was largest in 2014, with -0.62 percentage points. In the remaining years it varied between

Figure D.15: Detailed composition of FDI assets, 1999 to 2017

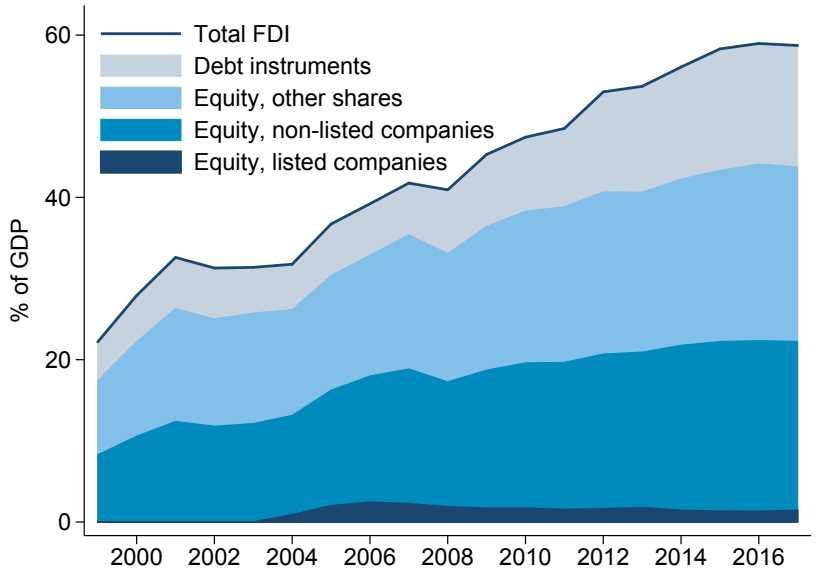

Notes: Composition of German FDI assets. Data from Bundesbank.

0.2 and -0.36 percentage points.[28]

D8 Book values for FDI assets

As discussed in Subsection 4.2.2, German FDI assets are recorded at book value before 2004, which could bias our results. This is mainly relevant for equity in listed companies because for equity in non-listed companies there is no direct estimate of market prices available anyways. Figure D.15 shows that equity in listed companies only makes up a small part of German FDI assets. Instead, the equity assets are roughly equally split between equity in non-listed companies and other equity, which includes real estate. Since market prices are less relevant for the non-listed shares and are applied to real estate assets, the valuation at book values should not affect our results too much.

Nevertheless, we now check our results for sensitivity to these valuation effects for FDI. For this purpose, we consider the counterfactual where there are only valuation changes due to exchange rates. Technically when assets are recorded at book value there can be no price adjustments, this means that if we record valuation changes other than those due to exchange rate movements these must be write-offs or due to mismeasurement. Now we can assume that all other changes observed are due to mismeasurement and set total valuation changes before 2005 to the estimated valuation

[28]The positive difference occurs in 2008 when the nominal return was negative because excluding Target2 increases the absolute size of the returns.

Table D.9: FDI returns with adjusted valuation changes

	Standard return				Adjusted return			
	(1) Baseline	(2) Val. FX	(3) Risk	(4) Geo.	(5) Baseline	(6) Val. FX	(7) Risk	(8) Geo.
Germany dummy	-3.34*** (1.25)	-3.61*** (1.27)	-3.23** (1.30)	-2.90** (1.36)	-2.44** (1.02)	-2.71*** (1.00)	-2.24** (1.06)	-1.91* (1.12)
Valuation ch. due to ex. rates, FDI		0.38*** (0.10)	0.36*** (0.11)	0.37*** (0.11)		0.39*** (0.10)	0.37*** (0.11)	0.38*** (0.11)
3-year rolling std. dev., FDI			0.15 (0.20)	0.13 (0.21)				
3-year rolling std. dev. adj. return							0.14 (0.19)	0.12 (0.20)
Advanced Europe				-0.12 (0.15)				-0.13 (0.15)
Advanced Non-Europe				-0.10 (0.18)				-0.12 (0.18)
Emerging &Developing				-0.10 (0.15)				-0.12 (0.14)
Constant	9.83 (5.98)	11.69** (5.18)	10.56** (5.09)	21.33 (17.06)	9.97* (5.93)	11.88** (5.12)	10.93** (4.89)	23.44 (16.71)
Observations	339	339	328	328	339	339	328	328
Adjusted R^2	0.14	0.18	0.19	0.18	0.15	0.19	0.20	0.19
No. countries	12	12	12	12	12	12	12	12
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Columns (1) to (4) reproduce the results of Table 4.8 in Section 4.5. These are compared to the results using adjusted returns in columns (5) to (8) (excluding any valuation changes not due to exchange rates). Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Net foreign assets and financial account balance included in the regressions, but not shown.

changes due to exchange rates. To get a conservative estimate we do this only for Germany.

Table D.9 shows the regression using the returns on FDI as dependent variable from Section 4.5. We compare results with the standard return measure to those with the adjusted return, measured as described above. While the difference between Germany and the other countries becomes smaller in absolute terms with the adjusted returns, it is still negative and significant. Thus, our results hold up even if we exclude all valuation changes except for those due to exchange rate movements. It should be noted that this analysis only adjusts the German returns. If other countries also report book values, their returns are likely to be underestimated as well.

Bibliography

Abberger, Klaus and Klaus Wohlrabe (2006). "Einige Prognoseigenschaften des ifo Geschäftsklimas – Ein Überblick über die neuere wissenschaftliche Literatur". *ifo Schnelldienst 59*, 19–26.

Altavilla, Carlo, Luca Brugnolini, Refet S. Gürkaynak, Roberto Motto, and Guiseppe Ragusa (2019). "Measuring euro area monetary policy". *Journal of Monetary Economics 108*, 162–179.

Altavilla, Carlo and Domenico Giannone (2017). "The Effectiveness of Non-Standard Monetary Policy Measures: Evidence from Survey Data". *Journal of Applied Econometrics 32* (5), 952–964.

Angeletos, George-Marios and Jennifer La'O (2013). "Sentiments". *Econometrica 81* (2), 739–779.

Bachmann, R., Steffen Elstner, and Eric R. Sims (2013). "Uncertainty and Economic Activity: Evidence from Business Survey Data". *American Economic Journal: Macroeconomics 5* (2), 217–249.

Bachmann, Rüdiger (2019). "Comments on 'Monetary policy announcements and expectations: Evidence from German firms'". *Journal of Monetary Economics 108*, 64–68.

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.

Bachmann, Rüdiger, Benjamin Born, Steffen Elstner, and Christian Grimme (2019). "Time-varying business volatility and the price setting of firms". *Journal of Monetary Economics 101*, 82–99.

Bachmann, Rüdiger and Steffen Elstner (2015). "Firm optimism and pessimism". *European Economic Review 79*, 297–325.

Bachmann, Ruediger and Peter Zorn (2018). "What Drives Aggregate Investment? Evidence from German Survey Data". CEPR Discussion Paper No. 12710.

Baldi, Guido and Börn Bremer (2015). "The Evolution of Germany's Net Foreign Asset Position". *DIW Economic Bulletin 5* (22/23), 303–309.

Balleer, Almut and Peter Zorn (2019). "Monetary Policy, Price Setting, and Credit Constraints". mimeo.

Barsky, Robert B. and Eric R. Sims (2012). "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence". *American Economic Review 52* (4), 1343–1377.

Bauer, Michael D. and Glenn D. Rudebusch (2014). "The Signaling Channel for Federal Reserve Bond Purchases". *International Journal of Central Banking 10* (3), 233–289.

Beaudry, Paul and Franck Portier (2006). "Stock Prices, News, and Economic Fluctuations". *American Economic Review 96* (4), 1293–1307.

Becker, Sasha O. and Klaus Wohlrabe (2008). "Micro Data at the Ifo Institute for Economic Research - The 'Ifo Business Survey' Usage and Access". *Schmollers Jahrbuch-Zeitschrift für Wirtschafts- und Sozialwissenschaften 182* (2), 307–319.

Bénétrix, Agustin S., Philip R. Lane, and Jay C. Shambaugh (2015). "International Currency Exposures, Valuation Effects and the Global Financial Crisis". *Journal of International Economics 96*, S98–S109.

Bernanke, Ben S. (2015). "Germany's Trade Surplus is a Problem". Brookings online, 03.04.2015.

Bholat, David, Nida Broughton, Janna Ter Meer, and Eryk Walczak (2019). "Enhancing central bank communications using simple and relatable information". *Journal of Monetary Economics 108*, 1–15.

Binder, Carola C. (2017a). "Consumer forecast revisions: Is information really so sticky?" *Economics Letters 161*, 112–115.

——— (2017b). "Fed speak on main street: Central bank communication and household expectations". *Journal of Macroeconomics 52*, 238–251.

——— (2017c). "Measuring uncertainty based on rounding: New method and application to inflation expectations". *Journal of Monetary Economics 90*, 1–12.

Board of Academic Advisors (2019). "Wirtschaftspolitische Probleme der deutschen Leistungsbilanz". Report for the German Ministry of Economic Affairs and Energy.

Boneva, Lena, James Cloyne, Martin Weale, and Tomasz Wieladek (2020). "Firms' Expectations of New Orders, Employment, Costs and Prices: Evidence from Micro Data". *Economic Journal*, forthcoming.

Breeden, Douglas T., Michael R. Gibbons, and Robert H. Litzenberger (1989). "Empirical Tests of the Consumption-Oriented CAPM". *The Journal of Finance 44* (2), 231–262.

Buchheim, Lukas and Sebastian Link (2017). "The Effect of Disaggregate Information on the Expectation Formation of Firms". CESifo Working Paper No. 6768.

Bundesbank (1998). *50 Jahre Deutsche Mark. Monetäre Statistiken von 1948 bis 1997 auf CD-ROM*. Verlag Vahlen und C.H. Beck.

Bundesbank (2008). "Das deutsche Auslandsvermögen seit Beginn der Währungsunion: Entwicklung und Struktur". *Monatsbericht Oktober*, 15–33.

——— (2012). "Zahlungsbilanzstatistik September 2012". *Statistisches Beiheft 3 zum Monatsbericht September*, 98–103.

——— (2014). "Diskrepanz zwischen der Veränderung des Auslandsvermögens und des kumulierten Saldos der Kapitalbilanz: kein geeigneter Indikator für Vermögensverluste". *Monatsbericht Mai*, 52–54.

——— (2017). "Die deutsche Zahlungsbilanz für das Jahr 2016". *Monatsbericht März*, 17–34.

——— (2018). "Die deutsche Auslandsposition: Höhe, Rentabilität und Risiken der grenzüberschreitenden Vermögenswerte". *Monatsbericht Dezember*, 47–68.

Caliendo, Marco and Sabine Kopeinig (2008). "Some practical guidance for the implementation of propensity score matching". *Journal of Economic Surveys 22* (1), 31–72.

Campbell, Jeffrey, Charles L. Evans, Jonas Fisher, and Alejandro Justiniano (2012). "Macroeconomic Effects of Federal Reserve Forward Guidance". *Brookings Papers on Economic Activity 2012* (SPRING), 1–80.

Campbell, John Y. (1999). "Asset Prices, Consumption, and the Business Cycle". *Handbook of Macroeconomics*. Ed. by John B. Taylor and Michael Woodford. Vol. 1C. North Holland: Amsterdam. Chap. 19, 1231–1303.

Carlstrom, Charles T., Timothy S. Fuerst, and Matthias Paustian (2015). "Inflation and output in New Keynesian models with a transient interest rate peg". *Journal of Monetary Economics 76*, 230–243.

Carrillo, Paul E. and Shahe M. Emran (2012). "Public information and inflation expectations: Microeconometric evidence from a natural experiment". *Review of Economics and Statistics 94* (4), 860–877.

Carroll, Christopher D. (2003). "Macroeconomic Expectations of Households and Professional Forecasters". *The Quarterly Journal of Economics 118* (1), 269–298.

Claus, Edda and Viet Hoang Nguyen (2020). "Monetary policy shocks from the consumer perspective". *Journal of Monetary Economics*, forthcoming.

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". NBER Working Paper No. 26106.

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), 2644–2678.

Coibion, Olivier and Yuriy Gorodnichenko (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 Saten Kumar (2018). "How Do Firms Form Their Expectations? New Survey Evidence". *American Economic Review 108* (9), 2671–2713.

Coibion, Olivier, Yuriy Gorodnichenko, and Tiziano Ropele (2020). "Inflation Expectations and Firm Decisions: New Causal Evidence". *The Quarterly Journal of Economics 135* (1), 165–219.

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

Conrad, Christian and Michael J. Lamla (2010). "The High-Frequency Response of the EUR-USD Exchange Rate to ECB Communication". *Journal of Money, Credit and Banking 42* (7), 1391–1417.

Crump, Richard K., Stefano Eusepi, Andrea Tambalotti, and Giorgio Topa (2015). "Subjective intertemporal substitution". FRB of New York Staff Report No. 734.

Curcuru, Stephanie E., Tomas Dvorak, and Francis E. Warnock (2008). "Cross-border Return Differentials". *The Quarterly Journal of Economics 123* (4), 1495–1530.

——— (2013). "On Returns Differentials". *Journal of International Money and Finance 36*, 1–25.

Curtin, Richard and Pierrick Dechaux (2015). "University of Michigan's Survey of Consumers: Measuring and Interpreting Economic Expectations". 7th Joint EC/OECD Workshop Business and Consumer Surveys (available at `https://data.sca.isr.umich.edu/fetchdoc.php?docid=57828`).

Curtin, Richard T. (1982). "Indicators of consumer behavior: The University of Michigan surveys of consumers". *Public Opinion Quarterly 46* (3), 340–352.

D'Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber (2019a). "Cognitive abilities and inflation expectations". *AEA Papers and Proceedings 109*, 562–66.

D'Acunto, Francesco, Daniel Hoang, and Michael Weber (2018). "Unconventional Fiscal Policy". *AEA Papers and Proceedings 108*, 519–523.

D'Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber (2019b). "Exposure to Daily Price Changes and Inflation Expectations". NBER Working Paper No. 26237.

D'Acunto, Francesco, Ulrike Malmendier, and Michael Weber (2020). "Gender Roles and the Gender Expectations Gap". NBER Working Paper No. 26837.

Das, Sreyoshi, Camelia M. Kuhnen, and Stefan Nagel (2020). "Socioeconomic Status and Macroeconomic Expectations". *The Review of Financial Studies 33* (1), 395–432.

Dedola, Luca, Georgios Georgiadis, Johannes Gräb, and Arnaud Mehl (2018). "Does a big bazooka matter? Central bank balance-sheet policies and exchange rates". ECB Working Paper No. 2197.

Dehejia, Rajeev H. and Sadek Wahba (1999). "Causal Effects in Nonexperimental Studies: Reevaluating the Evaluation of Training Programs". *Journal of the American Statistical Association 94* (448), 1053–1062.

——— (2002). "Propensity Score-Matching Methods For Nonexperimental Causal Studies". *The Review of Economics and Statistics 84* (1), 151–161.

Del Negro, Marco, Marc Giannoni, and Christina Patterson (2012). "The Forward Guidance Puzzle". Federal Reserve Bank of New York Staff Reports No. 574.

den Haan, Wouter (2013). "Introduction". *Forward Guidance: Perspectives from Central Bankers, Scholars and Market Participants*. Ed. by Wouter den Haan. Centre for Economic Policy Research, 1–21.

den Haan, Wouter, Martin Ellison, Ethan Ilzetzki, Michael McMahon, and Ricardo Reis (2016). "The Danger of Germany's Current Account Surpluses: Results of the CFM and CEPR Survey". VoxEU, 27.10.2016.

Deutsche Bank (2013). "Deutsche Leistungsbilanzüberschüsse – Kritik nicht stichhaltig". Deutsche Bank Research – Standpunkt.

Dräger, Lena and Michael J. Lamla (2018). "Is the Anchoring of Consumers' Inflation Expectations Shaped by Inflation Experience?" CESifo Working Paper No. 7042.

Dräger, Lena and Giang Nghiem (2020). "Are Consumers' Spending Decisions in Line With an Euler Equation?" *The Review of Economics and Statistics*, forthcoming.

EBDC-BEP (2017). "Business Expectations Panel 01/1980 – 06/2017, LMU-ifo Economics & Business Data Center, Munich". *doi: 10.7805/ebdc-bep-2017*.

Eggertsson, Gauti B. and Michael Woodford (2003). "The Zero Bound on Interest Rates and Optimal Monetary Policy". *Brookings Papers on Economic Activity 2003* (1), 139–211.

Ehrmann, Michael, Damjan Pfajfar, and Emiliano Santoro (2017). "Consumers' Attitudes and Their Inflation Expectations". *International Journal of Central Banking 13* (1), 225–259.

Enders, Zeno, Franziska Hünnekes, and Gernot J. Müller (2019a). "Firm expectations and economic activity". CESifo Working Paper No. 7623.

——— (2019b). "Monetary policy announcements and expectations: Evidence from German firms". *Journal of Monetary Economics 108*, 45–63.

Enders, Zeno, Michael Kleemann, and Gernot J. Müller (2017). "Growth Expectations, Undue Optimism, and Short-Run Fluctuations". Bundesbank Discussion Paper No. 11/2017.

European Commission (2016). *Country Report Germany 2016*.

Fiedler, Salomon, Holger Görg, Cecília Hornok, Nils Jannsen, Stefan Kooths, Léa Marchal, and Galina Potjagailo (2018). "Zur wirtschaftlichen Bedeutung deutscher Direktinvestitionen im Ausland und deren Relevanz für den deutschen Leistungsbilanzsaldo". Report for the German Ministry of Finance.

Frey, Rainer, Ulrich Grosch, and Alexander Lipponer (2014). "Fallstricke bei der Bestimmung von Vermögensverlusten deutscher Anleger im Ausland". *Wirtschaftdienst 94* (11), 806–812.

Fuest, Clemens (2017). "The Politics of Germany's External Surplus". Project Syndicate, 17.07.2017.

Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer (2015). "Expectations and Investment". *NBER Macroeconomics Annual 2015, Volume 30*. NBER Chapters. National Bureau of Economic Research, 379–431.

German Council of Economic Experts (2017). *Economic Forecast 2017 and 2018*.

German Ministry of Finance (2017). "The German Current Account in the Context of US-German Trade". Position Paper with the Ministry of Economic Affairs, 10.05.2017.

Gorodnichenko, Yuriy and Byoungchan Lee (2017). "A Note on Variance Decomposition with Local Projections". NBER Working Paper No. 23998.

Gourinchas, Pierre-Olivier and Hélène Rey (2007a). "From World Banker to World Venture Capitalist: U.S. External Adjustment and the Exorbitant Privilege". *G7 Current Account Imbalances: Sustainability and Adjustment*. Ed. by Richard H. Clarida. University of Chicago Press. Chap. 1, 11–66.

——— (2007b). "International Financial Adjustment". *Journal of Political Economy 115* (4), 665–703.

——— (2014). "External Adjustment, Global Imbalances, Valuation Effects". *Handbook of International Economics, Vol. 4*. Ed. by Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff. Elsevier. Chap. 10, 585–645.

Graf, Christina and Christian Grimme (2017). "Der Einfluss der Vermögenseinkommen auf den deutschen Leistungsbilanzüberschuss". *ifo Schnelldienst 22*, 38–44.

Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson (2005). "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements". *International Journal of Central Banking 1* (1), 55–93.

Habib, Maurizio M. (2010). "Excess Returns on Net Foreign Assets – The Exorbitant Privilege from a Global Perspective". ECB Working Paper No. 1158.

Hall, Robert (2011). "The long slump". *American Economic Review 101* (2), 431–469.

Heckman, James, Hidehiko Ichimura, Jeffrey Smith, and Petra Todd (1998). "Characterizing Selection Bias Using Experimental Data". *Econometrica 66* (5), 1017–1098.

Henzel, Steffen R. and Sebastian Rast (2013). "Prognoseeigenschaften von Indikatoren zur Vorhersage des Bruttoinlandsprodukts in Deutschland". *ifo Schnelldienst 66*, 39–46.

Huber, Martin, Michael Lechner, and Andreas Steinmayr (2015). "Radius matching on the propensity score with bias adjustment: tuning parameters and finite sample behaviour". *Empirical Economics 49*, 1–31.

IBS-IND (2018a). "Ifo Business Survey Industry 1/1980 – 6/2018, LMU-ifo Economics & Business Data Center, Munich". *doi: 10.7805/ebdc-ibs-ind-2018a*.

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

Imbens, Guido W. and Donald B. Rubin (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.

IMF (2016). *Germany: IMF Country Report No. 16/202*.

Jarociński, Marek and Peter Karadi (2020). "Deconstructing Monetary Policy Surprises – The Role of Information Shocks". *American Economic Journal: Macroeconomics*, forthcoming.

Jia, Chengcheng (2019). "The Informational Effect of Monetary Policy and the Case for Policy Commitment". Federal Reserve Bank of Cleveland, Working Paper No. 19-07.

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

Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M. Taylor (2019). "The Rate of Return on Everything, 1870-2015". *The Quarterly Journal of Economics 134* (3), 1225–1298.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor (2017). "Macrofinancial History and the New Business Cycle Facts". *NBER Macroeconomics Annual 2016*. Ed. by Martin Eichenbaum and Jonathan A. Parker. Vol. 31. University of Chicago Press, 213–263.

Kamdar, Rupal (2019). "The Inattentive Consumer: Sentiment and Expectations". mimeo.

Keynes, John M. (1936). *The General Theory of Employment, Interest, and Money*. Orlando, Florida: Edition Harvest/Harcourt Brace 1964.

Krishnamurthy, Arvind and Annette Vissing-Jorgensen (2011). "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy". *Brookings Papers on Economic Activity 2011* (FALL), 215–287.

Krugman, Paul R. (1998). "It's baaack: Japan's slump and the return of the liquidity trap". *Brookings Papers on Economic Activity 1998* (2), 137–205.

——— (2017). "Germany's Real Sin". New York Times Opinion Pages, 27.05.2017.

Kumar, Saten, Hassan Afrouzi, Olivier Coibion, and Yuriy Gorodnichenko (2015). "Inflation Targeting Does Not Anchor Inflation Expectations: Evidence from Firms in New Zealand". *Brookings Papers on Economic Activity 2015* (FALL), 151–208.

Kydland, Finn E. and Edward C. Prescott (1982). "Time to Build and Aggregate Fluctuations". *Econometrica 50* (6), 1345–1370.

Lamla, Michael J. and Sarah M. Lein (2014). "The role of media for consumers' inflation expectation formation". *Journal of Economic Behavior & Organization 106*, 62–77.

Lamla, Michael J. and Thomas Maag (2012). "The Role of Media for Inflation Forecast Disagreement of Households and Professional Forecasters". *Journal of Money, Credit and Banking 44* (7), 1325–1350.

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

Lane, Philip R. (2001). "Do International Investment Income Flows Smooth Income?" *Review of World Economics 137* (4), 714–736.

Lane, Philip R. and Gian Maria Milesi-Ferretti (2003). "International Financial Integration". *IMF Third Annual Research Conference*. Vol. 50. IMF Staff Papers. International Monetary Fund, 82–113.

——— (2007a). "A Global Perspective on External Positions". *G7 Current Account Imbalances: Sustainability and Adjustment*. Ed. by Richard H. Clarida. University of Chicago Press. Chap. 2, 67–102.

——— (2007b). "The External Wealth of Nations Mark II: Revised and Extended Estimates of Foreign Assets and Liabilities, 1970-2004". *Journal of International Economics 73* (2), 223–250.

——— (2009). "Where Did All the Borrowing Go? A Forensic Analysis of the U.S. External Position". *Journal of the Japanese and International Economies 23* (2), 177–199.

Lane, Philip R. and Jay C. Shambaugh (2010). "Financial Exchange Rates and International Currency Exposures". *American Economic Review 100* (1), 518–540.

Lechner, Michael (2001). "Identification and estimation of causal effects of multiple treatments under the conditional independence assumption." *Econometric Evaluation of Labour Market Policies*. Ed. by M. Lechner and F. Pfeiffer. Heidelberg: Physica, 1–18.

Lechner, Michael, Ruth Miquel, and Conny Wunsch (2011). "Long-Run Effects of Public Sector Sponsored Training". *Journal of the European Economic Association 9* (4), 742–784.

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

Lewis, Michael (2010). *The Big Short*. W. W. Norton & Company.

Lorenzoni, Guido (2009). "A Theory of Demand Shocks". *American Economic Review 99* (5), 2050–2084.

Lucas, Robert E. Jr. (1972). "Expectations and the Neutrality of Money". *Journal of Economic Theory 4* (2), 103–124.

——— (1973). "Some International Evidence on Output-Inflation Tradeoffs". *American Economic Review 63* (3), 326–334.

Maćkowiak, Bartosz and Mirko Wiederholt (2009). "Optimal Sticky Prices under Rational Inattention". *American Economic Review 99* (3), 769–803.

——— (2015). "Business Cycle Dynamics under Rational Inattention". *The Review of Economic Studies 82* (4), 1502–1532.

Malmendier, Ulrike and Stefan Nagel (2016). "Learning from inflation experiences". *The Quarterly Journal of Economics 131* (1), 53–87.

Mankiw, N. Gregory and Ricardo Reis (2002). "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve". *The Quarterly Journal of Economics 117* (4), 1295–1328.

Massenot, Baptiste and Yuri Pettinicchi (2018). "Can firms see into the future? Survey evidence from Germany". *Journal of Economic Behavior & Organization 145*, 66–79.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson (2016). "The Power of Forward Guidance Revisited". *American Economic Review 106* (10), 3133–3158.

Meissner, Christopher M. and Alan M. Taylor (2008). "Losing our Marbles in the New Century? The Great Rebalancing in Historical Perspective". *Global Imbalances and the Evolving World Economy*. Ed. by Jane Sneddon Little. Vol. 51. Conference Series. Federal Reserve Bank of Boston, 53–129.

Melosi, Leonardo (2017). "Signalling Effects of Monetary Policy". *The Review of Economic Studies 84* (2), 853–884.

Moessner, Richhild, David-Jan Jansen, and Jakob de Haan (2017). "Communication About Future Policy Rates in Theory and Practice: A Survey". *Journal of Economic Surveys 31* (3), 678–711.

Mortensen, Dale and Christopher Pissarides (2009). "Job creation and job destruction in the theory of unemployment". *The Review of Economic Studies 63* (2), 397–415.

Muth, John F. (1961). "Rational Expectations and the Theory of Price Movements". *Econometrica 29* (3), 315–335.

Nakamura, Emi and Jón Steinsson (2018). "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect". *The Quarterly Journal of Economics 133* (3), 1283–1330.

Nerlove, Marc (1983). "Expectations, Plans, and Realizations in Theory and Practice". *Econometrica 51* (5), 1251–1279.

Nimark, Kristoffer P. (2014). "Man-Bites-Dog Business Cycles". *American Economic Review 104* (8), 2320–2367.

Obstfeld, Maurice and Kenneth Rogoff (1996). *Foundations of International Macroeconomics*. The MIT Press.

Pagan, Adrian (1984). "Econometric issues in the analysis of regressions with generated regressors". *International Economic Review 25* (1), 221.

Pfajfar, Damjan and Emiliano Santoro (2013). "News on Inflation and the Epidemiology of Inflation Expectations". *Journal of Money, Credit and Banking 45* (6), 1045–1067.

Picault, Matthieu and Thomas Renault (2017). "Words are not all created equal: A new measure of ECB communication". *Journal of International Money and Finance 79*, 136–156.

Pigou, Arthur C. (1927). *Industrial Fluctuations*. Macmillan.

Ramb, Fred and Alfons J. Weichenrieder (2005). "Taxes and the Financial Structure of German Inward FDI". *Review of World Economics 141* (4), 670–692.

Romer, Christina D. and David H. Romer (2000). "Federal Reserve Information and the Behavior of Interest Rates". *American Economic Review 90* (3), 429–457.

——— (2004). "A New Measure of Monetary Shocks: Derivation and Implications". *American Economic Review 94* (4), 1055–1084.

Rosenbaum, Paul R. and Donald B. Rubin (1983). "The central role of the propensity score in observational studies for causal effects". *Biometrika 70* (1), 41–55.

Rubin, Donald B. (2001). "Using Propensity Scores to Help Design Observational Studies: Application to the Tobacco Litigation". *Health Services & Outcomes Research Methodology 2*, 169–188.

Sauer, Stefan and Klaus Wohlrabe (2018). "The New ifo Business Climate Index for Germany". *CESifo Forum 19* (2), 59–64.

——— (2019). "CEO or Intern – Who Actually Answers the Questionnaires in the ifo Business Survey?" *CESifo Forum 20* (2), 29–31.

Schmitt-Grohé, Stephanie and Martín Uribe (2012). "What's news in business cycles?" *Econometrica 80* (6), 2733–2764.

Schuknecht, Ludger (2014). "The Empire Strikes Back: Why Germany's Exports and Current Account Surpluses Benefit Other Countries". *The International Economy Fall*, 50–72.

Shapiro, Adam Hale, Moritz Sudhof, and Daniel Wilson (2019). "Measuring News Sentiment". Federal Reserve Bank of San Francisco Working Paper No. 2017-01.

Sims, Christopher A. (1998). "Stickiness". *Carnegie-Rochester Conference Series on Public Policy, 49*, 317–356.

——— (2003). "Implications of Rational Inattention". *Journal of Monetary Economics 50* (3). Swiss National Bank/Study Center Gerzensee Conference on Monetary Policy under Incomplete Information, 665–690.

Sinn, Hans-Werner (2017). "How to Reduce Germany's Surplus". Project Syndicate, 21.07.2017.

Sørensen, Bent E. and Oved Yosha (1998). "International Risk Sharing and European Monetary Unification". *Journal of International Economics 45* (2), 211–238.

Stock, James H. and Mark W. Watson (2018). "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments". *The Economic Journal 128* (610), 917–948.

Swanson, Eric T. (2017). "Measuring the Effects of Federal Reserve Forward Guidance and Asset Purchases on Financial Markets". NBER Working Paper No. 23311.

Tanaka, Mari, Nicholas Bloom, Joel M. David, and Maiko Koga (2020). "Firm performance and macro forecast accuracy". *Journal of Monetary Economics*, forthcoming.

Taylor, Alan M. and Jeffrey G. Williamson (1994). "Capital Flows to the New World as an Intergenerational Transfer". *Journal of Political Economy 102* (2), 348–371.

Vellekoop, Nathanael and Mirko Wiederholt (2019). "Inflation expectations and choices of households". SAFE Working Paper No. 250.

Weidmann, Jens (2014). "External Imbalances in the Euro Area". Speech at the International Business Cycle Conference at the Kiel Institute for the World Economy, 17.03.2014.

——— (2019). "Herausforderungen für das deutsche Wirtschaftsmodell ". Speech at the Sparkassentag (Meeting of German savings banks), 16.05.2019.

Wiederholt, Mirko (2015). "Empirical Properties of Inflation Expectations and the Zero Lower Bound". mimeo.

Wieland, Johannes and Mu-Jeung Yang (2020). "Financial Dampening". *Journal of Money, Credit and Banking 52* (1), 79–113.

Woodford, Michael (2002). "Imperfect Common Knowledge and the Effects of Monetary Policy". *Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*. Ed. by Philippe Aghion, Roman Frydman, Joseph Stiglitz, and Michael Woodford. Princeton University Press.

——— (2003). *Interest and Prices*. Princeton University Press.