

EXPECTATIONS AND BEHAVIOR OF FIRMS IN MACROECONOMICS AND FINANCE

Manuel Menkhoff



Dissertation
Munich, 2025

Expectations and Behavior of Firms in Macroeconomics and Finance

Inaugural-Dissertation

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

2025

vorgelegt von
Manuel Menkhoff

Referent:

Prof. Dr. Andreas Peichl

Koreferent:

Prof. Dr. Mirko Wiederholt

Promotionsabschlussberatung:

16. Juli 2025

Datum der mündlichen Prüfung: 30.06.2025

Namen der Berichterstatter: Andreas Peichl, Mirko Wiederholt, Martin Schneider

Acknowledgements

I am deeply grateful to the many people who have supported me throughout my Ph.D. journey. First and foremost, I am especially grateful to my supervisor, Andreas Peichl, for his continuous support from day one. He not only guided me academically, but also helped me connect with researchers within and beyond the ifo network, which greatly enriched my experience.

I would also like to express my special thanks to my second and third supervisors, Mirko Wiederholt and Martin Schneider, as well as my fourth letter writer, Benjamin Born, for their valuable insights and guidance during the Ph.D. and especially throughout the job market process.

I also want to thank all my other co-authors—Rüdiger Bachmann, Almut Balleer, Lea Best, Kai Carstensen, Zeno Enders, Isabel Gödl-Hanisch, Sebastian Link, Gernot Müller, Knut Niemann, Paul Schüle, and Peter Zorn—from whom I learned a lot during my Ph.D. while working on several projects, some of which are not part of the thesis. I greatly enjoy working with you.

Many thanks go to Klaus Wohlrabe and his survey team at ifo, including Katrin Demmelhuber, Jonas Hennrich, Felix Leiss, Agnesa Nimanaj, and Daria Schaller, for their support and valuable discussions on the survey implementations. My sincere thanks also go to Sebastian Wichert and his team at the EBDC—Heike Mittelmeier, Valentin Reich, and Michael Rindler—for their tireless technical support, including countless exports and imports. I gratefully acknowledge financial support from the German Academic Exchange Service (DAAD) and the Joachim Herz Foundation through an Add-On Fellowship.

In addition, I would like to thank Julia Baarck, Maximilian Blömer, Mathias Dolls, Lilly Fischer, Friederike Fourné, David Gstrein, Justyna Klejdysz, Kevin Kloiber, Felicitas Koch, Simon Krause, Carla Krolage, Max Lay, Robert Lehmann, Sascha Möhrle, Marcel Schlepper, Tiphaine Wibault, Timo Wollmershäuser, Pascal Zamorski and Lara Zarges for advice, discussions, and for sharing many enjoyable moments along the way.

This thesis is dedicated to Sarah and to my parents, from whom I have always received love and unconditional support.

Contents

Preface	xiii
1 Downward Revision of Investment Decisions after Tax Hikes	1
1.1 Introduction	2
1.2 Institutional background and data	5
1.2.1 The German local business tax	6
1.2.2 Firm-level data on revisions of investment plans	8
1.3 Empirical strategy	11
1.4 Results	15
1.4.1 Revision of investment plans after tax hikes: Main results	15
1.4.2 Magnitude of effect size in comparison to the literature	21
1.4.3 State dependence and heterogeneity	27
1.5 Conclusion	32
2 Firms' Pass-Through Dynamics:	
A Survey Approach	35
2.1 Introduction	36
2.2 Empirical results	40
2.2.1 Data	40
2.2.2 Five facts about pass-through dynamics	42
2.3 Survey experiments	51
2.3.1 Experimental design	51
2.3.2 Results	53
2.4 General equilibrium price-setting model	58
2.4.1 Model setup	58
2.4.2 Calibration	60
2.4.3 Results	61
2.5 Conclusion	64
3 Firm Expectations and News:	
Micro v Macro	67
3.1 Introduction	68

3.2	Framework	72
3.2.1	Baseline specification	72
3.2.2	Microfoundations	74
3.3	Price expectations	78
3.3.1	Survey of Inflation and Growth Expectations	78
3.3.2	Results for price expectations	81
3.4	Production expectations	85
3.4.1	ifo Business Climate Survey	85
3.4.2	Results for production expectations	88
3.4.3	Firm-level variation	91
3.4.4	Real effects	94
3.5	Aggregate implications	96
3.6	Conclusion	98
4	The Devil is in the Tail:	
	Macroeconomic Tail Risk Expectations of Firms	101
4.1	Introduction	102
4.2	Nature of macroeconomic tail risk expectations	109
4.2.1	ifo Business Survey	109
4.2.2	Firms' macroeconomic tail risk expectations	112
4.2.3	Firms' expected exposure to macro tail events	116
4.3	Macroeconomic tail risk expectations and firm decisions	120
4.3.1	Stylized model framework	120
4.3.2	Micro approach	124
4.3.3	Experimental approach	130
4.3.4	Macro approach	134
4.4	Quantitative heterogeneous firm model	141
4.4.1	Model setup	141
4.4.2	Model results	143
4.4.3	Policy implications	145
4.5	Conclusion	148
A	Appendix to Chapter 1	149
A.1	Business taxation in Germany	149
A.2	Data appendix	152
A.2.1	Administrative data at the municipality level	152
A.2.2	ifo Investment Survey	157
A.2.3	Construction and descriptive statistics of the merged dataset	161
A.3	Supplementary figures and tables	168

A.4	Back-of-the-envelope calculation	182
A.5	Calculation of effective tax rates	184
B	Appendix to Chapter 2	191
B.1	ifo Business Survey	191
B.1.1	Survey timing and inflation dynamics	191
B.1.2	Supplementary questions (translated)	191
B.1.3	Supplementary questions (original)	193
B.1.4	Additional results	196
B.2	Pass-through estimates in the literature	203
B.3	ifo Management-Survey	204
B.3.1	Original questionnaire	204
B.3.2	English translation	210
B.3.3	Additional results	214
B.3.4	Quantifying the pass-through dynamics and implications for the Phillips curve	220
B.4	Additional model details and results	222
B.4.1	Calibration	222
B.4.2	Simulated method of moments	222
B.4.3	Shape of the demand and profit functions	223
B.4.4	Results of estimation	224
C	Appendix to Chapter 3	227
C.1	Stylized model – details and proofs	227
C.1.1	Households	227
C.1.2	Solution	227
C.1.3	Proof of Proposition 1	229
C.1.4	Price expectations	232
C.2	Survey on Inflation and Growth Expectations	233
C.3	ifo Business Climate Survey	240
D	Appendix to Chapter 4	251
D.1	ifo Business Survey: Survey questions (translated)	251
D.2	Additional figures	253
D.3	Additional tables	264
D.4	Drivers of level differences between firms’ macroeconomic tail risk expec- tations	272
D.5	Firms’ exposure beliefs: Evidence from an open-ended survey question . . .	273
D.6	Model appendix	282

D.7 Additional quantitative model information	287
D.8 Back-of-the-envelope calculation to map the empirical estimates to the model	290
D.9 Additional LLM information	291
Bibliography	296

List of Figures

1.1	Variation in local business tax rates (1980-2018)	7
1.2	Relationship between planned and realized investment	11
1.3	Event study: Investment revision effect after a tax hike	19
1.4	Effect of tax hike on investment plans, realizations, and revisions	22
1.5	Investment revisions after a tax hike: State dependence	28
1.6	Testing for further heterogeneity	31
2.1	Pass-through dynamics over time	42
2.2	Cross-sectional variation in pass-through within and between industries	45
2.3	Expected duration of the shock and pass-through in April 2022	46
2.4	Event study: Expected duration of the shock and probability to increase prices	48
2.5	Pass-through until October 2022 and uncertainty	49
2.6	Intensive margin of pass-through until October 2022 and nominal rigidities	50
2.7	Pass-through of cost shocks: Main results	54
2.8	Pass-through of aggregate vs. idiosyncratic cost shocks after six months	56
2.9	Data vs model: Pass-through of a permanent aggregate shock (targeted)	62
2.10	Data vs model: Pass-through of a transitory aggregate shock (untargeted)	63
3.1	Basic statistics for SIGE	81
3.2	Response of price expectations to current and past news	85
3.3	Basic statistics for the ifo survey	88
3.4	Response to news over time	91
3.5	Firm-by-firm estimates	92
3.6	Relation between macro and micro coefficients at the firm level	93
3.7	How island illusion alters macro dynamics: Model simulation	98
4.1	Survey responses on investment and change in tangible fixed assets	111
4.2	Distribution of firms' macro tail risk expectations over time	113
4.3	Expected firm exposure scenarios in case of a 5% decline in GDP	118
4.4	Optimal capital stock and macro tail risk expectations	123
4.5	Timing of questions in the ifo Business Survey	125

4.6	Firms' macro tail risk expectations and investment plans	127
4.7	Heterogeneity across firm characteristics	129
4.8	Firms' investment choice in the hypothetical macro outlooks	133
4.9	Macroeconomic beliefs over time based on newspaper data	137
4.10	Investment dynamics after an increase in macro tail risk expectations (<i>MacroTail</i>)	139
4.11	Impact of macro tail risk expectations (<i>MacroTail</i>) on investment: Split by change in dispersion	140
4.12	Baseline simulation of a macro tail risk news shock	145
4.13	Macro tail risk news shocks and rule-based policy counterfactuals: Dy- namics of the aggregate capital stock	146
A.1	Timing of tax hike news	150
A.2	Time series of local scaling factors by municipality	154
A.3	Share of municipalities increasing the LBT over time	155
A.4	Number of tax hikes and distribution of tax changes	155
A.5	Predictability of tax hikes as a function of past tax hikes in the same municipality	156
A.6	Distribution of firms by number of employees	164
A.7	Time series of investment plans and realizations	164
A.8	Time series of investment revisions	165
A.9	Time series of share of large revenue drops	165
A.10	Distribution of the log revision ratio	167
A.11	Long event study: Effect of tax hike on investment plans, realizations, and revisions	168
A.12	Investment revision effect after a tax hike: Alternative estimators	169
A.13	Event study: Expenditures and revenues of municipalities	170
A.14	Investment revisions after a tax hike: Permutation test	171
A.15	Collectively bargained wage growth in manufacturing	171
A.16	Obstacles to investment by firm size	172
A.17	Time series of average interest rate on loans for firms	186
A.18	Present discounted value of depreciation: 7% vs time-varying interest rate	187
A.19	Relation of changes in LBT rate and changes in effective tax rates	188
A.20	Relation of changes in effective tax rates and changes in user cost of capital	189
B.1	Producer and consumer prices over time	191
B.2	Binned scatterplots of pass-through and expected effects of gas supply shortages	196
B.3	Limiting factors for pass-through	196

B.4	Heterogeneous pass-through across firms (2022/2023)	197
B.5	Heterogeneous pass-through across firms (2021/2022)	197
B.6	Gradual increase of pass-through over time (sector split)	198
B.7	Cross-sectional variation in pass-through within and between sectors: Drop short time spent	198
B.8	Cross-sectional histogram time spent	199
B.9	Time-series of average firm-level uncertainty	199
B.10	Box plots of pass-through estimates in the literature	203
B.11	Threshold of cost increases for pass-through	214
B.12	Heterogeneity of pass-through across sectors and firm size	215
B.13	Aggregate vs idiosyncratic pass-through across sectors and firm size	216
B.14	Relative demand and profit functions varying extent of micro real rigidities	223
B.15	Values of the weighted function	224
B.16	Effect of nominal demand shock on output	225
C.1	Representativeness of survey and sample	235
C.2	Firm-level regressions	239
C.3	Average forecast revisions and production growth	241
C.4	Distribution of firm-level responses to news	242
C.5	Response to concurrent and lagged news	243
D.1	Firms' macroeconomic tail risk expectations and mean GDP growth ex- pectations	253
D.2	Relation of subjective probability of macro tail event occurrence and firm exposure beliefs	254
D.3	Relation of macro tail risk to unemployment rate expectations	255
D.4	Relation of macro tail risk to unemployment rate uncertainty	256
D.5	Firms' macro tail risk expectations and investment plans: All four macro tail risk scenarios	257
D.6	Firms' macro tail risk expectations and investment plans; only within firm variation - split by average exposure	258
D.7	Business expectations - only within firm variation	258
D.8	Robustness check: Business expectations and uncertainty for different horizons	259
D.9	Robustness check: First and second moment macro controls (unemploy- ment rate)	259
D.10	Vignette: Baseline outlooks	260
D.11	Vignette: Shifted outlooks	261

D.12	Local projections: Impact on production of investment goods (other macro moments)	262
D.13	Local projections: Impact on production of investment goods (robustness)	263
D.14	Drivers of level differences between firms' macroeconomic tail risk expectations	273
D.15	Effect of different degrees of risk aversion on optimal capital stock	284
D.16	Effect of bankruptcy risk in the stylized model	286
D.17	Macro tail risk news shocks: Dynamics of the aggregate capital stock – sensitivity analysis	289

List of Tables

1.1	Tax hikes across municipalities and firms: Summary statistics	9
1.2	Difference-in-differences: Investment revisions after a tax hike	16
1.3	Investment revisions after a tax hike: Effective tax rates	25
1.4	Summary of estimated (semi-)elasticities	26
2.1	Pass-through over time across industries	44
2.2	Estimation results for parameters	61
3.1	Survey of Inflation and Growth Expectations	79
3.2	Response of price expectations to news	82
3.3	ifo Business Climate Survey	87
3.4	Response of production expectations to news	89
3.5	Accounting for firm-level heterogeneity	94
3.6	Over- and underreaction to news and real activity	95
3.7	Overreaction and underreaction to news: Monte Carlo evidence	97
4.1	Relation of macro tail risk expectations and business outlooks	115
4.2	Expected worst-case exposure and realized exposure to the Great Recession and the COVID-19 Pandemic	119
4.3	Investment choice and firm exposure to a macro tail event	133
A.1	Distribution of firms in the IVS by industry and size	158
A.2	Summary statistics of firms in the sample	163
A.3	Balance statistics of firms in the treatment and control group	163
A.4	Information content of investment plans for realized investment	167
A.5	Robustness: Baseline estimates excl. reunification period	173
A.6	Treatment effect heterogeneity: State dependence	174
A.7	Difference-in-differences: Investment revisions after a tax hike	175
A.8	Treatment effect heterogeneity: Volatility of revenue growth	176
A.9	Treatment effect heterogeneity: Current revenue growth I	177
A.10	Treatment effect heterogeneity: Current revenue growth II	178
A.11	Treatment effect heterogeneity: Financial constraints	179
A.12	Treatment effect heterogeneity: Firm size and settlement structure	180

A.13	Treatment effect heterogeneity: Tax hike dynamics	181
B.1	Pass-through until 10/2022 and return on sales in 2019 and 2022	200
B.2	Pass-through until 10/2022 and uncertainty at industry level	201
B.3	Pass-through until 10/2022 and uncertainty at firm level	202
B.4	Pass-through estimates across studies	203
B.5	Factors limiting firm’s price setting (open-ended question)	217
B.6	Expected duration of the shock and pass-through in April 2022	218
B.7	Change in expected duration of the shock and pass-through until April 2023	219
B.8	Implied coefficients of Phillips curve slope $\mathbf{K}$	220
B.9	Estimates of $\mathbf{K}$ in the literature	221
B.10	Calibration for the menu-cost model with micro real rigidities	222
C.1	Relevant questions from SIGE	233
C.2	Macro news and forecast revisions	234
C.3	Response of price expectations to news: Robustness	236
C.4	Additional regression results from the SIGE	237
C.5	Relevant questions from ifo survey	240
C.6	Macro news and forecast revisions	242
C.7	Production expectations—robustness	243
C.8	Business situation as outcome variable (ifo survey)	244
C.9	Prices as outcome variable (ifo survey)	244
C.10	Alternative specifications	245
D.1	Relation of macro tail risk expectations and sales expectations	264
D.2	Distribution of firms’ expected exposure to a 5% decline in GDP	265
D.3	Macroeconomic tail risk expectations and investment decisions	266
D.4	Revisions in macroeconomic tail risk expectations and investment decisions	267
D.5	Expected probability of a large decline in GDP and quantitative investment	268
D.6	Correlations of macro beliefs based on newspaper data and other indices .	268
D.7	Macro expectations and news topics	269
D.8	Dictionary for classifying articles into news topics	270
D.10	Codebook	275
D.9	Classifications	281
D.11	Parameter calibration in the model (quarterly level)	288
D.12	Prediction power of planned decrease in investment on quantitative changes in investment	291

Preface

This thesis analyzes firms' expectations and behavior in macroeconomics and finance. A common theme across all chapters is the use of firm survey data which informs managerial decision-making in finance and, ultimately, macroeconomic models and policymaking. Firm survey data provides two key advantages for economic research.

First, survey data enhances our understanding of how managers form expectations and how these expectations influence their decisions. Nearly all major economic decisions are made under uncertainty and thus heavily rely on subjective expectations about the future. However, existing research on expectation formation has primarily focused on households and professional forecasters. Only recently have efforts been made to collect and analyze quantitative expectations from high-level decision-makers in firms (e.g., Altig et al., 2022; Coibion, Gorodnichenko and Kumar, 2018; Kumar, Gorodnichenko and Coibion, 2023; Bachmann et al., 2024*b*). In the third chapter of this thesis, I examine firms' expectation formation process and in the fourth chapter, I analyze the effects of firms' macro tail risk beliefs on their investment decisions.

Second, survey methods enable new possibilities to understand the behavior of firms at the micro level. Firm surveys generate novel identifying variation, providing insights into key economic mechanisms that would otherwise be difficult to analyze—a point recently emphasized by Stantcheva (2023). In the first chapter, I use expectations data to identify the causal effect of corporate taxes on investment. When formulating survey questions, a key trade-off is ensuring a clear mapping to economic models while maintaining clarity and simplicity to minimize survey fatigue. In addition, surveys should be designed to remain relevant across firms with different industry characteristics and decision-making structures. In the second and fourth chapters of the thesis, I introduce new questions to firm surveys, including hypothetical scenarios and open-ended text questions to generate novel data.

This thesis consists of four self-contained chapters that analyze how firms adjust investment after tax hikes (Chapter 1), how firms adjust prices after cost shocks (Chapter 2), how firms' expectations respond to news (Chapter 3), and how firms' investment is shaped by beliefs about macroeconomic tail risk (Chapter 4).

Investment and corporate taxes In the first chapter, I estimate the causal effect of corporate tax hikes on firms' investment. The chapter is based on joint work with Sebastian Link, Andreas Peichl, and Paul Schüle.

The identification builds on two pillars. First, we exploit the decentralized design of the German local business tax. Each year, municipalities decide on the statutory tax rates, resulting in significant variation across time and space. Second, we use panel data on both planned and realized investment volumes among a large and representative survey of German manufacturing firms. The unique feature is that investment plans for the subsequent year are reported in fall and tax hikes are typically announced in December. Thus, firms are surprised by the tax hike while the investment plans incorporate all other pieces of firms' information. We therefore focus on the revision of investment plans.

We find that firms revise their investment plans downward both economically and statistically significantly in response to a corporate tax hike. The share of firms investing less than previously planned increases by 3 percentage points (p.p.) following a tax hike. In terms of magnitude, a 1 p.p. increase in the tax rate reduces investment by 2.6–3.5%.

The effect sizes on investment are about twice as large during recessions. This state dependence underscores the Keynesian argument against raising taxes during recessions. We discuss three potential explanations for the stronger effects during recessions: greater cash flow sensitivity, heightened uncertainty about expected returns to investments, and fewer opportunities to shift the tax burden.

Price-setting and cost changes The second chapter, co-authored with Isabel Gödl-Hanisich, turns towards firms' price-setting decisions, a core question in monetary economics. Specifically, we use a new survey approach to study how firms adjust prices after cost increases, by directly eliciting the price pass-through over multiple horizons in response to the same initial shock.

A main advantage of this approach is that we are combining a quantitative measure of pass-through at the firm level with managers' beliefs captured in the survey and the usage of experimental methods.

We apply this survey approach with German firms in two different settings. First, we exploit the global supply shock in 2021/2022 as a source of cost-push shock and elicit pass-through in the field. Second, we elicit pass-through in survey experiments where we confront firms with hypothetical cost increases. This allows us to vary the nature of the shock—holding all other factors fixed.

We find that firms adjust prices only slowly after cost increases. An important driver is that firms want to avoid price deviations from competitors, so called micro real rigidities. For instance, firms raise prices multiple times in response to a permanent aggregate shock and they list competitive pressure as the most important limiting factor.

Furthermore, we find a particularly low pass-through when the cost increase is expected to be only temporary. This shows that firms’ expectations about shock duration is key for the understanding of pass-through dynamics.

A key feature of our direct elicitation method is that it enables a direct mapping to impulse responses in general equilibrium price-setting models. We discipline such a model with our survey estimates and find an important role of nominal and micro real rigidities that imply a persistent impact of monetary policy on the macroeconomy.

Expectations and news While the first two chapters mainly use existing survey data and new survey methods to understand economic mechanisms at the firm level, the third chapter analyzes the expectation formation process in firms. In collaboration with Benjamin Born, Zeno Enders, Gernot Müller, and Knut Niemann, we study how firms’ expectations respond to news. Crucially, we distinguish between micro news (that is, firm-specific developments) and macro news (information about the aggregate economy).

Our empirical analysis is guided by a stylized general-equilibrium model a la Lucas (1972). Firms live on islands and have imperfect information about the aggregate economy. Firms are not necessarily rational and may suffer from ‘island illusion’. That means that they systematically underestimate the extent to which the overall economy affects them. The departure from rational expectations is governed by a single parameter that jointly restricts the impact of micro and macro news: under island illusion, firms overreact to micro news and underreact to macro news.

We then test this idea using firm surveys from Italy and Germany. These surveys differ significantly in terms of questions, timing, and even variables of interest. But the results

are the same. Firms consistently overreact to micro news and underreact to macro news. We also find that the degree of over- and underreaction moves together across firms and over time, just like the model predicts.

Finally, we study how island illusion affects dynamics in the aggregate economy. We use a quantitative model from Lorenzoni (2009) and calibrate the degree of island illusion to match our empirical findings. For example, we find that inflation rebounds more gradually after a technology shock with island illusion.

Macro tail risk and investment The fourth chapter, which is single-authored, examines firms’ beliefs about the likelihood and exposure to future macroeconomic tail events—defined as a GDP drop of at least 5%—including events such as the Great Recession and COVID-19.

In the first part, I measure macro tail risk beliefs by eliciting the subjective probabilities of occurrence of four specific macro tail events. I collect these probabilities at three points in time in a large panel of German firms. Changes in these subjective probabilities co-move strongly with changes in firms’ expectations and uncertainty about their own business. I then show by using structured and open-ended survey questions that managers are highly uncertain about the impact of such an aggregate disaster on their own business. This introduces an additional layer of uncertainty that is typically not accounted for in the literature.

In the second part, I study how macro tail risk beliefs affect investment decisions. For that I use novel (i) firm-level observational variation, (ii) experimental variation from hypothetical scenarios, and (iii) time-series variation derived from newspaper articles. I consistently find that a higher subjective probability of a macro tail event leads to substantially lower investment. It turns out that around half of the relationship persists even after conditioning on firms’ subjective expectations and uncertainty about their own business. This is a key finding of the paper and suggests that macro tail risk beliefs matter in its own right, beyond adjustments in 1. and 2. moments. The relationship is most pronounced for firms expecting to be heavily exposed.

In the third part, I propose a model of firm dynamics that connects idiosyncratic and aggregate uncertainty. Motivated by my empirical evidence, firms are highly uncertain about their idiosyncratic exposure to macro tail events in the model. Specifically, I assume that the exposure is ambiguous. As the subjective probability of a macro tail event

increases, this idiosyncratic exposure uncertainty becomes more important. Firms incorporate this uncertainty into their decision making: they act even more cautiously than the higher macro tail risk itself would imply.

I then calibrate a quantitative model to match the survey evidence. I find that fiscal policy is particularly effective in stimulating investment when tail risk increases. In this policy counterfactual, the government commits to a switch to symmetric taxation of profits in macro tail events. Intuitively, this is very similar to generous loss carrybacks, implying transfers in the case of losses. Thereby, the policy provides a partial insurance against idiosyncratic exposure uncertainty.

Chapter 1

Downward Revision of Investment Decisions after Tax Hikes

This chapter is based on co-authored work with Sebastian Link, Andreas Peichl, and Paul Schüle and is published in the American Economic Journal: Economic Policy. See Link et al. (2024*b*) for the full reference. Regarding the use of generative models, I am fully responsible for the selection, adoption and all results of the generated output. I used DeepL and ChatGPT for spell checking and linguistic refinement of the text.

1.1 INTRODUCTION

The effect of corporate taxes on firm investment is a central question in macroeconomics and public finance. Corporate tax reforms like the US Tax Cuts and Jobs Act (TCJA) are often motivated by the argument that high corporate tax rates inhibit firm investment and growth (CEA, 2017). Standard theories of corporate taxation indeed predict that firms cut on investment projects if their after-tax net present value is reduced by tax increases (Hall and Jorgenson, 1967). To what degree corporate taxation affects investment, however, is ultimately an empirical question. Credible evidence on it is still scarce, as estimating the causal effect of corporate taxes on investment is challenging.

On the one hand, attributing cross-country discrepancies in investment behavior to differences in corporate tax rates is difficult to justify, as the timing of tax reforms often correlates with other macroeconomic determinants of firm investment. On the other hand, studies exploiting within-country variation need a valid control group and face the problem that many national-level tax reforms such as the TCJA change several parameters of the tax system simultaneously. For these reasons, quasi-experimental evidence on the response of investment to changes in the corporate tax burden originates predominantly from targeted tax deductions, which provide exogenous variation in exposure to tax decreases across firms of different size or in different industries (e.g., Zwick and Mahon, 2017; Garrett, Ohrn and Suárez Serrato, 2020; Ohrn, 2018). However, to what extent the effects of such specific policies generalize to changes in the corporate tax *rate* remains unclear.

This paper addresses this gap by combining the specific system of business taxation in Germany with unique data on firm-level investment plans and their realizations. Our identification strategy builds on two pillars. First, we exploit the decentralized design of the German local business tax (LBT): while tax base and liability criteria are set by the federal government, municipalities each year autonomously decide on the statutory tax rates.¹ We can therefore distinguish tax rate variation from potential changes in the tax base. Furthermore, municipalities adjust their taxes frequently. Restricting the analysis to tax increases, which are much more common than tax cuts, our identifying variation consists of 1,443 tax hikes between 1980 and 2018. The large number of tax hikes allows us to control for potentially heterogeneous time trends across regions or industries.

Second, we estimate the investment response of firms to these tax changes by leveraging panel data on both planned and realized investment volumes among a large, representative

¹This variation has been used by Fuest, Peichl and Siegloch (2018) to study the wage incidence of corporate taxation and Isphording et al. (2021) to assess the effects on R&D spending.

survey of on average 1,500 German manufacturing firms. The unique feature of our data is that each fall, firms report the planned volume of investment for the subsequent year. Municipalities announce tax changes for the subsequent year typically in December, i.e., after firms have reported their investment plans. In consequence, firms are surprised by the tax changes and have not included this information in their investment plans. At the same time, investment plans arguably incorporate all other (partially unobserved) private and public information of the firms that determine investment in the subsequent year.

Focusing on the revision of investment plans, i.e., the difference between the investment volume planned prior to the tax change and the investment volume ultimately realized, is advantageous from several perspectives.² Most importantly, investment revisions allow us to estimate the effect of corporate taxes on firm investment under weaker assumptions than usually possible. Because investment plans incorporate all relevant firm-level information, our results would still be unbiased if, for example, the occurrence of tax hikes were endogenous to local economic conditions. Moreover, considering revisions avoids problems with sensitivity in estimates due to the lumpy nature of investment, and hedges against potential bias in two-way fixed effects models (de Chaisemartin and D’Haultfœuille, 2022).

Our results show economically large and statistically significant investment responses for firms experiencing a tax increase. On average, the share of firms that invest less than previously planned increases by approximately 3 percentage points after a tax hike. In terms of magnitudes, a 1 percentage point increase in the LBT rate is associated with a decrease in the ratio of realized over planned investment by 2.3-3.8 percent, depending on the empirical specification. As firms on average invest approximately as much as previously planned, this maps into a semi-elasticity of investment with respect to the LBT rate of around 3. The corresponding elasticity of investment with respect to the net-of-tax rate is of similar magnitude. We verify our identification approach with an event study design, demonstrating that firms only deviate from the baseline probability for revising an investment decision in the year of the tax hike. While our baseline specification exploits variation in statutory tax rates (as previous literature on LBT in Germany did, see, e.g., Fuest, Peichl and Siegloch, 2018; Isphording et al., 2021), we find similar effects when relying on effective tax rates that are more common in studies for other countries and settings.

The magnitude of the investment response varies substantially over the business cycle. Compared to our baseline estimates, the share of firms that invest less than previously

²Comparing planned to realized quantities connects to the macro literature exploiting deviations from forecasts for identification (e.g., Romer and Romer, 2004).

planned in response to a tax hike is twice as large if taxes are increased during a recession. We discuss three potential explanations for this state dependence of tax shocks, relating to uncertainty about expected returns to investments, cashflow sensitivity, and tax incidence.

Our main contribution is to investigate the impact of hikes in the corporate tax rate on firm investment.³ While we are not the first to study this important question, we add to the literature along two dimensions. First and foremost, by using a novel identification strategy based on revisions of investment plans to investigate firms' investment response, we can eliminate concerns about omitted variable bias that have not been fully resolved in most of the previous literature. When using realized investments as outcome variable instead, results could be biased if tax policy responds to economic conditions. For example, Giroud and Rauh (2019) and Ivanov, Pettit and Whited (2024) investigate the effects of changes in US state level taxes on firm level outcomes.⁴ To be precise, the former paper studies the effects of tax changes on the reallocation of labor and capital across states while the latter looks at corporate leverage as the main outcome. In additional analyses, both papers also investigate—among others—effects on firms' capital stock. For identification, both papers rely on parallel trends between US states.⁵ While in our context the variation is on a more local level (municipalities within states), the key advantage is that assuming parallel trends is well justified by the fact that firms' ex ante planned volume of investment—i.e., the counterfactual level of investment in absence of a tax hike—should incorporate all firm-level information besides the tax shock that is relevant for investment in the subsequent year. That is, we do not require flat pre-trends in realized investment levels but only in terms of revisions of investment plans, which is much less demanding.

An alternative way to overcome endogeneity concerns is to focus on targeted tax deductions or accelerated depreciation allowances, giving rise to arguably exogenous variation in exposure to tax decreases for firms in different sectors and industries (Zwick and Mahon, 2017; Xu and Zwick, 2022; Garrett, Ohn and Suárez Serrato, 2020; Ohn, 2018; House and Shapiro, 2008; Curtis et al., 2021; Maffini, Xing and Devereux, 2019; Guceri and Albinowski, 2021). However, the extent to which the effects of such specific policies

³Other studies investigate firm-level responses to the corporate income tax along other margins (e.g., Auerbach, 2006; Suárez Serrato and Zidar, 2016; Fuest, Peichl and Siegloch, 2018; Garrett, Ohn and Suárez Serrato, 2020; Ljungqvist and Smolyansky, 2018). We add to this literature by providing new evidence on the investment response.

⁴Mertens and Ravn (2013) use aggregate data and combine a narrative approach with a structural VAR model to exploit changes in US federal corporate taxes.

⁵Both studies provide an extension using a narrative approach in the spirit of Romer and Romer (2010) that classifies arguably exogenous tax changes. Yet, these approaches only exclude a small number of potentially endogenous tax changes from the analysis and hence some concerns remain. Furthermore, the sample in Giroud and Rauh (2019) is restricted to large multi-state firms, and Ivanov, Pettit and Whited (2024) study tax decreases, whereas we focus on tax hikes.

generalize to changes of the tax rate remains unclear. Studying these targeted policies can therefore not substitute for a direct evaluation of the investment effects of changing the corporate tax rate, which affects all corporate firms at the same time and independently of their investment behavior. To the best of our knowledge, the only other paper using firm-level data and quasi-experimental variation to study the investment responses to a change in the universal corporate income tax rate is Harju, Koivisto and Matikka (2022).⁶ However, as the Finish corporate tax cut also entailed an increase in dividend taxation, they cannot consistently disentangle the effects of both channels. Moreover, the German setting has the advantage to offer substantially larger variation, especially in terms of the number of tax rate changes.

In addition, our findings of higher investment responses during recessions relate to an ongoing debate about the state dependence of fiscal multipliers (Auerbach and Gorodnichenko, 2013; Ghassibe and Zanetti, 2022; Ramey and Zubairy, 2018) and the state dependence of investment effects in response to tax changes more specifically (Demirel, 2021; Hayo and Mierzwa, 2021; Jones, Olson and Wohar, 2015; Ljungqvist and Smolyansky, 2018; Winberry, 2021). We complement this macroeconomic evidence by means of firm-level microdata and a distinct research design, showing that investment reacts much stronger to tax increases during recessions.

The remainder of the paper is structured as follows. Section 1.2 describes the municipality-level data on local business tax rates and the survey data on firm-level investment plans and their realizations. Section 1.3 presents our empirical strategy, while Section 1.4 documents the results. Section 1.5 concludes.

1.2 INSTITUTIONAL BACKGROUND AND DATA

To investigate the effects of corporate tax rate changes on firm investment, we merge municipality-level data on local business tax rates with unique data on firm-level investment plans and their realizations.

⁶In the German context, Dobbins and Jacob (2016) compare the differential investment responses of domestically and foreign-owned firms after a cut in the federal corporate tax rate in 2008. Lerche (2022) estimates the effects of an investment tax credit in East Germany on firms' production behavior.

1.2.1 The German local business tax

Institutional background The local business tax (LBT) is one of three types of taxes on business income in Germany. It is applied to the operating profits of both corporate and non-corporate firms. While tax base and liability criteria of the LBT are set at the federal level, municipalities decide autonomously on the tax rate. The tax rate consists of two components: a basic rate, which is determined by the federal government, and a local scaling factor, which is set at the municipal level. Each year, the municipal council has to vote on next year's scaling factor, even if it remains unchanged. As it is common practice to decide on next year's local scaling factor jointly with the adoption of the budget in the year's last meeting of the municipal council, tax changes are typically announced in December.⁷ Municipalities in our sample are approximately ten times more likely to increase rather than decrease their local scaling factor. In consequence, the identifying variation in our setting is too weak to consistently estimate the effect of tax decreases on investment.⁸ We thus restrict the analysis to tax changes induced by municipalities increasing their local scaling factors, henceforth referred to as a tax hike. This implies that the tax reforms exploited in this paper affect investment exclusively via increases in the tax rate, and not via changes in the tax base. Taxable profits of firms with establishments in more than one municipality are divided between municipalities according to formula apportionment based on the payroll share. Appendix A.1 provides additional details on the institutional setting.


Variation in business tax rates We use information on municipal tax scaling factors from the Statistical Offices of the German Federal States for the years 1980 to 2018. We enrich these data with information on municipality budgets and local economic conditions from several administrative data sources, leaving us with a panel of all German municipalities with extensive information on taxes, revenues and expenditures. To avoid capturing structural changes of the German reunification, and as data for East Germany are only available since 1990, we restrict our sample to West German municipalities (excl. West-Berlin). We furthermore exclude the few municipalities that underwent a municipal merger during the period of consideration, as we cannot determine their exact tax rates.⁹

⁷Appendix Figure A.1 substantiates this empirically, showing that newspaper coverage of LBT hikes indeed peaks each year in December.

⁸The number of tax decreases that could in principle be used in the analysis is very low. Combining the municipality-level data on LBT rates and the firm-level data from the IVS, our analysis could only exploit 236 firm-year observations (0.7% of all observations) that face a tax drop in a given year despite spanning a time frame of almost four decades.

⁹Municipal mergers were very frequent in East Germany after 1990 and this rule would also lead to an exclusion of many municipalities in East Germany.

Figure 1.1 – Variation in local business tax rates (1980-2018)



Notes: This figure shows the cross-sectional and time variation in municipal scaling factors of the German local business tax (LBT). Panel (A) plots the average LBT rate (in percent) induced by different scaling factors for the period 1980-2018. Panel (B) indicates the number of tax hikes, defined as an increase of the scaling factor. Municipalities in light grey areas are dropped from the sample as they are either located in East Germany or underwent a change of boundaries due to a merger. Moreover, we exclude observations where a tax hike was followed or preceded by another tax hike in the next or last two years.

There is substantial variation in LBT rates across municipalities and over time.¹⁰ As shown in Panel (A) of Figure 1.1, average tax rates differ strongly between municipalities, ranging from 12 to 34 percent. Panel (B) displays the identifying variation we rely on, i.e., the number of tax hikes between 1980 and 2018. Only few municipalities never increased the LBT in this period, while the median municipality increased the LBT rate three times and the median duration between two tax hikes in our sample is 13 years. The distribution of tax hikes is rather stable over time in terms of average size and dispersion (see Appendix Figure A.4). Importantly, past increases in the LBT contain very little predictive power for future tax hikes, as shown in Appendix Figure A.5.

After combining the municipality-level data on LBT rates with the firm-level investment data described in Section 1.2.2, we can exploit large parts of this variation in LBT rates.


¹⁰See Appendix A.2.1 for a more detailed description and investigation of the variation in LBT rates.

As summarized in the left panel of Table 1.1, our empirical strategy outlined in Section 1.3 relies on 1,443 tax hikes in 802 municipalities. The average tax hike amounts to 0.92 percentage points, corresponding to a 6 percent increase on average. The right panel summarizes the variation in tax hikes across firms. On average, approximately 7 percent of firms are exposed to a tax hike each year.

1.2.2 Firm-level data on revisions of investment plans

We use micro data on firms' investment behavior from the ifo Investment Survey (IVS, 2019). The IVS is conducted biannually (spring and fall) by the ifo Institute on behalf of the European Commission and covers a representative sample of incorporated firms in the German manufacturing sector.¹¹ The main purpose of the IVS is to obtain timely information on investment activity at disaggregated industry levels.¹² To achieve this goal, the IVS does not only elicit quantitative information on ex post realizations, but also on the planned volume of investment for the subsequent year. Thus, the panel structure of the IVS allows measuring how firms have revised their investment plans. In addition, survey participants provide quantitative information on revenues and the number of employees. The survey is usually completed by high-level management personnel at the firms' controlling departments (Sauer, Schasching and Wohlrabe, 2023).¹³

The timing of the survey is as follows:



¹¹Appendix A.2.2 and Sauer, Schasching and Wohlrabe (2023) provide additional information on the purpose and design of the survey, its representativeness, data access, and the wording of the survey questions used in the paper. The IVS micro data have been extensively used in recent research, e.g., Bachmann, Elstner and Hristov (2017); Bachmann and Zorn (2020); Link et al. (2023).

¹²The German Federal Statistical Office releases information on realized investment at the levels of disaggregated industries only with a time lag of two years.

¹³As noted on the cover letter of the survey, the ifo Institute guarantees compliance with strict data security criteria, that the data is evaluated in anonymous form, and that the survey results are only made available at the aggregate industry level. The participants therefore know that the firm-specific information reported to the IVS neither is available to their stakeholders, nor can be related to the municipality they are located in—and where the LBT is set. Therefore, incentives to strategically misreport the company's future investment activity are limited and unrelated to the variation in the LBT.

Table 1.1 – Tax hikes across municipalities and firms: Summary statistics

	Municipalities			Firm Observations			
	with Tax Hikes			with Tax Hikes		without Tax Hikes	
	N	Mean	SD	N	Share of Downward Revisions	N	Share of Downward Revisions
1980-1984	119	1.05	0.83	265	0.50	2655	0.45
1985-1989	131	0.92	0.53	340	0.49	4940	0.45
1990-1994	266	1.09	0.54	546	0.58	4831	0.54
1995-1999	228	0.94	0.47	385	0.51	4560	0.51
2000-2004	178	1.06	0.55	269	0.60	4711	0.58
2005-2009	106	0.78	0.51	161	0.66	4446	0.59
2010-2014	263	0.74	0.42	413	0.58	4118	0.58
2015-2018	152	0.69	0.38	248	0.63	2422	0.60
Full Sample	1443	0.92	0.54	2627	0.56	32683	0.54

Notes: This table reports summary statistics of the final sample used in the main analysis, i.e., after combining the municipality-level data on LBT rates and the firm-level data from the IVS. The left panel depicts the number of tax hikes at the municipality level that can be exploited in the empirical analysis along with the average size and standard deviation of these hikes. The right panel summarizes the number of firm observations that face a tax hike in a given year or not, as well as the average share of downward revisions of investment plans $((I_{i,t})/(E_{i,t-1}[I_{i,t}]) < 1)$ for each of these groups.

In the fall of year t_{-1} , firms report how much they plan to invest in equipment and buildings (in Euro) in the subsequent year t_0 , denoted $E_{t-1}[I_{t0}]$. The realized investment volume of year t_0 , I_{t0} , is elicited in the spring and fall survey of year t_1 .¹⁴ By comparing planned investment $E_{t-1}[I_{t0}]$ to realized investment I_{t0} , we observe whether firms in year t_0 invested more, less, or the same amount as previously planned. As municipalities announce the LBT rate for year t_0 at the end of year t_{-1} , i.e., after the fall survey, firms' investment plans for year t_0 reported to the IVS do not include information about changes in the LBT.

The investment data of the IVS have been shown to be very accurate. For instance, Bachmann and Zorn (2020) show that aggregate investment growth calculated from the microdata of the IVS is highly correlated with manufacturing investment growth reported by the Federal Statistical Office, and Sauer, Schasching and Wohlrabe (2023) report that

¹⁴Following Bachmann, Elstner and Hristov (2017), we take the average of I_{t0} if firms report it in both waves of year t_1 and drop the observation if these reports deviate more than 20% from the mean (see Appendix A.2.3 for details). The results are similar once we restrict the analysis to I_{t0} reported in the fall wave.


the average absolute deviation of the former from the latter is less than two percentage points. Moreover, Bachmann, Elstner and Hristov (2017) present a series of stylized facts on the cross-sectional and time-series properties of revisions of investment plans, i.e., the difference between ex ante planned and ex post realized investment volumes, showing that these deviations are meaningful along many dimensions. For example, they document that the overall distribution of revisions is not systematically skewed, while their cross-sectional average is procyclical.¹⁵ This indicates that participants provide accurate investment plans given their current level of knowledge at the time of the survey.

We restrict the sample to firms that report both planned and ex post realized volumes of investment, referring to all years t_0 for which we have information on LBT rates in the municipality of their location available. Following the protocol deposited in Appendix A.2.3, we keep those firms for which we can observe revisions in investment plans for at least five years and, following Fuest, Peichl and Siegloch (2018), drop firms with legal forms that are exempted from the LBT. Our final sample consists of 35,310 firm-year observations in years $t_0 \in \{1980, 2018\}$ that are spread across 1,192 municipalities in West Germany. According to the descriptive statistics presented in greater detail in Appendix A.2.3, the median firm in our sample is a typical representative of the “German Mittelstand” employing 264 workers, generating annual revenues of 45 million Euro (CPI inflation-adjusted and—if denominated in German marks—converted to 2015 Euros), and investing 1.4 million Euro each year. For each firm, we can rely on information on reported planned and realized investment volumes in, on average, 17 years. In the final sample, firms report zero investment in only 0.7% of all observations.

Importantly, investment plans for the next year contain a large amount of information that is highly predictive for the level of investment that is subsequently realized and that is changing within firms from year to year. The binned scatter plot depicted in Figure 1.2 demonstrates that the relationship between ex ante planned and ex post realized volumes of investment, i.e., $E_{t-1}[I_{t0}]$ and I_{t0} (both in logs), is highly linear and virtually corresponding to the 45 degree line. As depicted in Appendix Table A.4, 84% of the unconditional variation in (log) realized investment is explained by the investment plans for the respective year. Appendix A.2.3 presents a more detailed investigation of this relationship that, *inter alia*, demonstrates that investment plans, $E_{t-1}[I_{t0}]$, are much more strongly correlated with ex post realized investment, I_{t0} , than the realized level of investment in the previous year, I_{t-1} , and that these patterns even hold when controlling

¹⁵Relatedly, Appendix Figure A.8 shows that the investment plans are more frequently and more strongly revised downward during recessions.

Figure 1.2 – Relationship between planned and realized investment



Notes: This figure shows a binned scatter plot between ex ante planned and ex post realized levels of investment in year t_0 (both in logs) as reported by firms to the IVS in years t_{-1} and t_1 , respectively. The red line depicts a linear fit of the data. The sample is restricted to observations in years without tax changes.

for firm fixed effects. Taken together, investment plans contain accurate information on subsequent year's investment that goes beyond the extrapolation of the level of investment that was realized in the year these plans are reported to the IVS.

The raw data provide a first indication of the main result of the paper, i.e., that firms revise investment decisions downwards after tax hikes. For each five-year interval of the data, the right panel of Table 1.1 depicts the average share of downward revisions of investment plans separately for firms in municipalities with and without tax hikes. The share of downward revisions is—at least weakly—larger among treated firms than untreated firms in each time interval.¹⁶ We investigate this effect more systematically in the remainder of the paper.

1.3 EMPIRICAL STRATEGY

Research design We seek to identify the average treatment effect of an increase in the statutory LBT rate on firm investment. We consider a firm as treated in year t_0 if residing in a municipality that increased its LBT scaling factor from year t_{-1} to t_0 . The hypothesis

¹⁶In the pre-treatment year t_{-1} , the averages of both main outcome variables (Log Revision Ratio and Downward Revision Indicator) are not statistically different for firms that eventually are affected by a tax hike in year y_0 and firms ending up in the control group, see Panel B of Appendix Table A.3.

guiding our analysis is that firms surprised by the announcement of a tax hike in December of t_{-1} will on average invest less in year t_0 than previously planned. We therefore expect downward revisions of planned investment to be more frequent in municipalities that increased their local scaling factors. At the same time, firms’ investment plans elicited in the fall should incorporate all other, potentially unobserved, information influencing investment in the subsequent year.

Our identification strategy thus eliminates concerns about omitted variable bias. When using realized investment as outcome variable, results could be biased if tax policy responded to economic conditions, even after controlling for unit and time fixed effects, violating the parallel trends assumption. In our context this is quite different, as we observe the ex ante planned volume of investment—i.e., the counterfactual level of investment in absence of a tax hike—in addition to the ex post realized level of investment directly in our data. Using investment revisions instead of realized investment, we have a strong theoretical argument why we can extrapolate a (flat) pre-trend into the post-treatment period.

Hence, compared to using realized investment as dependent variable, our analysis only requires the weaker assumption that there are no unobserved factors that are both (i) correlated with investment and local tax policy in year t_0 , and (ii) not in the information set of the firm when forming investment plans in the fall of year t_{-1} . The only scenario that could violate this assumption would be a *local* shock that hits after firms have reported their investment plans, and that induces municipalities to implement a tax hike *within a few weeks*. Given the municipal decision structures and the “speed” of German bureaucracy, however, such an immediate response is highly unlikely. Relatedly, Fuest, Peichl and Siegloch (2018) show that changes in the LBT are typically not triggered by shocks to economic variables, and Blesse, Doerrenberg and Rauch (2019) demonstrate that tax setting of the municipalities substantially deviates from theoretically optimal behavior. As in the US (Robinson and Tazhitdinova, 2022), regional variation in corporate tax rates seems to be to a large extent idiosyncratic and not readily explained by standard theories of tax setting. Overall, we are therefore confident that omitted variables do not threaten identification in our setting.

Instead, a potential limitation of our identification strategy is that some firms may put a positive probability on the scenario that taxes will be increased in the subsequent year, whereas our analysis implicitly assumes that firms expect taxes to remain constant. To the extent that this was not true, there would exist two potential sources of bias pointing in opposite directions. A downward bias originating from the treatment group (where some firms revise investment less strongly), and an upward bias originating from the

control group (where some firms upward-revise investment if taxes are *not* increased). As long as the expected probability of a tax hike does not differ between treatment and control group, both biases will cancel out on average. Furthermore, any systematic and time-constant differences in the expected probability to be treated between firms and municipalities are inconsequential once we include firm fixed effects. However, private information about the likelihood of tax hikes may lead to systematically different beliefs in t_{-1} , the year before the LBT is increased (Riedel and Simmler, 2021). In particular, some firms may receive signals on the likely occurrence of a tax hike in the subsequent year even before investment plans are reported in the fall. In this case, they will—at least partially—incorporate this information into their investment plans and hence revise their investment decisions less strongly on average thereafter. If the private information helped firms to better predict the occurrence or absence of tax hikes, we hence should, if anything, tend to underestimate the investment response to a tax hike.¹⁷ In our data, however, we do not find evidence for a downward bias: Private information about future tax hikes should be more prevalent in smaller municipalities, where social ties to the municipality council are more likely, but treatment effects are not significantly different between cities and rural municipalities (see Figure 1.6, Panel B).

Measurement and estimation We use two variables to measure investment revisions. The first is an indicator for revising investment decisions downwards, defined as:

$$\text{Downward Revision:} \quad \mathbb{1} \left(\frac{I_{i,t}}{E_{i,t-1}(I_{i,t})} < 1 \right)$$

The downward revision indicator is attractive due to its robustness against outliers and non-linear investment responses. The second variable is the log revision ratio and takes the magnitude of each revision into account. It is defined as the natural logarithm of the ratio between realized and planned investment volumes:

$$\text{Log Revision Ratio:} \quad \ln \left(\frac{I_{i,t}}{E_{i,t-1}(I_{i,t})} \right)$$

We choose the logarithmic form due to the lumpy nature of investment, which means that the distribution of investment volumes is skewed and the revision ratio can get very large

¹⁷Moreover, the investment response might be underestimated due to the fact that firms pay the LBT according to the payroll share attributable to each municipality. As firm investment reported in the IVS refers to all domestic plants, the tax hike variation is measured with error for firms with plants in multiple municipalities. Although the IVS data lacks information on the prevalence and the payroll share of multi-establishments, the resulting attenuation bias is arguably small as only 7% of firms in the manufacturing sector were operating in multiple municipalities in 2017 according to aggregated, administrative LBT data (*Gewerbesteuerstatistik*). Furthermore, Panel A of Figure 1.6 demonstrates that treatment effects do not differ by firm size, a proxy for being a multi-plant company.

for small denominators. Moreover, the resulting estimates directly translate into the semi-elasticity of investment with respect to the tax rate, the relevant quantity of interest that we can directly compare to other estimates in the literature. As firms invest approximately as much as previously planned, the ratio of realized over planned investment is equal to one on average. As furthermore the revision ratio and realized investment are measured in logarithmic form, a tax hike which decreases the log revision ratio by 0.01 implies that both the revision ratio and realized investment decrease by 1%.¹⁸

In our main analysis, we estimate the following linear model by OLS:

$$\text{InvestmentRevision}_{i,m,t} = \gamma \text{TaxHike}_{m,t} + \mu_i + \phi_{l,t} + \psi_{s,t} + \varepsilon_{i,t}, \quad (1.1)$$

which explains the investment revision of firm i in municipality m and year t by municipality level tax hikes $\text{TaxHike}_{m,t}$ that take one of the following two forms:

$$\begin{aligned} \text{Tax Hike Indicator:} & \quad \mathbb{1}(\Delta tax_{m,t} > 0) \\ \text{Tax Hike in Percentage Points:} & \quad \Delta tax_{m,t} \end{aligned}$$

The tax hike indicator equals one if at time t municipality m increased the LBT. In addition, $\Delta tax_{m,t}$ denotes the tax change in percentage points. As discussed above, the focus on deviations of realized investment from the planned value should by itself rule out omitted variable bias.¹⁹ Still, some specifications additionally include firm fixed effects (μ_i) and year fixed effects at the level of industries ($\psi_{s,t}$) and federal states ($\phi_{l,t}$) to flexibly control for any time-invariant heterogeneity or systematic time trends in the probability of investment revisions and the frequency of tax hikes. In these specifications, we obtain a (generalized) Difference-in-Differences (DiD) estimate.²⁰ Standard errors are clustered at the municipality level.

¹⁸In more recent years, firms tend to invest on average slightly less than previously planned (compare Appendix Figure A.8). For this reason, the constant for the log revision ratio in Table 1.2 is not exactly zero, and the constant for the share of downward revisions is slightly larger than 0.5. As such, the semi-elasticity of the revision ratio with respect to an LBT hike will slightly understate the semi-elasticity of investment with respect to an LBT hike.

¹⁹While Section 1.2.2 demonstrates that the investment plans reported to the IVS contain valuable information that is highly predictive for ex post realized investment volumes, these variables might be elicited imprecisely. The resulting measurement error in the dependent variable should thus decrease the precision of our estimates without resulting in attenuation bias.

²⁰Note that using investment revisions as the outcome of interest implies that treatment effects realize exclusively in the treatment period. Due to this lack of treatment effect dynamics, the recent concerns about bias in two-way fixed effects models (e.g., de Chaisemartin and D’Haultfœuille, 2022) do not apply in this setting, as discussed below in more detail. It is, however, relevant in a setting when using realized

1.4 RESULTS

We present our results in three steps: first, we show our baseline results and their robustness along various dimensions. Second, we discuss how the magnitude of the effects relates to other estimates in the literature, before third, documenting effect heterogeneity over the business cycle.

1.4.1 Revision of investment plans after tax hikes: Main results

The baseline results presented in Table 1.2 reveal that firms affected by a tax hike strongly downward revise their investment decisions in the year this change is enacted. Panel (A) displays the estimates for the downward revision indicator. In Column (1), we compare the share of firms investing less than previously planned between municipalities where a tax hike is enacted and municipalities where the LBT rate did not change, without including any controls. We find that the share of firms that revise their investment decisions downwards is 2.7 percentage points higher in affected municipalities (Panel A1). The estimates presented in the remaining columns demonstrate that the point estimates for the tax hike indicator are barely affected by sequentially adding fixed effects at various dimensions, indicating that firms' investment plans already largely absorb regional and industry-specific shocks. For the size of the tax change (Panel A2), the inclusion of fixed effects tends to slightly increase the estimated coefficients. In Column (5), where we impose the most restrictive set of fixed effects, the effects of the tax hike indicator and the percentage change in the LBT on the probability of downward revising investment decisions are estimated at 3.3 and 2.4 percentage points, respectively.

Panel (B) repeats the analysis using the log revision ratio as dependent variable. The estimated coefficients are negative in all specifications of Panel B1, indicating that firms invest less than previously planned in response to a tax hike. While the effects are estimated less precisely compared to Panel A1, the point estimates are largely unaffected by the choice of the control vector. Again focusing on the most restrictive specification in Column (5), we find that the ratio of realized over planned investment decreases by 3.6 percent in response to a tax hike. Taking the magnitude of tax changes into account in Panel B2, the estimate in Column (5) implies that a 1 percentage point increase in the LBT rate is associated with a decrease in the revision ratio by 3.8 percent. Since in the absence of a tax hike firms invest approximately as much as they have planned, the ratio

investment (instead of revisions) as the outcome variable in Section 1.4.2. Inspired by Dube et al. (2023), firms are then assigned to another firm identifier in the middle between two tax hikes in order to ensure that there is only one treatment for each unit and to allow for different long-run trends.

Table 1.2 – Difference-in-differences: Investment revisions after a tax hike

	(1)	(2)	(3)	(4)	(5)
<i>Panel (A): Downward Revision</i>					
A1: Tax Hike Indicator: $\mathbb{1}(\Delta tax_{m,t} > 0)$					
	0.027	0.028	0.026	0.028	0.033
	(0.011)	(0.010)	(0.011)	(0.011)	(0.011)
Constant	0.536	0.536	0.536	0.536	0.535
	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
A2: Tax Hike in Percentage Points: $\Delta tax_{m,t}$					
	0.012	0.018	0.017	0.021	0.024
	(0.009)	(0.009)	(0.010)	(0.009)	(0.010)
Constant	0.537	0.536	0.536	0.536	0.536
	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
Observations	35310	35310	35310	35310	35310
<i>Panel (B): Log Revision Ratio</i>					
B1: Tax Hike Indicator: $\mathbb{1}(\Delta tax_{m,t} > 0)$					
	-0.031	-0.033	-0.025	-0.029	-0.036
	(0.016)	(0.015)	(0.017)	(0.016)	(0.017)
Constant	-0.033	-0.033	-0.033	-0.033	-0.032
	(0.006)	(0.006)	(0.001)	(0.001)	(0.001)
B2: Tax Hike in Percentage Points: $\Delta tax_{m,t}$					
	-0.023	-0.032	-0.028	-0.034	-0.038
	(0.014)	(0.013)	(0.014)	(0.014)	(0.016)
Constant	-0.033	-0.033	-0.033	-0.033	-0.032
	(0.006)	(0.006)	(0.001)	(0.001)	(0.001)
Observations	34421	34421	34421	34421	34421
Firm FE	-	-	✓	✓	✓
Year FE	-	✓	-	✓	-
Year × State FE	-	-	-	-	✓
Year × Industry FE	-	-	-	-	✓

Notes: This table reports estimates from linear regressions of Equation (1.1). “Downward Revision” is an indicator that is one if the fraction of realized investment over planned investment is below one. “Log Revision Ratio” is the natural logarithm of this ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. Industry fixed effects refer to the ifo industry classification, comparable to two-digit NACE industries. Standard errors in parentheses are clustered at the municipality level.

of realized over planned investment is close to one (and the log of the ratio is close to zero, as visible from the constant). Hence, our estimates directly map into a semi-elasticity of investment with respect to the LBT of around 3.

Overall, we find a clear and statistically significant negative investment response of firms to increases in corporate tax rates in all estimated models.

Economic size of the investment response The estimated investment response is economically sizable. To illustrate this, we conduct a back-of-the-envelope calculation, described in detail in Appendix A.4. According to our estimated semi-elasticity of 3, each additional Euro of tax revenues raised comes with a loss in firm investment of 2.12 Euro in the first year after a tax hike. If we also consider that lower firm investment reduces tax revenues in the medium run due to lower profits, the approximated investment loss for each additional Euro of tax revenue increases to a range between 2.14 Euro and 2.28 Euro, depending on the assumed strength of the second-round effect. While these projections rely on a series of simplifying assumptions, they still illustrate that the foregone volume of investment is non-negligible.

From this approximation of the behavioral response, we can also derive the marginal value of public funds (MVPF) in the spirit of Hendren and Sprung-Keyser (2020), given as:

$$MVPF = \frac{\text{Beneficiaries' Willingness to Pay}}{\text{Net Cost to Government}}$$

In our setting, firms are the beneficiaries and their willingness to pay is equal to the change of the tax burden. The net cost of the government equals the change of tax revenues plus the additional revenue changes via the behavioral response. According to this, our estimates point at a MVPF in the range between 1.01 and 1.08, i.e., slightly above one.²¹ However, given that investment is not the only margin of adjustment through which firms may react to increases in the LBT, the true MVPF is presumably larger. For example, Fuest, Peichl and Siegloch (2018) find that workers bear about half of the total burden of the LBT via lower wages, while employment is unaffected. The corresponding loss in payroll tax revenues would show up in the denominator of the MVPF formula and thus further raise the MVPF. Likewise, the cross-base elasticity with respect to the corporate

²¹The calculation is: $(19,800)/(19,800 - 139) = 1.01 \geq MVPF \leq 1.08 = (19,800)/(19,800 - 1,386)$, where the value of 19,800 EUR refers to the increase in overall tax revenues paid by the median firm in response to a one percentage point in increase in the LBT and the values of 139 EUR and 1,386 EUR denote the assumed upper and lower bound of the behavioral response, as described in Appendix A.4.

income tax base is presumably positive, so that corporate income tax revenues accruing to the federal government will decline as well. The MVPF should be even higher once these fiscal externalities were taken into account.

Validity of the identifying assumptions Next, we estimate an event study to test a central implication of our identifying assumptions: an increase in investment revisions should only occur in the year of the tax hike (t_0), while no effect should be visible in the years before, when the tax hike could not have been anticipated, i.e., we should observe parallel trends before the tax change. Moreover, a tax hike implemented in January of year t_0 should be known to the firm (a) when reporting its investment plans for year t_1 to the IVS (in the fall wave of year t_0) and (b) when reporting the actual volume of investment for year t_1 (in the spring or fall wave of year t_2). Hence, investment revisions should also not be systematically higher in any period after year t_0 .

The results of the event study regression presented in Figure 1.3 confirm that investment revisions occur immediately in t_0 when the tax hike is enacted. In contrast, the point estimates are (close to) zero in all other years, supporting the validity of our identifying assumptions. Panel A of Appendix Figure A.11 shows that these patterns also hold when extending the time window covered by the estimation to four years prior and post treatment.²²


Moreover, recent research in econometrics calls for caution when estimating two-way fixed effects models in generalized DiD settings with multiple treatment groups and periods (see, e.g., the survey by de Chaisemartin and D’Haultfoeuille, 2022), as these estimators only provide an unbiased DiD estimate if the treatment effect is constant between groups and over time. This problem is less relevant in our setting given that the estimated treatment effects are not dynamic, i.e., do not evolve over time (as shown above). Still, to demonstrate that the recent critique does not apply here, we repeat the event study using the imputation estimator proposed by Borusyak, Jaravel and Spiess (2024) as well as the interaction-weighted estimator by Sun and Abraham (2021). As shown in Appendix Figure A.12, results are very similar to Figure 1.3.

Robustness of main results Next, we demonstrate that our main results are robust along various dimensions. We start by highlighting that the fact that the estimates summarized in Table 1.2 are barely affected by sequentially adding fixed effects at various


²²The reason why we restrict ourselves to two pre- and post-event periods in Figure 1.3 is sample size. Because we always require that no other tax change happened in the pre- and post-event periods, extending the number of periods would shrink the size of the estimation sample considerably.

Figure 1.3 – Event study: Investment revision effect after a tax hike

(A) Year and firm-level fixed effects



(B) Full set of fixed effects



Notes: This figure shows the estimates of the following event-study regression: $\text{InvestmentRevision}_{i,t} = \sum_{j=-2}^2 \gamma_j \text{TaxHike}_{m,t}^j + \varepsilon_{i,t}$. In Panel (A), we additionally include year and firm fixed effects. In Panel (B), industry-year, state-year, and firm fixed effects are included. The reference period is t_{-1} . The dependent variable is based on the ratio of realized investments over planned investments (elicited in the fall of the previous year). “Downward Revision” is an indicator that is one if the ratio is below one. “Log Revision Ratio” is the natural logarithm of this ratio. Industry fixed effects are at the ifo industry classification level that is comparable to two-digit NACE industries. The confidence intervals refer to the levels of 90% (thick line) and 95% (thin line).

dimensions provides a first indication for the robustness of our main results. If confounding local shocks were important, estimates should vary across these different specifications,

which they do not.²³ This pattern suggests that firms' investment plans already incorporate shocks along various dimensions that might simultaneously affect firm investment and the municipalities' decisions to increase the LBT. Hence, focusing on deviations of realized investment from investment plans reported prior to the tax hike should by itself rule out many potential channels of omitted variable bias.

Nevertheless, attributing investment revisions to increases in the LBT could be problematic if tax hikes were accompanied by changes in municipality expenditures. If municipalities re-invested the additional tax revenue in local infrastructure, tax hikes would not only lead to higher tax payments on profits, but could also increase the value of local amenities for firms. If this created incentives for investment, this would counteract the direct effect of the tax hike and the true investment response would be underestimated. While this scenario is not implausible in general, we cannot detect concurrent expenditure shocks in our data. In line with evidence from Fuest, Peichl and Siegloch (2018) and Isphording et al. (2021), Appendix Figure A.13 shows that, on average, municipalities do not increase their expenditures jointly with the LBT.

Moreover, our results are robust to excluding the years after the German reunification from our sample. Although we only focus on firms located in West Germany, many of these firms were affected by this particularly turbulent economic time and their investment decisions were potentially affected by many investment subsidies that were introduced with the aim to foster investment in East Germany. Indeed, the estimated effect size is slightly, but not substantially, larger when excluding the time period after the German reunification (Appendix Table A.5).

As a final, more general robustness check, we conduct a permutation test by randomly assigning tax hikes to municipalities and, for each permutation, estimate Model (1) with both dependent variables, the downward revision indicator and the log revision ratio, along with the full set of fixed effects. Appendix Figure A.14 plots the cumulative distribution function of these placebo treatment coefficients. The non-parametric p-values obtained from this exercise are 0.0005 for the downward revision indicator and 0.0115 for the log revision ratio, and thus in the same order of magnitude as in our baseline regression.

²³Relatedly, Appendix Table A.9 shows that treatment effects are not heterogeneous for firms experiencing large revenue drops (compared to those firms who do not) as a proxy for a local (or even firm-specific) shock. This provides further empirical evidence that local shocks are not driving our results. This is not surprising, given that we analyze manufacturing firms whose products are tradable across municipalities and hence are less reliant on local markets.

1.4.2 Magnitude of effect size in comparison to the literature


While previous literature focuses on the effect of tax changes on the realized level of investment, a key novelty of our paper is studying the revision of investment plans. In order to facilitate the comparison of our result to previous findings, this section first demonstrates that our results regarding the downward revision of investment decisions can indeed be interpreted in terms of a reduction in realized investment of equal size. In a second step, we convert the identifying variation in the statutory LBT rate to changes in the net-of-tax rate, the effective tax rate, or the user cost of capital to show how our results compare to studies that rely on these frequently used specifications.

Effect of tax hikes on realized investment As argued in Section 1.3, we can directly interpret a one percent decrease in the log revision ratio as a one percent decrease in the realized level of investment because firms on average invest as much as previously planned. In order to demonstrate this empirically, Figure 1.4 plots the coefficients of three event study regressions using either the log revision ratio (red), the log realized investment volume (yellow), or the log planned investment volume (green) as dependent variable. As expected, the point estimates of both realized investment and the revision ratio are of comparable size in t_0 , i.e., the year of the tax hike, and indicate a drop by approximately 3%. Accordingly, investment plans fall to a persistently lower level only one period later in t_1 , when firms have incorporated the tax hike of year t_0 into their information set.

Note that Figure 1.4 is estimated on a different sample than our baseline estimates: Because treatment effects on investment levels are persistent and do not return to zero after one year, heterogeneous treatment effects may bias the estimates without further restrictions (e.g., de Chaisemartin and D’Haultfoeulle, 2022). Inspired by Dube et al. (2023), we therefore assign firms that experienced several tax hikes into distinct episodes to distinct firm identifiers, such that each spell contains a window around one tax hike only. For each of these spells, we assume that treatment effects have stabilized after three years and trim more distant periods. We further require that no other tax hikes took place during these periods, which—in total—reduces the sample size by 40%. In consequence, the coefficients are estimated less precisely than in our baseline specification. Note that despite these necessary sample restrictions, the coefficient for the log revision ratio in Figure 1.4 is still identified under much weaker assumptions than the coefficients for investment plans and realizations.

Overall, the results suggest that increases in the LBT have lasting effects on firms’ investment decisions. While—given that the tax hike is also incorporated into firms’ ex

Figure 1.4 – Effect of tax hike on investment plans, realizations, and revisions



Notes: This figure shows event-study estimates of log planned investment (green, short dashed lines), log realized investment (orange, long dashed lines), and the log revision ratio (red, solid lines) on the tax hike indicator and fixed effects at the levels of firm identifiers and years. The reference period is $t - 1$. In addition, end-periods $t - 3$ and $t + 3$ are binned and not shown. The sample is trimmed outside the event window. “Log Revision Ratio” is the natural logarithm of the ratio between the ex-post realized and ex-ante planned volume of investment. Inspired by Dube et al. (2023), when estimating the effects with respect to log planned and realized investment, firms are assigned to another firm identifier after the year that is in the middle between two tax hikes in order to ensure that there is only one treatment for each unit and to allow for different long-run trends. The confidence intervals refer to the significance levels of 90% (thick lines) and 95% (thin lines).

ante investment plans for years t_1 and thereafter—the coefficient on the log revision ratio returns to zero in the years following t_0 , the point estimates regarding the level of investment remain negative. Panel B of Appendix Figure A.11 shows that these patterns also hold when extending the time window covered by the estimation to four years before and after the treatment, which, by construction, relies on substantially fewer observations.

Effect sizes expressed in terms of net-of-tax rates A common quantity of interest in the public finance literature is the elasticity of investment with respect to the net-of-tax rate. To interpret our finding through the lens of this literature, Appendix Table A.7 re-estimates our baseline results regarding the log revision ratio after expressing the variation in the LBT in terms of changes in the net-of-tax rate, defined as $\log(1 - \tau_t) - \log(1 - \tau_{t-1})$. The resulting elasticity of investment with respect to the net-of-tax rate ranges between 2

and 3, depending on the choice of the control vector. In the most restrictive specification, which includes the full set of fixed effects, a 1% increase in the net-of-tax rate increases the log revision ratio and thus investment by 3%.

Effect sizes expressed in terms of effective tax rates Our main specification estimates the investment response to changes in statutory marginal tax rates, i.e., the LBT parameter that is directly set by municipal policymakers and that hence can be evaluated empirically without imposing further assumptions. However, large parts of the literature estimate treatment effects in relation to changes in effective marginal tax rates, i.e., also accounting for deductions – including depreciation rules – and other exemptions or tax credits (of which there are very few in the German context). To better compare our estimates with these studies, we thus also run such alternative specifications of our baseline estimation.

Our procedure to calculate effective marginal tax rates τ_{eff} , which is described in detail in Appendix A.5, follows the framework of Hall and Jorgenson (1967), as, e.g., recently applied by Furno (2022). Under this framework, the effective marginal tax rate is given by $\tau_{eff} = 1 - ((1 - \tau)/(1 - z * \tau))$, and depends only on the present discounted value (PDV) of the depreciation z and the statutory LBT rate τ .²⁴ To obtain z , we rely on information on depreciation schedules for machinery and buildings obtained from the Oxford Corporate Tax Database. Due to lack of the respective information, expressing changes in the LBT in terms of changes in effective marginal tax rates requires additional assumptions on, *inter alia*, (i) firms’ discount rate and (ii) the distribution of the total volume of investment among categories subject to different depreciation schedules, i.e., investment in machinery or buildings. The choice of the adequate discount rate is not innocuous in our setting, given that our analysis covers a period of almost four decades during which interest rates have fluctuated strongly (see Appendix Figure A.17). Moreover, the composition across investment categories can only be roughly approximated by either relying on yearly aggregates of the entire manufacturing sector or using time-invariant shares of firm-level investment in machinery and buildings.

To account for this, we compute four different versions of τ_{eff} based on two sets of assumptions regarding the discount rate and the relative share of investment in machinery and buildings, each. In the first and second specification, we follow Zwick and Mahon

²⁴This simplified version of the formula can be applied in our setting because there are no relevant tax credits in the German LBT that would complicate the calculation.

(2017) in assuming a time-constant discount rate of 7% when calculating the present discounted value of depreciation, while the remaining specifications use time-varying interest rates on loans for discounting.

Further, the first and third specification rely on information on the average share of investment in machinery and buildings obtained from aggregate data from the Federal Statistical Office of Germany, while the others use the firm-specific share of investment in machinery and buildings reported to the ifo Investment Survey whenever available.²⁵ Appendix Figure A.19 shows that across all specifications, the variation captured by changes in effective tax rates is strongly associated with the underlying changes in the LBT rate. This is not surprising, as all tax base rules of the LBT are set at the federal level (and potential changes of those over time are largely absorbed by year fixed effects) and no specific tax credits exist. Hence, apart from its scale, the identifying variation exploited in the empirical estimation does not differ strongly between the different approaches.

The results presented in Panel A of Table 1.3 show that the estimated effect sizes are largely comparable across the different specifications of τ_{eff} . For better comparison with our baseline results, the estimates have to be rescaled, as the effective tax rates are on average much smaller than the statutory ones. In the first specification, rescaling takes into account that the average statutory marginal tax rate of 16.79% is considerably larger than the average effective marginal tax rate (assuming a discount rate of 7%). The size and precision of the estimated coefficient shown in Column (1) is remarkably close to the respective baseline specification using variation in the statutory LBT rate (compare $-0.132 * 3.82/16.79 = -0.030$ vs. -0.038). When relying on time-varying interest rates, e.g., in Column (3), the rescaled point estimate implies a slightly lower effect size, which is, however, still in the same order of magnitude ($-0.136 * 2.9/16.79 = -0.023$).

Panel B of Table 1.3 repeats this exercise using hikes in the tax term of the user cost of capital as explanatory variable, again aiming to produce estimates comparable to other parts of prior literature. Given the linear relationship between the user cost of capital and the effective tax rate depicted in Appendix Figure A.20 ($\tau_{eff} = 1 - UserCost^{-1}$), this approach does not impact the results apart from rescaling the coefficients. Across all specifications, the estimated effects of a hike in the user cost of capital on the log revision ratio are statistically significant and range between -0.107 and -0.121 .

²⁵We use the firm-specific mean across all years if firms reported machinery and building investments at least three times to the IVS and replace missing values by the aggregate data used in the first specification, see Appendix A.5.

Table 1.3 – Investment revisions after a tax hike: Effective tax rates

	Log Revision Ratio			
	(1)	(2)	(3)	(4)
<i>Panel A: Variation in Effective Tax Rate</i>				
Effective Tax Hike	-0.132 (0.059)	-0.124 (0.057)	-0.136 (0.067)	-0.133 (0.066)
Constant	-0.033 (0.001)	-0.033 (0.001)	-0.033 (0.001)	-0.033 (0.001)
<i>Panel B: Variation in User Cost of Capital</i>				
User Cost Hike	-0.120 (0.054)	-0.107 (0.050)	-0.121 (0.061)	-0.111 (0.057)
Constant	-0.033 (0.001)	-0.033 (0.001)	-0.033 (0.001)	-0.033 (0.001)
Assumptions:				
Interest Rate	0.07	0.07	time-varying	time-varying
Spec. Investment Share	I	II	I	II
Observations	34421	34421	34421	34421
Firm FE	✓	✓	✓	✓
Year × State FE	✓	✓	✓	✓
Year × Industry FE	✓	✓	✓	✓

Notes: This table reports estimates from linear regressions of the log revision ratio on the size of the tax changes. In Panel A, the estimation is based on variation in effective tax rates (τ_{eff}) calculated as described in Appendix A.5. Panel B runs separate regressions exploiting changes in the user cost of capital (multiplied by 100). The respective specifications rely on different assumptions regarding the calculation of the present discounted value of depreciation, either assuming a time constant interest rate of 7% (as, e.g., in Zwick and Mahon, 2017), or based on time-varying interest rates on firm loans as depicted in Appendix Figure A.17. Further, τ_{eff} (and $UserCost$) is either calculated based on the average share of investment in machinery and buildings based on aggregate data from the Federal Statistical Office of Germany (Specification “I”) or on the firm-specific share of investment in machinery and buildings reported to the IVS whenever available (Specification “II”). All regressions apply firm fixed effects, as well as industry-by-year and state-by-year fixed effects. Standard errors in parentheses are clustered at the municipality level.

Comparison to the literature How do the investment effects documented in our paper compare to findings in other studies? The earlier public finance literature (e.g., surveyed by Hassett and Hubbard, 2002) typically estimated the effect of changes in the tax term of the user cost of capital on investment, measured relative to the lagged capital stock (I/K). For the sake of comparability, the coefficient of -0.12 depicted in Panel (B) of Table 1.3, which can be interpreted as the effect on log investment, as demonstrated above, hence needs to be expressed in terms of I/K . As information on the capital stock is not available in our data, we use the information on I/K documented by Zwick and

Table 1.4 – Summary of estimated (semi-)elasticities

	Log Revision Ratio	I/K
Change in Tax Rate	-0.038 (0.016)	-
Change in Net-of-Tax Rate (in %)	3.032 (1.312)	-
Change in Effective Tax Rate	-0.133 (0.066)	-
Change in User Cost of Capital	-0.111 (0.057)	-1.11 (0.57)

Notes: This table summarizes the main (semi-)elasticities reported in Section 1.4.2, relating the change in corporate tax rates to changes in the log revision ratio, and, as explained in the main text, to changes in realized investment. We always report the results of the most restrictive specification that controls for firm fixed effects as well as industry-year and state-year fixed effects. The I/K estimate is not directly estimated in our data, but obtained by rescaling the estimate for the log revision ratio with mean I/K of 0.1 reported by Zwick and Mahon (2017).

Mahon (2017) to rescale our estimate. Accordingly, a one unit change in the user cost of capital is associated with a decrease in the ratio of investment over the lagged capital stock by 1.2 percentage points in our setting.²⁶ Table 1.4 summarizes the main (semi-)elasticities reported in this section. The I/K transformation in the last column suggests that the investment response documented in our paper is slightly stronger compared to the estimates summarized in Hassett and Hubbard (2002) that range between -0.5 and -1 , but smaller than the comparatively large estimate of -1.6 found by Zwick and Mahon (2017).²⁷ Given that the investment stimulus studied in the latter paper was an explicitly recessionary policy, our evidence on state-dependence in the effect of tax hikes on firm investment delivers a potential explanation for the larger effect size found in their setting. Zwick and Mahon (2017) also find that loss-making firms are less responsive to tax shocks, in line with our evidence on smaller treatment effects among firms facing large revenue drops documented below (Appendix Table A.10). Furthermore, our semi-elasticity of 3 can be compared to Ohn (2018) who reports a semi-elasticity of 4.7 for a tax decrease induced by a specific tax provision for the manufacturing sector. Besides differences in research design and institutional setting, these larger responses documented in more recent studies might be due to targeted policies being more effective at stimulating investment than statutory tax rate cuts as rationalized by Chen et al. (2022).

²⁶Mean I/K in Zwick and Mahon (2017) amounts to 0.1. Hence, our estimate can be expressed in terms of I/K as follows: $-0.12/0.1 = -1.2$.

²⁷Ohn (2018, p. 296) also derives, after rescaling his estimates (Appendix J), an effect size for the same policy that is remarkably close to Zwick and Mahon (2017).

Two recent studies also estimate how firms respond to changes in universal corporate tax rates. Giroud and Rauh (2019) study how firm-level variables react to changes in US state level corporate taxes, including changes in the capital stock. They focus on a selected sample of large multi-state firms, for which they find a semi-elasticity of 0.24. As pointed out by the authors, this small elasticity might arise due to measurement error. Furthermore, the elasticity of capital can be temporarily lower than the elasticity of investment due to adjustment costs along the transition path in the firms' dynamic response. Mertens and Ravn (2013) use aggregate data and combine a narrative approach with a structural VAR model to exploit changes in US federal corporate taxes. They find semi-elasticities between 2.1 and 4, comparable to the responses documented in this paper.


In general, the elasticities reported in our paper are of course specific to the institutional setting of corporate taxation in Germany and the sample of firms affected. For instance, the elasticities we estimate for the manufacturing sector might differ from elasticities in other sectors, as manufacturing is more capital intensive than the overall economy. For the US, Cloyne, Kurt and Surico (2023) argues that tax cuts stimulate real investment activity mainly in the manufacturing sector, whereas firms in the services sector are more prone to adjust dividends, instead.

Finally, our results complement evidence from Isphording et al. (2021), showing that tax hikes in the German LBT reduce plant-level R&D spending by around 2 to 3 percent in the year of implementation. As R&D spending constitutes a (small) part of firm investment, we can directly compare the estimate to our semi-elasticity of investment of around 3. While both estimates suggest comparable effect sizes, our results are obtained for a different sample of firms and under less restrictive identifying assumptions.

1.4.3 State dependence and heterogeneity

State dependence Next, we exploit the long time dimension of our data to analyze potential heterogeneity in effect sizes over the business cycle. While a large literature in macroeconomics studies the state dependence of fiscal policy, there is not yet a consensus on whether effects of corporate tax changes are state dependent (Demirel, 2021; Hayo and Mierzwa, 2021; Jones, Olson and Wohar, 2015; Ljungqvist and Smolyansky, 2018; Winberry, 2021). As most quasi-experimental evaluations of the effect of corporate taxes on investment behavior rely on few tax changes or just a single tax reform, the treatment variation is typically not large enough to distinguish effect size heterogeneity along the

Figure 1.5 – Investment revisions after a tax hike: State dependence



Notes: This figure estimates how the probability of investing less than previously planned (left panel) or the log revision ratio (right panel) change in response to a tax hike separately for recession and non-recession years by including respective interaction terms in Equation (1). In Panel A, recession years are defined following the classification of the German Council of Economic Experts and refer to 1980-1982, 1992-1993, 2001-2003, and 2008-2009. Panel B classifies recessions as years with negative real GDP growth according to World Bank data (<https://data.worldbank.org/indicator/NY.GDP.MKTP.KD.ZG?locations=DE>), resulting in a smaller set of recession years (1982, 1993, 2002, 2003, and 2009). The estimation purges for firm fixed effects, as well as year fixed effects at the levels of federal states and industries. Standard errors are clustered at the municipality level. Confidence intervals refer to the 90% level. The full regression output is disclosed in the even columns of Appendix Table A.6.

business cycle. In contrast, the long time dimension of our data in combination with the occurrence of multiple local tax changes in each given year allows us to evaluate whether the treatment effect is state dependent.²⁸

The effect of tax hikes on revisions of investment plans are substantially stronger during recessions compared to normal times. Figure 1.5 presents the estimation results of interacting the tax hike treatment with indicators capturing periods of recession and normal

²⁸As shown in Appendix Figure A.3, municipalities are as likely to raise taxes in recessions as in normal times. The reasons why municipalities increase taxes (also in recession) are diverse, ranging from growing budget requirements to electoral cycles (Foremny and Riedel, 2014) and rent extraction (Langenmayr and Simmler, 2021); see also the discussion in Section 1.2.

times. To this end, Panel A classifies t_0 as a recession year if at least one quarter of that year is defined as a recession by the German Council of Economic Experts. The average effect that we estimated in Table 1.2 masks substantial heterogeneity over the business cycle. For instance, while in normal times the share of firms that invest less than previously planned increases by 2 percentage points in years with a tax hike, this figure triples to 6 to 7 percentage points in recessions. The same pattern also holds for the remaining specifications and the results tend to become even stronger when using a narrower classification of recession periods, defined as years with negative real GDP growth in Panel B.²⁹

Mechanisms of state dependence While our baseline estimates are in line with the predictions of standard theories of investment (Hall and Jorgenson, 1967), theory fails to explain why the effect of tax hikes should be state dependent. In the following, we discuss three channels which could explain the stronger effect during recessions.

The first channel relates to the fact that investment projects are risky. As investments are only partially deductible from the tax base, profits and losses are treated unequally by the tax authorities.³⁰ In expectation, tax hikes thus lead to stronger decreases in the net present value of those investment projects with a higher variance of expected returns as first formalized by Domar and Musgrave (1944).³¹ During recessions, the expected return to many investment projects becomes more uncertain, as it is unknown when the economy will recover again. Tax hikes should therefore lead to stronger behavioral responses in economic downturns when a higher share of planned investments is risky. While we cannot test this conclusively in our data, we can assess whether firms with more volatile revenue paths react stronger to tax hikes. For this purpose, we calculate the standard deviation of yearly revenue growth for each firm and construct an indicator for having above median volatility. Appendix Table A.8 shows the regression results when the tax hike effect is interacted with this volatility indicator. While the effects are estimated imprecisely and

²⁹Despite the fact that the effects during the recession period can only be estimated relatively imprecisely due to the small sample size, the estimated effects during recessions are statistically different from those during expansions in half of the specifications, while being close to approaching significance in the remaining specifications (see Appendix Table A.6).

³⁰As discussed by Fuest, Peichl and Siegloch (2018), costs of debt financing are usually fully deductible from the LBT, while costs of equity financing are not and loss offset is restricted. Moreover, due to depreciation rules, investment costs are split over several years while the revenues are fully taxed in each year.


³¹While Domar and Musgrave (1944) refer to the personal income tax, the same logic applies to the corporate tax and has been tested in the data. For state-level corporate tax rates in the US, Ljungqvist, Zhang and Zuo (2017) show that in response to a tax increase the average firm reduces risk as measured by their earnings volatility. Langenmayr and Lester (2017) find similar results in a cross-country panel and among small Spanish firms.

are sensitive to the specified model, they indeed show slightly larger responses of firms with more volatile revenue paths, suggesting that one reason for the state dependence of tax shocks may be the heightened uncertainty about returns to investment during recessions.

Second, firm investment is sensitive to cashflow (Almeida, Campello and Weisbach, 2004). Corporate taxes decrease the cashflow for profitable firms and therefore lower investment. At the same time, Almeida, Campello and Weisbach (2004) show that cashflow sensitivity is higher in recessions. During recessions, firms expect a higher probability of being cash constrained in the future and therefore retain more earnings for profitable investment opportunities. Taken together, this could give rise to an interaction effect, which reduces investment disproportionately if taxes are increased during recessions. Two regularities in our data support such a mechanism. First, we find that profitable firms react stronger to tax hikes during recessions. We use an indicator for a revenue drop by more than 10 percent compared to the previous year as a proxy for no longer being profitable. While firms that experience a large revenue drop in general revise investments downwards, the revision effect after a tax hike is smaller compared to firms without a large revenue drop during a recession (Appendix Table A.9). Firms with a large decline in revenues might still be profitable if they reduce their labor costs significantly. Appendix Table A.10 shows that the results hold when we exclude firms with a reduction in the number of employees by more than 5 percent as a robustness check. Second, if an adverse financing situation is reported to be a factor for a strong slowdown in investment volumes, the revision effect tends again to be larger in recessions (Appendix Table A.11). Both findings provide suggestive evidence that the stronger investment response in recessions may relate to cashflow sensitivity.

Finally, the stronger investment response to tax hikes in recessions could result from diminished possibilities to shift the tax burden to third parties. Fuest, Peichl and Siegloch (2018) show that workers bear approximately half of the incidence of the LBT in Germany. However, as wages are nominal downward rigid (e.g., Barattieri, Basu and Gottschalk, 2014), firms often cannot decrease wages in response to adverse economic conditions. This lower bound bites predominantly in recessions, especially given that collective bargaining agreements are still the norm in the German manufacturing sector and bargained wages usually slow down only with a considerable time lag as depicted in Appendix Figure A.15. This could suggest that during recessions, (cashflow sensitive) firms reduce their investment disproportionately as downward rigid wages do not allow shifting the tax burden on workers. Consistent with such a channel, Fuest, Peichl and Siegloch (2018) report a lower wage incidence for less profitable firms. In our data, we do not observe wages, preventing us from investigating this issue further.

Figure 1.6 – Testing for further heterogeneity



Notes: This figure estimates how the probability of investing less than previously planned (left panel) or the log revision ratio (right panel) change in response to a tax hike separately for different groups of firms by including respective interaction terms in Equation (1). Panel A provides separate estimates for small and large firms (split at the threshold of 250 employees). Panel B sorts firms according to their location using the definition of urban and rural areas of the Federal Institute for Research on Building, Urban Affairs, and Spatial Development (BBSR) that is mainly based on population density. In Panel C, we split the sample into municipalities with few (≤ 3) and many (> 3) tax hikes over the entire sample period. In Panel D, the tax hike treatment is split into cases where at least one tax hike has already occurred in the previous five years and where no tax hike occurred in the previous five years in the respective municipality. The estimation purges for firm fixed effects, as well as year fixed effects at the levels of federal states and industries. Standard errors are clustered at the municipality level. Confidence intervals refer to the 90% level. The full regression output is disclosed in the even columns of Appendix Tables A.12 and A.13.

Testing for further heterogeneity While the sample size of our data does not permit a comprehensive heterogeneity analysis, we perform the main estimation for a number of additional sample splits emphasized in the literature. For example, the investment effects of accelerated depreciation allowances in the corporate tax code have often been found to be much stronger among small (liquidity-constrained) firms (e.g., Zwick and Mahon, 2017). Figure 1.6 summarizes the results, showing that treatment effects do not differ by firm size when splitting our sample into firms with more or less than 250

employees (Panel A).³² We also fail to detect significant differences between rural and urban municipalities (Panel B), suggesting that a downward bias in our estimates due to private information of the firms—which should be more relevant in rural municipalities—is not a major concern.

In Panel C, we split the sample into municipalities with few (≤ 3) and many (> 3) tax hikes over the entire sample period. If firms in municipalities, which often increase the LBT, expected a tax hike with a higher probability, then downward revisions of investment should be less likely among these firms. However, effect sizes are again very similar for both groups. An alternative way to investigate whether tax setting dynamics at the municipality level correlate with the effect sizes is to split the sample by the occurrence of a tax hike in the last years. The results depicted in Panel D suggest that having experienced a tax hike in the last five years is plausibly associated with a larger investment response, although the estimates are not statistically different from each other in any specification (Appendix Table A.13). This result could be consistent with the notion that higher policy uncertainty triggers a stronger response after tax hikes.

Overall, Figure 1.6 demonstrates that other than the strong effect heterogeneity with respect to the business cycle, the effect of tax hikes on the investment behavior of firms is rather homogeneous across other important partitions of our data.

1.5 CONCLUSION

This paper provides novel empirical evidence on the effect of corporate taxation on firm investment. Our research design allows us to address several concerns that often complicate identification of an investment response. By considering 1,443 tax changes of the German local business tax between 1980 and 2018, we draw on extensive treatment variation and average out idiosyncratic characteristics of single tax reforms. By observing both planned and realized investment volumes, we can control for ex ante investment plans when estimating the effect of tax hikes on firm investment, eliminating a wide set of further potentially confounding factors.

³²Given that firm size is often been used as a measure of cashflow sensitivity, this finding might seem inconsistent with the aforementioned suggestive evidence that the stronger investment response in recessions may relate to cashflow sensitivity. There are at least two possible explanations for this discrepancy. First, the share of *very* small manufacturing firms, which arguably suffer most from adverse financing conditions, in our sample is small (Appendix Table A.2). Second, Appendix Figure A.16 demonstrates that small firms are not only more likely to be cashflow sensitive, but also more likely to underinvest due to a weak earnings situation. As the LBT taxes profits and less profitable firms should thus react *less* sensitively to tax hikes, the latter mechanism could counteract the former.

We find significant and economically large investment responses for firms experiencing a tax shock. The share of firms that invest less than previously planned increases by 3 percentage points after a tax hike, with strong heterogeneity along the business cycle. While in normal times the share of firms that revise their investment decisions downwards increases by 2 percentage points in response to a tax hike, this figure triples to over 6 percentage points if taxes are increased during a recession. These findings have direct policy implications that support the countercyclical Keynesian notion of “do not increase taxes during recessions”. While we find suggestive evidence that the state dependence of tax shocks could plausibly be related to uncertainty about expected returns to investments, cashflow sensitivity, and tax incidence, more research is needed to disentangle the channels behind this finding.

Overall, our results confirm the view that investment decreases substantially in the corporate tax burden. While our estimates were obtained for increases in the statutory corporate tax rate, prior studies have often evaluated targeted tax policies which were deliberately designed to stimulate investment. We look forward to future research comparing the effects of both types of policies within a unified framework.

Chapter 2

Firms' Pass-Through Dynamics: A Survey Approach

This chapter is based on co-authored work with Isabel Gödl-Hanisch. A previous version has been published as a CEPR Discussion Paper. See Gödl-Hanisch and Menkhoff (2023) for the full reference. Currently, we revise the paper for publication in the *Journal of Political Economy Macroeconomics*. Regarding the use of generative models, I am fully responsible for the selection, adoption and all results of the generated output. I used DeepL and ChatGPT for spell checking and linguistic refinement of the text.

2.1 INTRODUCTION

Price stickiness plays a crucial role in New Keynesian models and has been extensively studied empirically (e.g., Nakamura and Steinsson, 2008; Klenow and Kryvtsov, 2008). However, even with rich granular pricing data at hand, it is challenging to analyze price dynamics, especially at the firm level. Therefore, the understanding of frictions in price-setting is still incomplete. For instance, nominal rigidities, as measured by the frequency of price changes, are not sufficient to explain strong monetary non-neutrality. One way to generate more monetary persistence is to introduce micro real rigidities creating strategic complementarities that impose costs on firms when they set prices that deviate from those of their competitors (Ball and Romer, 1990).

While these micro rigidities are often used in theoretical models, their quantitative significance is still in question (e.g., Klenow and Willis, 2016). Moreover, the empirical literature provides little insight into the importance of expectations for pricing decisions, such as the expected persistence of shocks, even though future marginal costs are a key determinant in standard price-setting models.

We use a new approach to study price dynamics by *directly* eliciting pass-through of cost increases in a large-scale business survey. We integrate the advantages of a quantitative measure with the richness of firm characteristics and beliefs captured in the survey to provide a complete picture of how firms adjust prices after cost increases. Eliciting price adjustments across multiple horizons in response to the same initial shock allows for a comprehensive examination of pass-through *dynamics* and enables a direct mapping to impulse responses in general equilibrium price-setting models.

We proceed in three steps. First, we document five facts about pass-through in the field, emphasizing the sluggishness of pass-through due to a combination of observed nominal rigidities—infrequent price adjustments—and micro real rigidities—multiple small price increases due to competition. The sluggishness amplifies when firms believe that the shock is less persistent. Second, our tailored survey experiments provide new causal insights into the importance of real rigidities: firms raise prices multiple times in response to a permanent aggregate shock, and idiosyncratic shocks of the same size have a lower pass-through than aggregate shocks. Further, we find that firms’ pass-through is significantly lower for more transitory shocks but is similar across different economic environments. Third, we use our empirical estimates to discipline menu costs and real rigidities in a general equilibrium price-setting model. The observed sluggish pass-through requires a high degree of both nominal and real rigidities, bringing us a big step closer to the high persistence of monetary non-neutrality found in other studies.

We start with evidence from the ifo Institute’s business survey panel, a monthly survey covering 6,500 firms across all sectors of the German economy. We directly ask about the pass-through of higher prices for energy, raw materials, and intermediate inputs multiple times since the onset of the cost surge in 2021 in a backward- and forward-looking manner for the next six months. The survey also contains information on the business’s current standing and outlook and unique supplementary questions, such as subjective beliefs about the duration of supply chain disruptions.

We present five facts based on our novel survey approach. First, firms pass through cost increases only gradually over time. This is due to a combination of nominal and micro real rigidities. Competitive pressure is most often mentioned as a limiting factor for pass-through. Second, there is substantial variation across and especially within narrowly defined industries. Understanding this heterogeneity is the goal of the analysis that follows. In our third set of facts, we focus on the expected persistence of shocks. Specifically, we are the first to document that pass-through increases in the firm’s expected duration of the shock. In turn, a longer duration of the shock also increases the probability that firms actually raise prices. Fourth, we demonstrate that pass-through increases with the firm’s idiosyncratic uncertainty, shedding light on a theoretically ambiguous relationship (Vavra, 2014). Fifth, leveraging the survey’s panel dimension, we establish a strong, positive link between pass-through and ex-ante price-setting frequency.

To further investigate the causal link between the nature of the shock and pass-through, we conduct survey experiments in the ifo Management-Survey in the second part of the paper. We confront firms with hypothetical scenarios about facing an exogenous global supply shock raising costs by 20%, but varying the nature of the shock. In the style of supplementary questions in the first part of the paper, we ask firms about the extent to which they pass on these cost increases to customers over specific time horizons of one to 24 months. By exploiting within-firm variation, we discern the marginal effect of shock duration (permanent vs. transitory) and micro real rigidities, comparing idiosyncratic vs. aggregate shocks. Between-firm variation allows us to analyze the effects of economic conditions (low vs. high demand). We are the first paper to provide causal evidence on the relevance of these factors and the economic size of pass-through at different horizons.

We find a concave increase in pass-through over time for all hypothetical scenarios at the aggregate level. In the first three months, pass-through evolves similarly, independent of the duration of the shock, driven primarily by the extensive margin. In the medium term, the permanent shock scenario shows significantly higher pass-through compared to the temporary and uncertain shock duration scenario. Most firms exhibit gradual pass-through increases instead of immediate adjustments, with lower pass-through for

idiosyncratic shocks, supporting the presence of micro real rigidities. The results indicate no strong state dependence, with similar price dynamics for two assumed economic conditions: low and high aggregate demand. We can use our pass-through estimates to infer the slope of the Phillips curve using sufficient statistics from Auclert et al. (2024): the slope flattens significantly when accounting for real rigidities, but is still steeper than recent empirical estimates (Hazell et al., 2022). Although relying on a different method and survey, the results of the survey experiments corroborate our empirical analysis in the first part. This demonstrates the robustness of our direct elicitation approach. Both sets of results emphasize the existence of nominal and, especially, real rigidities, as well as the important role of the expected nature of the shock for pass-through.

In the third part, we introduce micro real rigidities into an otherwise standard general equilibrium price-setting model (Nakamura and Steinsson, 2013). We discipline the degree of nominal and real rigidities in the model by targeting the pass-through dynamics of a permanent aggregate shock and the price-setting frequency based on our empirical results. As an external validation, we compare the pass-through dynamics of transitory aggregate shocks and the difference in pass-through relative to idiosyncratic shocks. In both cases, the model predictions closely align with the empirical counterparts. The implied real effects of monetary policy in the estimated model are sizable. This is mainly due to micro real rigidities, which brings us already a long way to the estimated persistence of monetary policy in other studies (see, e.g., Christiano, Eichenbaum and Evans, 1999; Romer and Romer, 2004; Jarociński and Karadi, 2020). However, a remaining gap leaves room for other factors like macro real rigidities and imperfect and heterogeneous information about the shock’s nature.

Related literature Our results contribute to four strands of the literature. First, we provide new evidence to the literature that empirically characterizes price dynamics (Bils, Klenow and Malin, 2012; Beck and Lein, 2020; Eichenbaum et al., 2014; Eichenbaum, Jaimovich and Rebelo, 2011) and estimates pass-through (Auer and Schoenle, 2016; Garetto, 2016; Amiti, Itskhoki and Konings, 2019; Gopinath and Itskhoki, 2010; Dedola, Kristoffersen and Zullig, 2021; Joussier, Martin and Mejean, 2022). These papers estimate quantitative pass-through based on granular administrative price data and typically focus on exchange rate pass-through.¹ Mostly, they find incomplete and—if dynamics are estimated—sluggish pass-through. However, these studies frequently face limitations in analyzing underlying mechanisms due to missing additional data and challenges in precise pass-through estimation. Complementarily, there is a literature employing surveys to im-

¹Smets and Wouters (2007) estimate the degree of micro real rigidities using aggregate data.

prove the understanding of price-setting (Blinder et al., 1998; Loupias and Sevestre, 2013; Fabiani et al., 2005; Bachmann et al., 2019; Bunn et al., 2022). However, these papers do not provide dynamic and quantitative pass-through estimates. A key contribution of this paper is the ability to bridge both approaches by providing quantitative firm-level estimates of pass-through, while taking advantage of the survey approach to understand the observed dynamics. Importantly, this is the first paper leveraging experimental methods to inform the literature about the underlying mechanisms.²

Second, our findings on pass-through are of first-order importance for the degree of monetary non-neutrality. Our contribution to the literature is twofold. First, we provide evidence for the microfoundation in New Keynesian models that often rely on strong micro real rigidities to generate monetary non-neutrality (e.g., Smets and Wouters, 2007). Second, our hypothetical vignette allows us to infer the slope of the Phillips curve. Thereby, we contribute an estimate based on an entirely different approach compared to the literature (e.g., Hazell et al., 2022; Ball and Mazumder, 2011; Phillips, 1958; Coibion and Gorodnichenko, 2015*b*).

From a methodological perspective, we contribute to a third strand of the literature on using hypothetical scenarios to retrieve crucial economic mechanisms that are otherwise difficult to assess, see Fuster and Zafar (2023) for an overview. So far, most papers apply hypothetical vignettes in household surveys. For instance, this method is used to improve the estimates and understanding of the marginal propensity to consume (Fuster, Kaplan and Zafar, 2021; Jappelli and Pistaferri, 2014, 2020; Colarieti, Mei and Stantcheva, 2024), education choices (Wiswall and Zafar, 2021), late-in-life savings (Ameriks et al., 2020), subjective models of the macroeconomy (Andre et al., 2022), or the effect of monetary policy on consumption (Roth, Wiederholt and Wohlfart, 2023). Only few papers employ this approach in business surveys (Dibiassi, Mikosch and Sarferaz, 2021; Drechsel et al., 2022; de Bruin et al., 2023). We add to this emerging literature by analyzing pass-through dynamics after cost shocks using designed vignettes.

Lastly, we speak to a literature on modeling firms' price-setting in general equilibrium (Nakamura and Steinsson, 2013; Wang and Werning, 2022; Aruoba et al., 2022; Beck and Lein, 2020; Klenow and Willis, 2016; Golosov and Lucas, 2007; Blanco et al., 2022; Karadi and Reiff, 2019; Midrigan, 2011). We contribute by augmenting a standard menu-cost model with micro real rigidities and mapping it to our novel survey results.

²In parallel and independent work, de Bruin et al. (2023) study pass-through using a hypothetical vignette. In contrast to them, we elicit pass-through (i) directly, (ii) across multiple horizons, and (iii) distinguish between aggregate and idiosyncratic shocks. Importantly, our scenarios are independent of the current economic stance and vary the nature of the shock.

The remainder of this paper is structured as follows. Section 2.2 documents five facts based on the novel survey approach. Section 2.3 presents the survey experiment results. Section 4.4 estimates a general equilibrium price-setting model. Section 4.5 concludes.

2.2 EMPIRICAL RESULTS

This section documents five stylized facts about firms' pass-through over time and space in response to a historic supply shock.

2.2.1 Data

We build on the well-established ifo Institute's business survey panel data (IBS-IND, 2023; IBS-TRA, 2023; IBS-SERV, 2023).³ The monthly representative business survey covers approximately 6,500 firms across all sectors of the economy, containing information on the business's current standing and expectations about the business's outlook (Sauer, Schasching and Wohlrabe, 2023). At the heart of our analysis is a *supplementary* quantitative pass-through question posed at three points in time since the beginning of the surge in input prices as illustrated in Appendix Figure B.1 (June 2021, April 2022, and October 2022). The specific wording of the question is:

To what extent do you [the firm] pass through higher prices for energy, raw material, and intermediate inputs to your customers? —%

Notably, the October 2022 survey further differentiates between pass-through so far and planned pass-through over the next six months, asking in a backward- and forward-looking manner while holding the firm information set constant.⁴ Conceptually, we can summarize the elicited pass-through of cost increases at the firm level for the two dates, PT_t and $\mathbb{E}_t[PT_{t+6}]$, as follows:⁵

$$PT_t = \frac{\partial \log P_t}{\partial \log mc_t} = f(\mathbf{s}, \mathbb{E}_t[\mathbf{x}|\Omega_t]); \quad \mathbb{E}_t[PT_{t+6}] = \frac{\partial \log P_{t+6}}{\partial \log mc_t} = g(\mathbf{s}, \mathbb{E}_t[\mathbf{x}|\Omega_t])$$

³For example, Bachmann, Elstner and Sims (2013); Bachmann et al. (2019) use the ifo Institute's business survey panel data to study the role of firm-level uncertainty for price setting and economic activity.

⁴Gödl-Hanisch and Menkhoff (2022) summarize insights on the gradual pass-through in a policy report.

⁵We interpret $\partial \log mc_t$ as the increase in input prices and $\partial \log P_t$ as the respective price change since the onset of the supply shock in January 2021 at the firm level.


Pass-through is a function of $\mathbf{s}$, reflecting perfectly observable state variables (e.g., current marginal costs mc_t and the degree of competition), and $\mathbb{E}_t[\mathbf{x}|\Omega_t]$, the expected path of marginal costs and other relevant price determinants in the future. Importantly, a given firm reports pass-through for t and $t + 6$ based on the same state variables and information set Ω_t . This allows us to analyze pass-through at two points in time for the same underlying shock.

In addition, we explore other supplementary questions, such as subjective beliefs about the expected duration of supply chain disruptions. We refer to Appendix B.1 for more details.

Our survey approach has the key advantage of directly asking firms about the level of pass-through, as opposed to imputing pass-through measures from potentially imperfectly measured cost shocks and price estimates. Moreover, it allows to account for potential heterogeneity across firms and industries, which is challenging to gauge in other data sets, particularly micro-level pricing data. Nonetheless, a potential limitation to our approach is that firms might strategically underreport their true pass-through. We offer three reasons why this is unlikely to be the case in our survey. First, the survey is incentivized by providing granular survey results for participants whose own weight is too small to affect average results. Second, past supplementary questions in the survey suggest that firms do not answer strategically. For instance, Bachmann et al. (2022b) elicited potential production reductions in response to a 10% (or 50%) cut in gas supply. The median firm reported no decline in production (or a 25 % decline). These comparatively low numbers suggest no strategic reporting on average. Furthermore, if firms differed in the extent of strategic reporting, one would expect a positive correlation between a high reported decline in production in the case of a gas supply cut and a low reported pass-through in the cross-section. However, no such relationship is evident (see Appendix Figure B.2). Third, our estimates align with other pass-through estimates in the literature as shown in Appendix Figure B.10.⁶

⁶To further validate pass-through estimate reliability, we explore the connection between firm pass-through and return on sales (ROS) as a profitability indicator. We expect firms reporting low pass-through to show reduced ex-post ROS. Monopoly theory also predicts a negative correlation between ex-ante profitability and pass-through. Appendix Table B.1 indicates that pass-through is positively linked with ex-post ROS and negatively with ex-ante ROS, using October 2022 pass-through regression on 2019 and 2022 ROS data.

Figure 2.1 – Pass-through dynamics over time



Notes: The figure shows the average planned pass-through by April 2023 as a function of the level of pass-through by October 2022, conditional on cost changes until October 2022. The gray shaded area is the 95% confidence interval. Source: ifo Institute’s business survey.

2.2.2 Five facts about pass-through dynamics

Pass-through dynamics over time We start by examining the dynamics of aggregate pass-through over time. The average pass-through for the entire sample is around 34% in October 2022 and is expected to increase by 17 percentage points (p.p.) in the following six months (see Table 2.1). The gradual increase in average pass-through could be entirely driven by the extensive margin. To further illustrate the gradual increase in pass-through over time, Figure 2.1 shows the planned pass-through by April 2023 as a function of the level of pass-through by October 2022. The difference between the smoothed line and the 45-degree line corresponds to the planned change in pass-through. Firms with an incomplete pass-through by October 2022 expect to increase pass-through by April 2023. For instance, firms at a pass-through level of 20% in October 2022 plan to increase pass-through by another 20 p.p. to 40% over the next six months. Similarly, firms at a zero pass-through level plan to increase pass-through by almost 20 p.p. To put these numbers in perspective, a constant pass-through rate of 20 p.p. increase every six months would imply that firms would reach full pass-through over the course of 2.5 years.

The substantial share of zero adjustments in October 2022 indicates that firms face nominal rigidities such as menu costs. In standard price-setting models, firms base their pricing decisions on the discounted sum of expected marginal costs. While nominal rigidities can

lead to incomplete pass-through, as firms might expect only a transitory cost increase (see Fact 3), firms would not plan a gradual increase of pass-through given the same information set. Hence, the gradual and sluggish adjustment over time supports the additional presence of micro real rigidities, primarily affecting the intensive margin of adjustment and leading to a lower adjustment per price hike and over time. Micro real rigidities imply strategic complementarities in price setting between firms and thereby disproportional costs of deviating prices from competitors.

To gain a deeper understanding of incomplete pass-through, we ask the participants directly about the limiting factors in the October 2022 survey. We distinguish between competition, weak demand, long-term contracts, administrative effort/burden, regulatory reasons, and other factors. About 70% of firms report that competition limits their potential to adjust prices, a further indication of the importance of micro real rigidities (Appendix Figure B.3). Similarly, about 45% listed that weak demand limits pass-through. Also, factors related to nominal rigidities are present: about 35% report that long-term contracts limit pass-through, and administrative burden and regulation have a share of less than 10%.

Accordingly, we can summarize our first fact as follows:

Fact 1: Firms pass through cost increases only gradually over time due to nominal and real rigidities.

Pass-through dynamics within and across industries Direct elicitation in the representative survey allows us to study pass-through dynamics at a disaggregated level, both within and across industries. There is substantial heterogeneity in the degree of pass-through across firms: while about 30% of firms have not passed through any cost changes by October 2022, 12% have passed through 50% of cost changes, and less than 10% of firms fully pass through cost increases to their customers (Appendix Figure B.4). The pattern looks similar for the planned pass-through until April 2023, but we observe a right shift. Almost 30% of firms now expect to fully pass through cost increases to customers. At the same time, the share of firms with zero pass-through declined from 30% to 15%, indicating an important role of nominal rigidities, particularly in the short run. Appendix Figure B.5 delivers similar evidence for the survey rounds in June 2021 and April 2022.

Table 2.1 presents the average pass-through levels across broadly defined industries. We observe the lowest pass-through for services firms (19% until October 2022) and the highest pass-through for manufacturing firms (48% until October 2022). The difference

is mainly explained by the extensive margin: 91% of manufacturing firms pass through costs at least to some extent, while this is true for less than half of the firms in the services industry (46%). Revisiting planned pass-through until April 2023 as a function of pass-through until October 2022, we find that the pass-through shape holds for all industries: manufacturing, services, and trade, as shown in Appendix Figure B.6, with the trade sector exhibiting a slightly slower pass-through pace.

Table 2.1 – Pass-through over time across industries

	Total	Manufacturing	Trade	Services	Construction
October 2022	34%	48%	36%	19%	40%
April 2023	51%	65%	50%	33%	65%
N	6407	1960	1445	2177	825

Notes: The table displays the average pass-through of cost changes to sales prices in October 2022 and planned pass-through until April 2023 in percent by industry. Source: ifo Institute’s business survey.

Even within narrowly defined industries, there remains substantial cross-sectional variation. Decomposing the total variation into within-industry and between-industry variation shows that the former outweighs the latter in Figure 2.2. At the 2-digit industry level, between-industry variation explains roughly 20%, and within-industry the remaining 80%. By construction, the between-component gains slightly higher importance at the 4-digit industry level, but still, the within-industry variation explains roughly 75% of variation in pass-through. One potential concern is that measurement errors bias the within component upwards. To address this point, we restrict our sample to respondents who spent sufficient time on the question following Giglio et al. (2021). Appendix Figure B.7 shows similar estimates when cutting out the bottom 25% of the time spent distribution.⁷


Accordingly, we can summarize our second fact as follows:

Fact 2: There is substantial heterogeneity in pass-through across and within industries.

In the following, we focus on several factors to understand the significant differences in pass-through across and within industries.

⁷As there is no information about time spent for each question, we proxy this by the total time spent on the survey (only available for firms who participate online, around 3/4 of all participants). Appendix Figure B.8 shows the distribution of time spent in October 2022.

Figure 2.2 – Cross-sectional variation in pass-through within and between industries




Notes: The stacked bar figure shows the R^2 of pass-through in October 2022 explained by industry fixed effects at 2-digit and 4-digit levels (blue). The remaining variation is within industries (red). Source: ifo Institute’s business survey.

Pass-through and firm’s expectations about the duration of the shock As firms set prices based on their discounted sum of expected marginal costs, beliefs about the persistence of the shock are of first-order importance. To the best of our knowledge, we are the first to investigate this empirically at the firm level by examining the influence of firms’ beliefs about the expected duration of the supply shock for the extent of pass-through. For this analysis, we turn to firms in the manufacturing sector, where the expected duration of supply shortages in months has been elicited in October 2021 (mean: eight months) and June 2022 (mean: ten months). We view supply shortages as an instance of a supply shock since shortages imply higher marginal costs in inputs of production to clear the inputs market. There is large heterogeneity between firms regarding the expected duration, with a standard deviation of five (seven) months in October 2021 (June 2022). To estimate the effect of firms’ expected duration of the shock on pass-through, we regress pass-through, PT_i , of firm i in industry j on firms’ expected duration of the supply shortage, $Exp.Duration_i$, controlling for the range of orders in months, capacity utilization, the qualitative change in revenues, orders, production expectations, log employees, cost changes, and energy intensity (X_i), as well as 4-digit industry fixed effects α_j :⁸

⁸Including industry fixed effects also allows us to control for other factors, such as competition (e.g., Coibion, Gorodnichenko and Kumar, 2018; Afrouzi, 2023), that affect firms’ expectations about macro variables, the duration of supply chain disruptions, and thereby the pass-through.

Figure 2.3 – Expected duration of the shock and pass-through in April 2022



Notes: The figure reports estimates from linear regressions of the firm-level pass-through of cost changes to sales prices in April 2022 on the expected duration of supply shortages in months (elicited in October 2021 and June 2022) in the manufacturing sector. All specifications contain 4-digit industry fixed effects. The respective second specification controls for the range of orders in months, capacity utilization in percent, the qualitative change in revenues, qualitative change in orders, qualitative production expectations, and log employees. The respective third specification additionally controls for the energy intensity (measured as the share of energy costs in terms of revenues) and cost changes between March and May 2021. Appendix Table B.6 displays the respective coefficients and discusses the coefficient of the control variables. Appendix B.1 provides the wording of the specific questions. Standard errors are clustered at the 4-digit industry level. The whiskers reflect the 95% confidence interval. Source: ifo Institute's business survey.

$$PT_i = \alpha_j + \beta Exp.Duration_i + \gamma X_i + \varepsilon_i. \quad (2.1)$$

Figure 2.3 shows a robust positive relationship between the firm's expected duration of supply shortages and pass-through in April 2022. In terms of magnitude, pass-through increases by 6.8 to 8.3 p.p. to a ten-month increase in the expected duration of supply shortages, depending on the model specification. This result underlines the economic significance of the beliefs about the duration of the underlying shock for pass-through. In an alternative specification, we exploit the panel dimension and analyze whether an update of the expected duration at the firm level due to changes in the information set Ω translates into changes in pass-through. Indeed, we find that an increase in the expected duration of supply shortages from October 2021 to June 2022 is significantly associated with a higher pass-through until April 2023, conditional on the reported pass-through in April 2022, see Appendix Table B.7.

As the survey also contains information on the extensive margin of price decisions, we can establish a connection between the expected shock duration and the actual price decisions.⁹ By conducting an event study, we ensure that the expected shock duration is not systematically related to other firm characteristics that shape price decisions. Concretely, the panel dimension allows us to verify common pre-trends prior to the supply shock. We define an indicator for a realized price increase in the respective month t ($\mathbb{1}[P_{i,t} > 0]$) as outcome variable. For the period June 2019 to July 2022, we estimate the following linear probability model for the expected shock duration elicited in October 2021 and June 2022 separately for each month, controlling for the change in revenues, orders, production expectations, and log employees ($X_{i,t}$), as well as 4-digit industry fixed effects α_j :

$$\mathbb{1}[P_{i,t} > 0] = \alpha_{j,t} + \beta_t \text{Exp.Duration}_{i,2021m10/2022m6} + \gamma X_{i,t} + \varepsilon_{i,t}. \quad (2.2)$$

Figure 2.4 plots the estimated β_t over time, which informs us about the importance of expected duration for firms' probability of increasing prices in a given month. Reassuringly, the expected duration of the shock is not significantly related to the probability of increasing prices in 2019 and 2020, the period before the economy was exposed to the historic supply shock. When the supply shock hits in 2021, firms' expectations about the duration of the shock start to predict price increases. A 12-month increase in expected shock length is associated with an approximately ten p.p. higher probability of increasing prices. Hence, the expected shock duration affects price setting at the intensive as well as the extensive margin.

Accordingly, we can summarize our third fact as follows:


Fact 3: Pass-through increases in the expected duration of the shock.

Pass-through and firm's idiosyncratic uncertainty Another potential factor explaining pass-through differences across and within industries is firm-level uncertainty about overall business performance. Uncertainty can have ambiguous effects on firms' pricing decisions. On the one hand, higher uncertainty is associated with experiencing larger shocks, implying more flexible prices in the presence of menu costs Barro (1972).¹⁰ On the other hand, uncertainty could also trigger a "wait-and-see"-strategy, with firms

⁹The ifo business survey elicits the extensive margin of realized price changes compared to the previous month. In the manufacturing sector, the aggregated survey responses show a high correlation with the producer price index from the German Federal Statistical Office ($\rho = 0.78$), see Balleer et al. (2024) for more information and Appendix B.1 for the exact wording of the question.

¹⁰Barro (1972) proposes a theory directly linking the extensive and intensive margin of adjustment to idiosyncratic uncertainty.

Figure 2.4 – Event study: Expected duration of the shock and probability to increase prices




Notes: The figure shows point estimates β in Equation 2.2 and corresponding 90%-confidence bounds of monthly regressions of an indicator for a realized price increase in the respective month on the expected duration of the supply shock in months (blue: elicited in October 2021; red: elicited in June 2022) for the period June 2019 to July 2022. It is controlled for the change in revenues, orders, production expectations, and log employees, as well as 4-digit industry fixed effects. Standard errors clustered at the 4-digit industry level. The vertical dashed line indicates when the supply shock begins to hit the German economy in January 2021. Appendix B.1 lists the wording of the respective survey question. Source: ifo Institute’s business survey.

postponing further price adjustments (Vavra, 2014). Bachmann et al. (2019)—at the extensive margin—and Arndt and Enders (2023)—at the aggregate level—provide evidence for the former channel by relying on measures of volatility over time.

In contrast, we exploit the cross-sectional relation between direct and quantitative measures of uncertainty and pass-through across industries and firms. We measure subjective uncertainty at the firm level: respondents move a continuous slider between zero and 100 depending on the uncertainty about their business development in the next six months. Appendix Figure B.9 shows the average uncertainty over time. Uncertainty spikes from a value of 55 to almost 75 in the spring of 2020 when the pandemic reached Germany. A second spike is in March 2022 after Russia invaded Ukraine. A distinct feature of this measure is that it is not based on nominal values in contrast to other measures in the literature using revenue expectations (e.g., Altig et al., 2022; Bachmann et al., 2024b).

Figure 2.5 presents the relationship between pass-through of cost changes to sales prices until October 2022 and firms’ idiosyncratic uncertainty both at the industry and firm level. In line with Bachmann et al. (2019) and Arndt and Enders (2023), we find that industries and firms with a higher idiosyncratic level of uncertainty pass through cost changes by more. The relation is quantitatively substantial: an increase in uncertainty

Figure 2.5 – Pass-through until October 2022 and uncertainty



Notes: The figure relates the pass-through of cost changes to sales prices until 10/2022 to uncertainty of firms. Uncertainty is elicited on a continuous scale from 0 (low) to 100 (high). The left panel shows the relation as a scatter plot at the industry level (2 digits, at least 20 obs. per industry). The right panel shows the relation as a binned scatter plot at the firm level. It is controlled for the business situation and expectations in both panels. Appendix Tables B.2 and B.3 display the corresponding regression tables. Appendix B.1 lists the wording of the respective survey question. Source: ifo Institute's business survey.

by ten units is associated with a 13 p.p. higher pass-through at the industry level, which remains robust to controlling for the state of the business and expected business situation, see Appendix Table B.2. Appendix Table B.3 confirms the positive relation at the firm level. However, the coefficient becomes significantly smaller, even more so when variation is absorbed at the 2-digit or 4-digit industry levels. This highlights the role of more flexible prices at the industry level for pass-through, indicating the importance of price coordination among competitors.


Accordingly, we can summarize our fourth fact as follows:

Fact 4: Pass-through increases with idiosyncratic uncertainty.

Pass-through and price-setting frequency We next analyze the link between the intensive margin of pass-through and the average frequency of price changes. The panel dimension of the ifo survey allows us to calculate the average frequency of price changes between 2014 and 2019 at the firm level.¹¹ The median and mean are 11% and 19%, respectively, serving as a reduced form measure of price stickiness and indicating the degree

¹¹We take the pre-crisis average price change frequency because it is arguably exogenous to other factors, such as shock size, exposure to the pandemic, or current idiosyncratic shocks, that influence pass-through.

Figure 2.6 – Intensive margin of pass-through until October 2022 and nominal rigidities



Notes: The figure relates the intensive margin (> 0) of pass-through until 10/2022 to the average pre-crisis frequency of price changes. The left panel shows the average frequency (2014-2019) and pass-through until 10/2022 at the industry level (2 digits, ≥ 20 obs. per industry). Blue (red) points indicate the manufacturing (trade) sector. The right panel shows the non-parametric relation between the frequency (2014-2019) and pass-through until 10/2022 at the firm level. The gray shaded area is the 95% confidence interval. Only firms with ≥ 12 obs. in the period 2014-2019 are considered in the figure. Source: ifo Institute's business survey.

of nominal rigidities. Our focus is on the *intensive margin* of pass-through to examine the role of real rigidities and abstract from the mechanical relation of the extensive margin with the frequency of price changes.

Figure 2.6 depicts this relation at the industry level (left panel) and the firm level (right panel).¹² A higher frequency of past price changes is associated with a higher pass-through at both levels.¹³ Intuitively, the more often firms change their prices, the better the price coordination among them. Firms pass through cost increases to a larger extent if competitors are also more likely to be able to adjust, ensuring that the price is close to the competitors. Interestingly, firms in the trade sector—marked in red—have a comparatively low pass-through despite a high frequency of price changes. This indicates the high importance of micro real rigidities in these industries. The relation at the firm level is highly non-linear. While firms that almost never changed their prices between 2014-2019 have an average pass-through of around 45%, firms that changed their prices roughly twice a year in this period have an 8-10 p.p. higher pass-through on average. Firms that have changed their prices even more often exhibit a similar pass-through.

¹²This analysis excludes services as the realization of price changes was not elicited until 07/2018.

¹³Our results complement Gopinath and Itskhoki (2010) finding that the firms with a higher frequency of adjustment exhibit a higher long-run pass-through of exchange rate shocks. We document a similar pattern for a generic supply shock using an entirely different approach.

Accordingly, we can summarize our fifth fact as follows:

Fact 5: Pass-through increases in firms' past price-setting frequency.

2.3 SURVEY EXPERIMENTS

This section turns to the survey experiments to examine the causal link between the nature of the shock and pass-through. Our objective is to quantify how firms' pass-through of a cost shock changes over time depending on the nature of the shock and the underlying economic environment. We use the new insights from the survey experiments to bridge our previous empirical evidence from the field and theoretical predictions from the literature on nominal and real rigidities.

2.3.1 Experimental design

For the survey experiments, we build on the ifo Management-Survey (IMS, 2023). Since 2020, the Management-Survey asks decision-makers in German companies about current economic policy issues and changes in the macroeconomic environment (Sauer, Schasching and Wohlrabe, 2023). The majority of respondents are the owner or CEO of the firm. The ad hoc survey covers approximately 300 firms across all sectors of the economy.¹⁴

We conduct hypothetical vignettes on the link between the nature of the shock and pass-through in January and February 2023. We confront firms with the following hypothetical scenario as well as a graphical representation of the scenario, as described in detail in Appendix B.3:

Suppose that purchase prices for [you and your competitors] in the industry [permanently] increase by 20% due to a global supply shock (see figure). All other factors, such as interest rates and fiscal policy measures, remain unchanged. Assume [good economic conditions with normal capacity utilization] for your company and the overall economy.

¹⁴51% of the firms are in manufacturing, 27% in services, 12% in trade, and the remainder in construction. About 44% of the firms are classified as medium (50-249 employees), 26% as small with less than 50 employees, and 30% as large with more than 500 employees. We exclude decision-makers who do not deal with the firm's price-setting strategy (7.6%).

To what extent would you pass through the cost increase to your customers in this scenario? Indicate the level of pass-through at the respective point in time.

	1m	3m	6m	12m	24m
Pass-through in % after ...	☐	☐	☐	☐	☐

(no pass-through = 0%, full pass-through = 100%, disproportionate pass-through > 100%)

We vary the shock duration (permanent, transitory, uncertain duration) and the exposure to the shock (aggregate, idiosyncratic) within firms. The underlying general economic (low and high aggregate demand) is randomized across firms. Importantly, the setup ensures that (i) the cost increase is exogenous to the firm and (ii) the macroeconomy is held constant to control for general equilibrium effects and policy responses. We ask participants to provide quantitative pass-through estimates *across time horizons* and given different shock durations. Appendix B.3.1 presents the original survey questionnaire, including the hypothetical vignettes, and Appendix B.3.2 offers an English translation.

Our approach allows us to isolate the marginal effect of the nature of the shock, particularly the duration and exposure of the shock, in a *within*-firm setting by keeping all other factors the same. The within-firm design allows us to elicit quantitative pass-through dynamics over time at the firm level. This is crucial for understanding the role of real rigidities by analyzing the number of price adjustments to a permanent shock and directly comparing the pass-through between aggregate and idiosyncratic shocks. On top of that, our approach allows us to discern the importance of the economic environment by comparing responses *across* firms, again keeping all other factors the same. Thus, our approach has the advantage of abstracting from other channels that do not influence the findings.

Firms may choose not to change prices at all due to high price adjustment costs. To address this, we ask firms about their *general* threshold for adjusting prices before presenting the hypothetical scenarios. Appendix Figure B.11 displays the cumulative density function of this threshold, supporting the relevance of sizable menu costs in general but also that in our scenario, menu costs are unlikely to play a role for the initial extensive margin decision. At a 20% cost increase, over 90% of firms typically adjust prices. Furthermore, the results cannot be influenced by macro real rigidities similar to Basu (1995), as all firms in our hypothetical vignette *directly* face a 20% cost increase. In other words, this allows us to completely abstract from any dilution in pass-through from input-output linkages. Note that the shock information is fully available to everyone. Hence, we also abstract

from imperfect information (Phelps, 1970; Mackowiak and Wiederholt, 2009). Our setup has two critical advantages over existing studies: (i) the respective pass-through coefficients are estimated with high precision despite a small sample size, and (ii) by fixing the information set and firms' expectations, we can differentiate between the nature of the shock as well as the economic environment. This provides a clean setup to estimate the causal effect of a cost shock on pass-through dynamics.

2.3.2 Results

Figure 2.7 summarizes the main results of the survey experiment. The upper panels show the time path of the respective average pass-through of firms faced in three shock scenarios: (i) a *permanent* increase in costs by 20% (blue, solid lines) (ii) a *transitory* 20% cost increase in place for 12 months (red, dashed lines), and (iii) a 20% cost increase of *uncertain* duration with a 90% monthly survival rate (green, dotted lines), during high- and low-demand economic environments, respectively. The light-blue lines reflect Calvo-implied counterfactuals, subsequently explained in detail.

Aggregate pass-through dynamics The overall trend observed is a gradual, concave increase in the pass-through at the aggregate level, eventually approaching a level below 100%. As a cross-validation exercise, we compare our experimental estimates with previous studies based on micro price-level data, such as, Gopinath and Itskhoki (2010), Dedola, Kristoffersen and Zullig (2021). Our estimates are in the range but towards the upper end compared to the literature (see Appendix B.2). This could be explained by the fact that our vignettes provide full information about the shock process. Imperfect information—e.g. about the persistence of the shock—can lead to further coordination failure and more gradual pass-through.


Shock duration To what extent do expectations about the nature and duration of the shock matter for firms' pass-through of cost shocks? A comparison of the three shock scenarios—permanent, transitory, and uncertain duration—indicates no significant difference between them in the first three months but a consistently steep increase in pass-through of roughly 40%. This evidence is in line with the presence of *nominal* rigidities leading to a gradual increase in pass-through in the aggregate, as not all firms can change their prices immediately. Similarly, the middle panels in Figure 2.7 show that 63% do not increase prices at all in the first month supporting the importance of nominal rigidities at the beginning. Afterwards, the pass-through increases only slightly for the temporary shock and the shock with uncertain duration. For the permanent scenario, we observe

Figure 2.7 – Pass-through of cost shocks: Main results


(A) High demand environment

(B) Low demand environment


I: Average pass-through of cost shocks across horizons



II: Share of zero/incomplete/complete pass-through (permanent cost shock)



III: Number of changes in pass-through (permanent cost shock)



Notes: The upper panels plot the pass-through of cost shocks across horizons depending on the duration of the shock, differing between permanent, transitory, and uncertain duration shocks, in blue, red, and green colors, respectively. The whiskers reflect the 68% confidence interval (one standard error). The light-blue lines reflect the “Calvo counterfactual” (explained in the text). The middle panels plot the shares of zero pass-through, partial pass-through (between 0 and 100), and complete pass-through (equal or above 100) of cost-push shocks across horizons depending on the economic conditions for the permanent shock scenario. The dashed lines reflect the “Calvo counterfactual.” The lower panels plot histograms of the firm-level number of changes in pass-through in the permanent shock scenario. Source: ifo Institute’s Management-Survey.

further increases in pass-through that are significantly higher than for the other two scenarios highlighting the importance of *expectations* about the nature of the shock for the extent of pass-through. Appendix Figure B.13 confirms that the gradualness and the increasing pass-through in the shock duration are similar across sectors and along the firm size distribution. In line with the findings in the ifo business survey (see Section 2.2.2), pass-through is higher across all scenarios for firms in the manufacturing and construction sectors compared to services and trade. Similarly, Nakamura and Steinsson (2010) find a low magnitude of price changes for services consistent with the presence of higher micro real rigidities. There is no consistent pattern across the size distribution.

Underlying economic conditions We next examine whether the underlying economic environment affects the decision of firms to pass through cost shocks. Comparing pass-through in times of high and low demand (left and right panels, respectively), we generally observe a very similar pattern, even quantitatively.¹⁵ While there is a slightly lower and more gradual pass-through in the low-demand environment, we conclude that the state of the economy is not a main driving force for the characteristics of pass-through. This also demonstrates the generality of our other results.

Nominal and real rigidities Our framework enables us to *directly* assign the observed sluggish pass-through to the presence of nominal and real rigidities. To this end, we take a closer look at pass-through changes at the firm level. While one-time increases after a permanent shock favor nominal rigidities alone,¹⁶ several gradual increases would be consistent with additional micro real rigidities. The majority of firms change their pass-through level more than once in the given scenario, as shown in the lower panels of Figure 2.7.¹⁷ Hence, this is strong evidence for an important role of real rigidities in the pass-through of cost shocks. Note that real rigidities alone would not lead to several adjustments since coordination would work immediately after the cost increase. Only the combination of nominal and real rigidities can explain the observed pattern.

To further quantify the importance of real rigidities, we contrast pass-through under aggregate and idiosyncratic cost shocks. The underlying idea is such: a firm facing an idiosyncratic shock bears in mind that competitors may keep their prices constant, so any


¹⁵One possible reason could be that a lower ability to pass through prices due to low demand is accompanied by the need to raise prices to ensure solvency.

¹⁶Time- (e.g., Calvo) and state-dependent (e.g., menu-cost) models support a one-time adjustment.

¹⁷This is likely a rather conservative estimate due to potential survey fatigue. Further, this may be amplified by our data cleaning procedure, where we fill up/down the pass-through values over time if firms have not entered a value to all five horizons (applies to less than 10% of firms).

price adjustment on the firm’s side distorts relative prices, entailing possible customer resentment and disproportional demand drops. We confront firms with the same hypothetical scenario as before, except that firms face an idiosyncratic shock. Then, we elicit the pass-through level after six months. Figure 2.8 shows that the average pass-through of the idiosyncratic shock is significantly lower vs. the aggregate shock. This pattern arises robustly across the different shock scenarios and for the high- and low-demand environments. Quantitatively, the difference in pass-through of aggregate vs. idiosyncratic shocks is approx. 10-15 p.p. A particularly strong difference is visible for the transitory shock scenario. This might be due to the fact that the relative price difference exists only for a limited period of time, which further reduces the incentive to deviate from competitors. Appendix Figure B.12 shows that this pattern holds across sectors and along the firm size distribution.

Figure 2.8 – Pass-through of aggregate vs. idiosyncratic cost shocks after six months



Notes: The figure shows the pass-through of aggregate (blue) vs. idiosyncratic (red) cost shocks after six months depending on the shock duration, differing between permanent, transitory, and uncertain duration shocks. The whiskers reflect the 68% confidence interval (one standard error). Source: ifo Institute’s Management-Survey.

Connection to the empirical results The survey experiment allows us to rationalize the pass-through pattern in the cross-section and over time presented in Section 2.2. The results from the field show that firms with a longer expected duration have a higher pass-through. While, in reality, the distribution of beliefs about the duration of supply chain disruptions is continuous, the survey experiment maps this to two cases: permanent vs. transitory. Similar to the empirical results, we find a statistically significant increase in pass-through with the expected duration in the survey experiments, confirming the

importance of beliefs about the duration of the shock. Note that belief updating can also affect pass-through over time and explain an increasing pass-through (even in the absence of any real and nominal rigidities). More specifically, firms’ learning over time about the shock duration may lead to an adjustment in pass-through.

Likewise, the survey experiment indicates that nominal and real rigidities play an important role in the pass-through of shocks: about 60% of firms do not react on impact, and firms increase prices only gradually. These results from the experiment can be directly linked to the results from the field: (i) firms that adjust prices infrequently exhibit lower pass-through (Figure 2.6), and (ii) firms plan a granular increase of pass-through over six months (Figure 2.1). In addition, 41% of the participants in the survey experiment listed competition as a limiting factor for price setting and the main source of lack of pass-through in open-ended questions, as shown in Appendix Figure B.5, aligning fully with the survey panel evidence of 40% listing competitive pressure as the main factor.¹⁸

Comparison to Calvo model predictions How well would a canonical Calvo model fit the observed pass-through estimates? To compute the implied pass-through under a Calvo model, we calibrate the adjustment probability θ to match the fraction of zero pass-through on impact ($\theta = 0.37$) in the middle panels of Figure 2.7 and the cumulative pass-through to a permanent shock after 24 months.¹⁹ The predicted share of zero pass-through by the Calvo counterfactual, visualized by the dashed white line in the middle panels, closely follows the empirical share of zero pass-through over time. The light-blue, solid lines in the upper panels of Figure 2.7 reflect the pass-through dynamics under Calvo pricing for the permanent shock scenario. The implied Calvo estimate does not fit the observed average pass-through estimates since firms adjust prices gradually. The Calvo coefficient must be lowered to match the observed pass-through estimates. θ equal to 0.22 (light-blue, dashed lines) best fits the observed pass-through estimates and closely matches the survey experiment. Hence, the Calvo coefficient based on the extensive margin is 68% too high. This “Calvo counterfactual” is a reduced-form way to quantify the role of micro real rigidities.

Implications for the Phillips curve What do our pass-through results tell about inflation-output dynamics and monetary policy transmission? The pass-through esti-

¹⁸The open-ended survey question design has the advantage of not priming respondents to a specific set of answer categories, see Haaland et al. (2024) for a review on open-ended questions.

¹⁹We treat average pass-through of 85% after 24 months as complete and assume firms can substitute the remaining costs. We do not want to overstate incomplete long-run pass-through, as the median is at 100%.

mates can be directly mapped to the slope coefficient of the Phillips curve (PC) building on Auclert et al. (2024). In Appendix B.3.4, we show that the implied PC coefficient based on the nominal rigidities observed in the survey experiment is with 0.22 relatively steep. Once accounting for real rigidities it shrinks to 0.06 to 0.14, depending on the specification. Overall, our estimate is in the range of previous work, as summarized in Appendix Table B.9.

2.4 GENERAL EQUILIBRIUM PRICE-SETTING MODEL

This section introduces a model to rationalize the observed pass-through depending on the persistence and exposure of the shock and estimates the degree of nominal and real rigidities based on our empirical results. We take an off-the-shelf menu cost model à la Nakamura and Steinsson (2010) featuring nominal rigidities (menu costs) and idiosyncratic as well as aggregate shocks and extend it to include micro real rigidities in the form of Kimball demand (see e.g., Klenow and Willis, 2016; Aruoba et al., 2022; Beck and Lein, 2020). This class of models is consistent with the empirical evidence, particularly with firms' inaction to small cost changes and the importance of competition for price setting.

2.4.1 Model setup

The economy consists of different types of agents: households, producers, and a monetary authority.

Households A continuum of households consumes a composite consumption good, Y_t , composed of a continuum of differentiated goods y_t^i and assembled free of charge with a Kimball (1995) demand aggregator: $\int_0^1 G\left(\frac{y_t^i}{Y_t}\right) di = 1$, where G takes the following functional form depending on firm i 's relative demand $\frac{y_t^i}{Y_t}$:

$$G\left(\frac{y_t^i}{Y_t}\right) = 1 + (\bar{\theta} - 1) \exp\left(\frac{1}{\bar{\epsilon}}\right) \bar{\epsilon}^{(\frac{\bar{\theta}}{\bar{\epsilon}} - 1)} \left[\Gamma\left(\frac{\bar{\theta}}{\bar{\epsilon}}, \frac{1}{\bar{\epsilon}}\right) - \Gamma\left(\frac{\bar{\theta}}{\bar{\epsilon}}, \frac{y_t^i \frac{\bar{\epsilon}}{\bar{\theta}}}{Y_t \bar{\epsilon}}\right) \right], \quad (2.3)$$

with Γ reflecting the incomplete gamma function. $\bar{\theta}$ denotes the steady-state value of the demand elasticity, i.e., the elasticity of demand with respect to the price level, and $\bar{\epsilon}$ the super-elasticity, that is the derivative of the price elasticity with respect to the price

level. While the price elasticity of demand is constant (the super-elasticity is zero) in the constant elasticity of substitution (CES) demand case, the super-elasticity is larger than zero in the presence of real rigidities via Kimball demand.

The corresponding demand function for differentiated good i is:²⁰

$$y_t^i = \left[1 - \bar{\epsilon} \ln \left(\frac{p_t^i}{P_t} \right) \right]^{\frac{\bar{\theta}}{\bar{\epsilon}}} Y_t, \quad (2.4)$$

where demand for good i increases with overall demand for goods Y_t and decreases in the relative price $\frac{p_t^i}{P_t}$, the more so the larger the steady-state value of the super-elasticity $\bar{\epsilon}$. The demand also depends on the steady-state value of the demand elasticity $\bar{\theta}$.

In the presence of real rigidities, that is, $\bar{\epsilon} > 0$, the demand function becomes more elastic in the relative price. This relationship is particularly evident when looking at the demand and profit functions, shown in Appendix Figure B.14. The higher the degree of real rigidities, i.e., the larger $\bar{\epsilon}$, the steeper the demand function, and at very high values it becomes almost a step function. Intuitively, a steeper demand function means that consumers are more responsive to relative price changes.

Firms A continuum of firms indexed by i produces differentiated good, y_t^i , using labor and materials as input factors. Resetting prices is costly. Firms face menu costs, which are implemented as a fixed cost proportional to labor costs. Firms produce according to a standard production function depending on firm-specific productivity, A_t^i , and a diminishing returns to scale technology in labor l_t^i and materials m_t^i : $y_t^i = A_t^i (l_t^i)^{1-sm} (m_t^i)^{sm}$, where sm reflects the share of materials. Accordingly, firms maximize the sum of the future discounted per-period profits Π_t^i :

$$\Pi_t^i = p_t^i y_t^i - w_t l_t^i - m_t^i - \chi I_t^i, \quad (2.5)$$

where $p_t^i y_t^i$ denotes the firm's sale revenue, $w_t l_t^i$ the firm's wage bill, m_t^i the cost for materials, χ fixed costs that the firm incurs for adjusting prices, and I_t^i an indicator function that takes the value one if the firm decides to change prices and zero otherwise. The firm's idiosyncratic productivity A_t^i with persistence ρ_A evolves according to:

$$\log A_t^i = \rho_A \log A_{t-1}^i + \varepsilon_t^A, \quad (2.6)$$

²⁰The results are robust to using the alternative demand specification of Dotsey and King (2005).

where ε_t^A reflects an idiosyncratic productivity shock with $\varepsilon^A \sim \mathcal{N}(0, \sigma_A^2)$, denoted hereafter as idiosyncratic shock.

After setting up the firm's problem and specifying the demand functions, we discuss the interaction of micro real rigidities, menu costs, and the persistence of shocks. Intuitively, in the presence of menu costs, firms change prices (and pass through cost increases) only when shocks are sufficiently large or persistent, and firms do not have to face disproportional losses in market share (micro real rigidities). Introducing micro real rigidities increases firms' threshold to change prices (extensive margin), and firms raise prices more gradually (intensive margin) to minimize deviations from competitors' prices.

Monetary authority The monetary authority determines the level of nominal demand, S_t , which equals consumption spending, $S_t = P_t C_t$, in equilibrium and evolves according to an AR(1) process:

$$\log S_t = \mu + \rho_S \log S_{t-1} + \varepsilon_t^S, \quad (2.7)$$

where μ reflects trend growth in nominal demand, and ε_t^S a shock to nominal demand with $\varepsilon_t^S \sim \mathcal{N}(0, \sigma_S^2)$, labeled henceforth aggregate shock. Typically, nominal demand is assumed to follow a random walk, i.e., $\rho_S = 1$. In light of our empirical findings and given the mean-reverting nature of interest rates, we allow ρ_S to be less than 1, implying that aggregate shocks can also be transitory.

2.4.2 Calibration

We calibrate the standard parameters to values in the literature and Euro Area estimates following Beck and Lein (2020), Smets and Wouters (2003) and Nakamura and Steinsson (2010). Appendix Table B.10 presents the calibration for the menu-cost model with micro real rigidities.

Targeted moments The remaining two parameters, which determine the degree of nominal and real rigidities, are estimated based on our empirical results using the method of simulated moments. We target (i) the frequency of price changes and (ii) the level of the impulse response function to the aggregate shock. More specifically, fixed costs for changing prices, χ , are estimated to match the median frequency of price changes at the firm level in the ifo business survey panel (from 2014 and 2019): 11%, corresponding to price changes every nine months and similar to a recent estimate (12%) based on CPI micro data for euro area countries by Gautier et al. (2023). The degree of real rigidities measured by the super-elasticity $\bar{\epsilon}$ is estimated to match the average impulse response

dynamic of pass-through to a permanent cost increase of 20% for horizons from one to 24 months.²¹ Thereby, we implicitly assume that firms' pass-through to cost shocks is similar to nominal demand shocks—in our view, a reasonable assumption, especially since the nominal demand process is understood as a reduced-form way to capture aggregate shocks. The estimation results are shown in Table 2.2:²²

Table 2.2 – Estimation results for parameters

Parameter	Description	Value
$\bar{\epsilon}$	Super-elasticity of demand	6
χ	Fixed cost for changing price	0.0019

Notes: Fixed cost for changing price is expressed as a fraction of steady-state revenue.

Appendix Figure B.15 provides a heatmap of the target function values for a grid of menu costs and micro real rigidities. It clearly shows that a combination of both rigidities is necessary to match the data.

2.4.3 Results

We next shift the focus to studying the implied pass-through dynamics by the estimated degree of micro real rigidities using simulated data. We proceed in three steps: first, we compare the estimates of our simulated impulse response for the aggregate shock to its targeted empirical counterpart; second, we validate our results by comparing the untar-geted impulse response to a transitory shock to its empirical counterpart and contrasting pass-through of aggregate versus idiosyncratic shocks; finally, we study counterfactual pass-through dynamics and implications for the real economy in an environment without real rigidities ($\bar{\epsilon} = 0$). To do so, we simulate data with parameter calibrations according to Tables 2.2 and B.10 and estimate local projections of the price level to an aggregate or idiosyncratic shock.


We start by assessing the fit of the estimated model parameters and compare the pass-through dynamics of a permanent aggregate shock in the survey data versus the model. Figure 2.9 shows that the implied pass-through impulse responses almost perfectly mirror

²¹To ensure monetary neutrality in the long run, we rescale the pass-through values from the hypothetical vignette such that pass-through is 100% after 24 months.

²²See Appendix B.4.2 for details on the estimation.

each other. The results imply that the survey responses can be reconciled for a significant degree of real rigidities ($\bar{\epsilon} = 6$), leading to a sluggish pass-through and corroborating *Fact 1*. Our estimates of the super-elasticity parameter $\bar{\epsilon}$ are in line with the literature; see, e.g., Beck and Lein (2020) for a summary. However, assuming that no real rigidities are at play ($\bar{\epsilon} = 0$), as shown in the counterfactual, leads to an overestimation of pass-through. The average estimated price change frequency is 0.11, matching the median in the survey data.

Figure 2.9 – Data vs model: Pass-through of a permanent aggregate shock (targeted)




Notes: The dark-blue line reflects the pass-through impulse response to an aggregate permanent shock from the survey data. The pass-through is rescaled to 100% after 24 months to ensure monetary neutrality in the long run. The orange line is the model-implied pass-through impulse response to an aggregate permanent shock using the estimated χ (0.0019) and $\bar{\epsilon}$ (6). The light-orange line reflects counterfactual pass-through in an environment with no real rigidities $\bar{\epsilon}=0$.

Next, we turn to the untargeted moments and compare the fit of the estimated model parameters to a transitory shock ($\rho_S=.9$). Figure 2.10 shows that the model does a decent job matching the transitory shock. We observe a much more sluggish and dampened pass-through to a transitory shock in both the survey data and model, validating *Fact 3*. This is due to firms adjusting prices to the expected sum of marginal costs, which is lower in the case of transitory shocks and affects both the intensive and extensive margins of adjustment. As before, a model without real rigidities ($\bar{\epsilon} = 0$) has a hard time matching the observed survey data and overestimates pass-through. Interestingly, the pass-through

dynamics for the permanent and transitory shock are quite similar in the absence of real rigidities. Hence, there is an interaction effect when real rigidities are introduced: transitory shocks have particularly low pass-through in this case.²³

Another metric for evaluating our estimated model's performance is to analyze the pass-through of idiosyncratic shocks in the model. In line with the findings in the survey experiments, the pass-through of an idiosyncratic shock is significantly lower than that of an aggregate shock in the model. However, the difference in the model is even more pronounced compared to our empirical findings.

Figure 2.10 – Data vs model: Pass-through of a transitory aggregate shock (untargeted)



Notes: The dark-blue line reflects the pass-through impulse response to an aggregate transitory shock from the survey data. We compute the average of the transitory scenario (12 months) and the uncertain duration scenario (90% monthly survival rate) from the hypothetical vignettes, see Figure 2.7. The orange line is the model-implied pass-through impulse response to an aggregate transitory shock ($\rho_S = .9$) using the estimated χ (0.0019) and $\bar{\epsilon}$ (6). The light-orange line reflects counterfactual pass-through in an environment with no real rigidities $\bar{\epsilon}=0$.

The model additionally captures the other facts documented in Section 2.2.2. Pass-through increases with the degree of uncertainty (standard deviation of shocks), resembling *Fact 4*. Pass-through also increases in the frequency of price changes (lower menu

²³Note that the model abstracts from the possibility that firms learn about the nature of a shock over time. This is consistent with our hypothetical vignette design, where there is full information about the shock process. In reality, however, firms may learn about the persistence of a shock, leading to additional coordination failures.

costs), as indicated by *Fact 5*. Additionally, the observed heterogeneity in pass-through (*Fact 2*) can be explained by differences in expectations and real and nominal rigidities, as supported by the evidence presented earlier.

Which degree of monetary non-neutrality does the model imply? To answer this question, we study the impulse response of real output to a nominal demand shock. Appendix Figure B.16 shows that real effects based on the model without real rigidities are very short-lived. The introduction of real rigidities makes real effects significantly more persistent. For the permanent and transitory nominal demand shock, real effects die out after 1.5 years and nine months, respectively. However, estimates in the literature suggest that the real effects of monetary policy last three years or more (e.g., Gertler and Karadi, 2015; Romer and Romer, 2004; Christiano, Eichenbaum and Evans, 1999; Jarociński and Karadi, 2020). Hence, additional frictions are necessary to generate a stronger persistence of nominal shocks, in line with our implication for the Phillips curve in the previous section.

2.5 CONCLUSION

We find sluggish pass-through dynamics of cost shocks to customers using a novel firm-level survey approach. The gradual pass-through attributes an important role to nominal and micro real rigidities, limiting pass-through at the firm level. Significant heterogeneity in pass-through across and within narrowly defined industries suggests that additional idiosyncratic factors play a role. In particular, we highlight the importance of firm-specific expectations about the duration of the shock and its interaction with nominal and real rigidities. Additional evidence from hypothetical vignettes confirms these results and allows a causal quantification of frictions and a direct mapping to the slope of the Phillips curve. We then use our empirical estimates to quantify the degree of nominal and real rigidities in a standard price-setting model and find a substantial degree of both.

Our results have implications for the propagation of shocks and the transmission of monetary policy. First, sluggish pass-through implies persistent effects of monetary policy. In fact, New Keynesian models that rely on a high degree of real rigidities to match the persistence of monetary shocks in the data, such as Smets and Wouters (2007), might be better empirically micro-founded than previously thought. Second, our estimates indicate that the Phillips curve becomes much flatter once accounting for the extent of real rigidities we observe in the data, bridging the wide range of previous estimates. The estimated flat Phillips curve also implies a worse trade-off between inflation and output

for monetary policy, making it more costly to reduce inflation. At the same time, our empirical findings indicate that increased uncertainty acts as a counterbalancing factor, bringing the economy towards a more flexible price equilibrium.

We deduct two avenues for future work. First, the hypothetical vignettes on dynamic pass-through could be augmented further; for instance, to learn more about potential non-linearities with respect to the shock size or asymmetries. Second, our evidence on the importance of the belief about the shock persistence could deliver another dimension of friction, leading to gradual pass-through and, in turn, a high degree of monetary non-neutrality. Here, communication about the duration of the shock could significantly affect and amplify aggregate transmission. Implementing imperfect information—and thereby heterogeneous expectations—about shock persistence into a price-setting model would give further intuition into this mechanism, especially in combination with endogenous information acquisition (Gorodnichenko, 2008; Melosi, 2014; Coibion, Gorodnichenko and Kumar, 2018; Afrouzi, 2023).

Chapter 3

Firm Expectations and News: Micro v Macro

This chapter is based on co-authored work with Benjamin Born, Zeno Enders, Gernot Müller, and Knut Niemann, and has been published as a CEPR Discussion Paper. See Born et al. (2024) for the full reference. Regarding the use of generative models, I am fully responsible for the selection, adoption and all results of the generated output. I used DeepL and ChatGPT for spell checking and linguistic refinement of the text.

3.1 INTRODUCTION

In theory, firms act in a forward-looking manner as they decide about production, investment and hiring as well as on prices (Lucas, 1972; Kydland and Prescott, 1982; Woodford, 2003; Mortensen and Pissarides, 2009). And yet, at an empirical level, little is known about how firms form expectations. A fundamental question for both theory and policy is what type of economic information is reflected in firm expectations and to what extent. In this paper, we analyze survey data to assess in detail how firm expectations adjust to news. We leverage earlier work on professional forecasters by Coibion and Gorodnichenko (2015a) and link forecast errors to forecast revisions as a measure of news perceived by firms, but zoom in on expectations about firm-level outcomes—the ultimate concern of firms—and allow them to be affected differently by micro and macro news. These news represent information about firm-specific developments and the aggregate economy, respectively.

Our empirical specification is microfounded in a stylized general-equilibrium model where firm expectations take center stage. Information is noisy and expectations are not necessarily rational: firms may suffer from ‘island illusion,’ meaning they systematically underestimate the importance of aggregate developments for their own performance. The departure from rational expectations is governed by a single parameter that jointly restricts the impact of micro and macro news: firm expectations overreact to micro news and underreact to macro news under island illusion; they do neither under rational expectations.

Our analysis uses data from two firm surveys: The Banca d’Italia’s ‘Survey on Inflation and Growth Expectations’ (SIGE) of Italian firms and the German ‘ifo Business Climate Survey,’ or ifo survey for short. They differ along many dimensions, including how answers are solicited, the time frequency, and even the variables of interest. Yet, we find robustly across all specifications that firm expectations overreact to micro news and underreact to macro news. Also, the extent of overreaction and underreaction co-moves positively both over time and across firms—just like the model would have it assuming there is island illusion. Island illusion also has aggregate implications: We establish this as we embed it in the noisy information model of Lorenzoni (2009) and calibrate it to match the survey evidence.

More in detail, we relate firms’ forecast errors about firm-level outcomes to news as processed by firms in real time, measured as forecast revision across survey rounds. These revisions may reflect both micro and macro news. We isolate the effect of the micro component by controlling for macro news. The SIGE is particularly suited for our anal-

ysis because it solicits firms’ expectations about their own prices but also their inflation expectations, based on which we compile macro news. The ifo survey, in turn, offers longer time series than the SIGE, particularly at the level of individual firms. It also asks firms about their production expectations but does not regularly inquire about their macroeconomic expectations. However, the ifo survey also provides the basis for the ifo *Business Climate Index*, an aggregate indicator of the German business cycle. It is widely watched, including by professional forecasters. And indeed, firms expectations about the macroeconomy have been shown to align well with those of professional forecasters (Link et al., 2023). Hence, to measure macro news, we may resort to the surprise component of the index relative to the corresponding consensus forecast of professional forecasters. Importantly, in both surveys, the responses are confidential, and the weight of individual responses is sufficiently small for strategic motives not to matter (Gemmi and Valchev, 2023).

Our main result emerges robustly across both surveys: Firm expectations overreact to micro news and underreact to macro news. Positive macro news, or new information about the overall economy, tends to lead to positive forecast errors, meaning that actual outcomes exceed expectations. More concretely, if firms revise inflation expectations upwards, their own prices down the road will be even higher than expected. Or, if the current ifo index surprises positively, a firm’s production will likely exceed its expectation. In this sense, firm expectations do not fully account for macro news as it becomes available: they underreact to macro news. Micro news, instead, has a negative effect on the forecast error: an upward revision of expectations tends to be followed by a lower-than-expected realization. Put differently, firm expectations respond too strongly to micro news, they overreact. These patterns are not only a distinct feature of both data sets but also emerge robustly for alternative definitions and across various specifications. Moreover, underreaction and overreaction are persistent over time, suggesting that they are not caused by measurement error that might otherwise partially rationalize overreaction (Juodis and Kucinskas, 2023).

Our baseline specifications pool firm-level observations for, in turn, the SIGE and the ifo survey. For the ifo survey, we also estimate the reaction to news for individual firms, taking advantage of the large number of consecutive observations available for most firms. It turns out that individual firms do not differ much in their overreaction to micro news, but the response to macro news is more dispersed. Most firms underreact, but how strongly depends on firm characteristics. Larger firms, for instance, underreact more strongly. This result may reflect a stronger impact of the macroeconomy on the production—and hence the forecast errors—of larger firms.

We also detect a pattern in the cross-section: Firms that tend to overreact strongly to micro news also tend to underreact strongly to macro news, as predicted by island illusion. Likewise, when we allow estimates to vary over time, we find that in periods when there is a strong underreaction to macro news, the overreaction to micro news is also strong, and vice versa. Lastly, we establish that the variation in the reaction to news across firms systematically correlates with firm-level outcomes. We find, in particular, that a stronger overreaction to micro news is associated with lower profits, and both overreaction to micro news and underreaction to macro news are associated with higher firm-level production and forecast-error volatility—consistent with earlier work showing that firm expectations matter for firm outcomes (Bachmann, Elstner and Sims, 2013; Enders, Hünnekes and Müller, 2022).

Our results are consistent with island illusion that we embed in the canonical noisy-information model of Lorenzoni (2009); first, in a stylized manner, to offer a microfoundation of our empirical specification, and later to explore the aggregate implications of island illusion. In this context, island illusion means that firms *systematically* underestimate the weight of the aggregate component in the signals that help them to learn about the state of the economy—hence, the departure from rational expectations. More broadly, island illusion is an instance of salience, which Taylor and Thompson (1982) define as “the phenomenon that when one’s attention is differentially directed to one portion of the environment rather than to others, the information contained in that portion will receive disproportionate weighing in subsequent judgments” and which has been shown to bias inference in various contexts (Bordalo, Gennaioli and Shleifer, 2022; Bordalo et al., 2023).¹

As we embed island illusion in the full Lorenzoni model, we first run a Monte Carlo exercise: using simulated data to estimate our empirical specification, we obtain overreaction to micro news and underreaction to macro news, just like in the data. In the process, we calibrate the extent of island illusion in the model via an indirect inference procedure. Finally, we use the model to show that island illusion alters aggregate dynamics in response to a productivity shock—aligning them more closely with the time-series evidence put forward in the classic study by Galí (1999).

The rest of the paper is organized as follows. In the remainder of the introduction, we place the paper’s contribution in the context of the literature. Section 3.2 introduces

¹Bianchi, Ludvigson and Ma (2022) use a machine-learning algorithm to estimate time-varying systematic expectational errors and find that survey respondents place too much weight on the private or judgmental component of their forecasts and too little weight on publicly available economic information. Also, at a more general level, direct experience has been found to impact (risk) perceptions more strongly than outcomes experienced by others (Smith et al., 2001; Viscusi and Zeckhauser, 2015).

our empirical framework alongside a model-based microfoundation. Sections 3.3 and 3.4 introduce the SIGE and ifo surveys and present results for price and production expectations, respectively. In Section 3.5, we perform a brief model-based assessment of the aggregate implications of island illusion. The conclusion outlines potential venues for future research.

Related Literature First, at an empirical level, our work relates to the literature which is concerned with macroeconomic expectations of firms; see, for instance, Coibion, Gorodnichenko and Ropele (2020); Andrade et al. (2022); Kumar, Gorodnichenko and Coibion (2023); Savignac et al. (2024), as well as the recent survey article by Candia, Coibion and Gorodnichenko (2023). In contrast, our focus is on firms’ expectations about their own performance. Here, only a limited number of studies have analyzed firm expectations about firm outcomes (see Born et al., 2023). Massenot and Pettinicchi (2018) and Barrero (2022) document ‘over-extrapolation’ for firm expectations based on the ifo survey and for U.S. managers, respectively, meaning that firm-level performance (rather than forecast revisions) predicts forecast errors. Enders, Hünnekes and Müller (2019), in turn, take a macro perspective and document that the response of firm expectations to monetary policy shocks is non-linear in the shock size. Neither of these studies distinguishes between the response to micro and macro news.

Second, our empirical specification extends the framework popularized by Coibion and Gorodnichenko (2015a); see Born, Enders and Müller (2024) for a survey. Coibion and Gorodnichenko (2015a) find for the *consensus* forecast of professional forecasters that forecast revisions predict forecast errors. This conflicts with the full information rational expectations benchmark but not necessarily with rational expectations because the consensus forecast is not available to individual forecasters in real time. Bordalo et al. (2020) show that forecast revisions predict forecast errors also *at the level of individual forecasters*.² As such they reject rational expectations but not rationality *per se*: Predictable forecast errors may emerge because of forecasters’ asymmetric loss function, specific constraints on information processing, or in a learning environment with parameter uncertainty (e.g., Elliott, Komunjer and Timmermann, 2008; Bachmann et al., 2024a; Farmer, Nakamura and Steinsson, 2023; Kohlhas and Robertson, 2024).³

²In fact, the consensus forecast underreacts to news, while individual forecasters overreact to news. See also Angeletos, Huo and Sastry (2021); Kucinskas and Peters (2022); Afrouzi et al. (2023); Broer and Kohlhas (2024) for evidence on the reaction to news of households, professional forecasters, or participants of experiments.

³However, we stress that models that abandon the full information assumption in favor of noisy information still predict that forecast errors should not be predictable at the level of individual forecasters (Bordalo et al., 2020). This includes models of rational inattention (e.g., Mackowiak and Wiederholt,

Lastly, our paper relates to theoretical work that accounts for behavioral aspects in expectation formation.⁴ Models of *level-K thinking*, *cognitive discounting* and *sticky expectations* can rationalize why there is underreaction to current news (e.g., Bouchaud et al., 2019; Farhi and Werning, 2019; García-Schmidt and Woodford, 2019; Carroll et al., 2020; Gabaix, 2020), while *constrained memory* may account for overreaction (Azeredo da Silveira and Woodford, 2019). Ba, Bohren and Imas (2023) show that bounded rationality at various stages of belief formation can lead to both over- and underreaction. Potentially unrepresentative media reporting or, more broadly, *narratives* may also distort the expectation formation process (Shiller, 2017; Chahrour, Nimark and Pitschner, 2021; Andre et al., 2024). Our model of island illusion is conceptually closely related to *diagnostic expectations* and *overconfidence* (Bordalo et al., 2019, 2020; Bianchi, Ilut and Saijo, 2024b; Broer and Kohlhas, 2024). It differs from these approaches in simultaneously accounting for under- and overreactions in a general-equilibrium setting.

3.2 FRAMEWORK

This section presents our empirical framework, placing it in the context of earlier work. It also offers a microfoundation based on a stylized general-equilibrium model where firm expectations are subject to ‘island illusion.’

3.2.1 Baseline specification

Our empirical analysis extends the approach of Coibion and Gorodnichenko (2015a), CG for short, and related studies in two ways. First, we consider firm expectations of firm-level outcomes: either a firm’s production or its price. Second, as we study the impact of news on expectations, we distinguish between different *types* of news: Firm expectations about firm-level outcomes may respond to a) news about firm-specific developments, that is, ‘micro news’ and b) news about relevant macro variables, that is, ‘macro news.’⁵

2009). Kohlhas and Walther (2021) put forward a model of asymmetric attention which rationalizes the observation that forecasts of output growth underreact to *average* forecast revisions (news) but overreact to recent realizations of output growth. They stress that asymmetric attention may arise in a fully rational framework.

⁴Under certain conditions, behavioral models and incomplete information models give rise to equivalent equilibrium effects (Angeletos and Huo, 2021).

⁵Broer and Kohlhas (2024) study systematically how professional forecasters respond to various real-time indicators, including macroeconomic variables. They do not, however, focus on the distinction between micro and macro news, as we do below.

To measure micro news, we rely on firms’ forecast revisions about firm-level outcomes because forecast revisions reflect the reaction to available news (Bordalo et al., 2020). Of course, such revisions may reflect both micro news and macro news. To isolate micro news, we pursue two alternative strategies. First, we explicitly control for macro news in the empirical analysis. Second, we purge forecast revisions of a macro component before the estimation. Specifically, in this case, we measure micro news as the residual of firm-level regressions, which relate forecast revisions to changes in relevant macro variables observable to firms in real time. To measure macro news about inflation and the business cycle, we use firms’ forecast revisions about macro variables (whenever they are available). Alternatively, we rely on the surprise component of relevant macro variables, measured relative to the consensus forecast.⁶ We provide more details on the construction of these variables as we turn to the data in the following sections.

Our empirical analysis is centered around the following specification (baseline):

$$x_{t+h,t}^i - F_t^i(x_{t+h,t}^i) = \beta_0 + \beta_1 \cdot \text{micro news}_t^i + \beta_2 \cdot \text{macro news}_t^i + v_t^i. \quad (3.1)$$

Here, $x_{t+h,t}^i - F_t^i(x_{t+h,t}^i)$ is the forecast error of firm i made at time t regarding firm-level outcome x_{t+h}^i , which is either production or prices, depending on the specification. In what follows, we refer to β_1 as the ‘micro coefficient’ and β_2 as the ‘macro coefficient’: under rational expectations, these coefficients are zero—independently of whether there are information frictions or not—because micro and macro news are part of a firm’s information set.⁷

Specification (3.1) is not simply a test of rational expectations. It also points towards specific alternative models of expectation formation. When positive news ($\text{news}_t > 0$) tends to be followed by positive forecast errors ($\beta > 0$), the forecast revision turns out to be too weak from an ex-post point of view. Hence, there is underreaction to news. Conversely, when positive news is, on average, followed by negative forecast errors ($\beta < 0$), the forecast revision is too strong from an ex-post point of view: There is overreaction to

⁶Surprises are measured relative to a firm’s information set, e.g., in period t , we subtract from the indicator release at the end of period $t - 1$ the corresponding consensus forecast.

⁷Coibion and Gorodnichenko (2015a) estimate a version of Specification (3.1) on the median (consensus) professional forecast for inflation and find a positive coefficient. This result is still consistent with rational expectations: It may simply reflect a failure of the full-information assumption. Intuitively, by estimating the specification on the consensus forecast, the econometrician uses information not available to the individual forecaster. Instead, if Specification (3.1) is estimated on data for individual forecasters, rational expectations imply that forecast errors are not predictable by news_t that forecasters observe in real time—independently of whether there is otherwise full information or not. Bordalo et al. (2020) estimate a version of Specification (3.1) at the forecaster level using a single news measure and find a negative coefficient, rejecting rational expectations.

news (Bordalo et al., 2020). Section 3.2.2 develops a stylized general-equilibrium model, providing a microfoundation to Specification (3.1). It also establishes conditions under which β_1 and β_2 indeed differ from zero.

Before that, we take up a potential complication. In principle, measurement error may mechanically induce a negative correlation between the forecast revision and the forecast error, a concern raised by Juodis and Kucinskas (2023).⁸ To see this, consider the possibility that firms do not report their actual expectations but, for whatever reason, deviate from the true value when reporting their expectations in the survey. Formally, let $\varepsilon_t^{i,rep.}$ denote some i.i.d. measurement error. Reported expectations are then given by $F_t^{i,rep.}(x_{t+h,t}^i) = F_t^i(x_{t+h,t}^i) + \varepsilon_t^{i,rep.}$. The observed forecast error, $x_{t+h,t}^i - F_t^i(x_{t+h,t}^i) - \varepsilon_t^{i,rep.}$, is then by construction *negatively* correlated with the reported forecast revision: $FR_t^{i,rep.} = FR_t^i + \varepsilon_t^{i,rep.} - \varepsilon_{t-1}^{i,rep.}$.

Hence, in principle, measurement error offers an explanation for overreaction. In our empirical analysis, we do not find overreaction in response to macro news, but in response to micro news. To rule out that this result merely reflects measurement error, we consider a number of alternative specifications. First, we consider a dynamic version of Specification (3.1) which relates forecast errors to concurrent and *lagged* news, including p lags for both micro and macro news:

$$x_{t+h,t}^i - F_t^i(x_{t+h,t}^i) = \beta_0 + \sum_{\rho=0}^p (\beta_{1,\rho} \cdot \text{micro news}_{t-\rho}^i + \beta_{2,\rho} \cdot \text{macro news}_{t-\rho}^i) + v_t^i. \quad (3.2)$$

If there is an overreaction not only to concurrent but also to lagged news ($\beta_{1,\rho} < 0$ for $\rho > 0$), it is unlikely to be caused by measurement error: As our derivation above shows, measurement error in lagged news is not correlated with the forecast error made at time t . In principle, if measurement error were positively autocorrelated over some periods, lagged news may still cause the time- t forecast error. However, prolonged periods of forecast revisions in a specific direction are unlikely to reflect measurement error, but rather actual expectations. As a second robustness check, we simply exclude outliers, which are especially prone to measurement error.

3.2.2 Microfoundations

We develop a stylized general-equilibrium model of firm expectations tailored to illustrate how expectations about firm-level outcomes depend on micro and macro news. This

⁸Hence, measurement error may explain some of the overreaction; in case of professional forecasters around one half according to the estimates of Juodis and Kucinskas (2023).

perspective sets the model apart from earlier work by Bordalo et al. (2020) and Broer and Kohlhas (2024). Specifically, we propose a particularly simple departure from rational expectations that allows us to map the model solution into Specification (3.1). This departure is captured by a single parameter, which, in turn, represents the degree of island illusion and imposes joint restrictions on β_1 and β_2 . In what follows, we keep the exposition brief and delegate details to the appendix.

Model outline There is a continuum of islands indexed by $i \in [0, 1]$, each populated by a representative firm. Firms operate under monopolistic competition and manufacture differentiated goods based on island-specific productivity, simultaneously driven by an aggregate component and an island-specific component, as in Lorenzoni (2009). Firms adjust production to meet the demand that arises at pre-set prices. A representative household buys from and supplies labor at the pre-set nominal wage to all firms.

The timing of events is as follows: firms start with expectations about their own output, which will depend on their price relative to the aggregate price index.⁹ Next, firms observe their own productivity but not the underlying components. They also receive a public signal about the aggregate price index that is noisy but correlates with aggregate productivity (macro news). These observations might prompt firms to revise their output forecasts. Once they have set prices, demand materializes and firms produce. A firm's forecast error is actual output minus the level of output expected by the same firm after receiving the news. Below, we evaluate firm expectations and relate them to forecast revisions and news according to Specification (3.1). Time subscripts are omitted for simplicity.

More in detail, production Y^i is linear in productivity A^i and labor L^i : $Y^i = A^i L^i$. Demand, in turn, is given by $Y^i = (P^i/P)^{-\gamma} M/P$, where $\gamma > 1$ is the elasticity of substitution between differentiated goods, P is the aggregate price index, and $M > 0$ is constant nominal demand. Using small letters to denote logs, productivity is the sum of the aggregate component ε and the island-specific component η^i : $a^i = \varepsilon + \eta^i$. Both components are i.i.d. random variables with mean zero, variances σ_ε^2 and σ_η^2 , and the cross-island average of η^i is zero.

Productivity on island i matters not only for production. It is also a signal about the aggregate component that, in equilibrium, is inversely related to the aggregate price level and, hence, determines a firm's relative price, $p^i - p$. Formally, ignoring constants in what

⁹For simplicity, we consider firms' level forecasts of output. Since the earlier realizations are known, this comes without loss of generality.

follows, we have $p = -\varepsilon$. In addition to their own productivity, firms receive a noisy signal of the aggregate price index, that is, macro news: $s = p + e = -\varepsilon + e$. The noise term e has variance σ_e^2 .

Island illusion As we turn to expectation formation, we further simplify the model and assume $\text{Var}(a^i) = \text{Var}(s)$.¹⁰ The parameter ω measures the weight of the idiosyncratic component in firm-level productivity such that $\sigma_\eta^2 = \omega \text{Var}(a^i)$. This implies $\sigma_\varepsilon^2 = (1 - \omega) \text{Var}(a^i)$ and, given our simplifying assumption, $\sigma_e^2 = \omega \text{Var}(s)$. The weight of the aggregate component perceived by firms potentially differs from the actual weight $1 - \omega$ and is given by $1 - \omega\Upsilon$. Under rational expectations: $\Upsilon = 1$. Values of $\Upsilon \geq 1$ capture island illusion, meaning the *perceived weight* (and the perceived value of σ_ε^2) is smaller than the actual weight. As firms correctly observe the sum of both components and its variance $\text{Var}(a^i)$, the perceived weight of the idiosyncratic component, $\omega\Upsilon$, is then larger than its actual value ω . Υ is bounded from above by $1/\omega$. The solution for a firm's expectation of the aggregate price level, that is, its forecast, is then given by (see Appendix C.1.2):

$$F^i(p) = \frac{1 - \omega\Upsilon}{2 - \omega\Upsilon} s - \frac{1 - \omega\Upsilon}{2 - \omega\Upsilon} a^i. \quad (3.3)$$

Intuitively, expectations about the aggregate price level decline in a firm's own productivity a^i and rise in the signal s about the aggregate price level. After a decrease in a^i , firms expect overall technology to have fallen, causing the price level to rise. Similarly, a higher s directly signals an elevated price level. This holds for rational expectations ($\Upsilon = 1$) and in the case of island illusion ($\Upsilon > 1$).

In the latter case, however, firms think of changes in a^i as more idiosyncratic (or island-specific) than they actually are, dampening their impact on the expectations of aggregate prices. For a given signal volatility $\text{Var}(s)$, island illusion also reduces the perceived signal-to-noise ratio of s , as its volatility $\text{Var}(s)$ is known but its informational content, given by σ_ε^2 , is thought to be lower than it actually is. Under island illusion, firms, therefore, react less to both signals than under rational expectations. As a result, the aggregate price index—and, hence, own sales—will, on average, be higher (lower) than expected, given $s > 0$ ($a_i > 0$).

¹⁰In the model, this holds if macro news is published by a forecaster which observes productivity on one island and makes this information public so that firms effectively observe two technology draws: their own and that of one competitor. In Appendix C.1.3, we relax the assumption $\text{Var}(a^i) = \text{Var}(s)$ and establish a (fairly mild) condition for our main result to go through.

Put differently, firms make systematic forecast errors in response to signals. Note, however, that firms do not err when it comes to the variances of s and a^i ; they just misjudge the relative importance of the driving forces.¹¹ Moreover, the implied probability distribution cannot be rejected based on a finite number of draws from the data-generating process—an instance of near-rational expectations (Woodford, 2010).¹²

Formally, the forecast error of firm i is the difference between actual output, y^i , and the firm forecast conditional on productivity and macro news $F^i(y^i|a^i, s)$. The forecast revision is the change in the forecast upon observing these signals relative to the initial expectation $F_0^i(y^i)$, which is based on the commonly known pre-shock value of aggregate technology: $FR^i = F^i(y^i|a^i, s) - F_0^i(y^i)$. Armed with these definitions, we state our main result (derived in Appendix C.1.3):

Proposition 1. *Consider the regression*

$$y^i - F^i(y^i) = \beta_1 FR^i + \beta_2 s + v^i, \quad (3.4)$$

where all terms are defined as introduced above, and v^i represents a potential error term. In the case of island illusion, that is, for $\Upsilon > 1$, we obtain

$$\beta_1 < 0 \quad \text{and} \quad \beta_2 > 0.$$

Furthermore, a higher degree of island illusion (a higher Υ) implies larger absolute values of β_1 and β_2 . For $\Upsilon = 1$, we obtain $\beta_1 = \beta_2 = 0$ (rational expectations).

Equation (3.4) is the counterpart to our baseline Specification (3.1) and Proposition 1 establishes stringent cross-parameter restrictions on a possible deviation from rational expectations. With island illusion, the model predicts a simultaneous overreaction to private signals, as reflected in the forecast revision FR^i net of the public signal s , and an underreaction to public information by individual firms—based on a single parameter that captures the departure from rational expectations.¹³ We obtain Proposition 1 by considering three stages of the information flow—forecast, news/forecast revision, production—in

¹¹Note that island illusion does not lead to unconditionally biased expectations but an underestimation of future output volatility (which depends on the underestimated volatility of ε). Both predictions are consistent with evidence established by Barrero (2022) based on a survey of managers.

¹²In this framework, placing an upper bound on the relative entropy (a measure of the discrepancy) between firms' subjective and objective conditional probabilities implies an upper bound for Υ , given the remaining parameter values.

¹³The same holds true if FR^i and s are included separately.

a single period, which simplifies the notation. In Section 3.5, we incorporate island illusion into a multi-period business cycle model. In this context, the information flow spans across periods, and hence island illusion affects how news impacts the intertemporal forecast error, consistent with Specification (3.1).

Model extension: From production to price expectations A relation isomorphic to Proposition 1 emerges for firms’ price expectations once we consider a simple model extension. In particular, assume that the aggregate component of technology follows a random walk while the idiosyncratic component is white noise. Furthermore, assume that firms operate for multiple periods. In this case, firms expect to re-adjust prices once they can (at the beginning of next period) to align them with the long-run technology level. Given that they overestimate (underestimate) the distance between their own and the aggregate technology in response to $a^i(s)$, they also overestimate (underestimate) their future price adjustment. In this way, the model can also rationalize over- and underreaction of price expectations to micro and macro news. Appendix C.1.4 provides more details.

3.3 PRICE EXPECTATIONS

In this section, we estimate our empirical specification on data for price expectations elicited in a survey of Italian firms. The data is uniquely suited for this purpose since it permits constructing direct measures of micro news and macro news, as processed by firms, and relating them to forecast errors about their own prices. We briefly describe the data before presenting the results.

3.3.1 Survey of Inflation and Growth Expectations

The quarterly *Survey of Inflation and Growth Expectations* (SIGE, 2023) among Italian firms is run by the Banca d’Italia. It was launched as a panel in 1999 and polls firms in the industrial sector, excluding construction, and non-financial services. The average response rate in the survey is 45 percent.¹⁴ Our sample starts in 2002, when questions about realized changes in firms’ own prices were introduced in the survey, and runs up to the end of 2022. Firm responses are generally in the form of growth rates, and the

¹⁴For the comparable Survey of Business Uncertainty in the United States, the response rate is around one-third (Altig et al., 2022).

Table 3.1 – Survey of Inflation and Growth Expectations

(A) Definition of variables		
Label	Name	Wording
$F_t^i(\pi_{t+4,t}^i)$	Exp. change in own prices	For the next 12 months, what do you expect will be the average change in your firm's prices?
$F_t^i(\pi_{t+4,t})$	Exp. inflation (12 months ahead)	In July, consumer price inflation, measured by the 12-month change in the harmonized index of consumer prices, was 8.4 percent in Italy and 8.9 percent in the euro area. ^a What do you think it will be in Italy in September 2023?
$\pi_{t,t-4}^i$	Realized change in own prices	In the last 12 months, what has been the average change in your firm's prices?

(B) Construction of micro news, macro news, and forecast errors		
Component	Definition	Alternative/robustness
Micro news	$F_t^i(\pi_{t+4,t}^i) - F_{t-1}^i(\pi_{t-1+4,t-1}^i)$	net of $\gamma_i \Delta \pi_t$
Macro news	$F_t^i(\pi_{t+4,t}) - F_{t-1}^i(\pi_{t-1+4,t-1})$	$\pi_{t-1} - \text{mean forecast}_{t-1}$
Forecast error	$\pi_{t+4,t}^i - F_t^i(\pi_{t+4,t}^i)$	–

Notes: Panel (A): definition of variables and questions in SIGE September 2022 questionnaire. All questions elicit growth rates in percentage points; t denotes quarter of the survey. Since 2009/2010, SIGE features additional questions about firms' inflation expectations six months and two years ahead, which we use to construct fixed-time macro news as robustness exercise. ^aPrevious sentence is information treatment that firms generally receive in survey. See Appendix Table C.1 for the complete set of questions that we use. Panel (B): definition of variables featured in baseline specification. γ_i indicates firm-specific factor loadings used to purge micro news of their potential macro component, i.e., the most recent change in inflation $\Delta \pi_t$.

wording of the relevant questions is documented in Panel (A) of Table 3.1.¹⁵ As a unique feature, firms are asked about expected changes in both their own prices and aggregate inflation, making it possible to measure micro and macro news straightforwardly.

As summarized in Panel (B) of Table 3.1, to construct forecast errors in quarter t about a firm's price change, we subtract from the realized change reported in period $t + 4$ its four-quarters-ahead expectation, as reported in period t . We measure macro news as the first difference of firms' four-quarter-ahead expectations for aggregate inflation across survey rounds. Next, we compute the forecast revision of a firm for its own price: the change in its four-quarters-ahead expectation relative to what it reported in the previous quarter. This is not a direct measure for micro news. As discussed above, forecast revisions may

¹⁵For further details on the SIGE, see Banca d'Italia (2022) and Grasso and Ropele (2018).

reflect not only micro but also macro news.¹⁶ Hence, we control for macro news as we estimate the impact of forecast revisions on forecast errors. In this way, we can isolate the effect of micro news.

For both forecast revisions and macro news, the overlap in forecasting periods is three out of four quarters, which we consider large enough for changes to reflect actual news (rather than changes in the forecast horizon). Still, in our robustness analysis, we consider alternative measures for which there is full overlap. Lastly, we emphasize that we measure forecast revisions, macro news, and forecast errors (all in percentage points) at the level of individual firms based on firms' survey responses.


In our baseline, we winsorize these survey-based variables at the top and bottom one percent to account for potential measurement error. To obtain a consistent sample for the entire analysis, we consider only observations for firms in the survey for at least 20 quarters.¹⁷ When asked about their expectations for aggregate inflation, firms are generally provided with Italy's most recent inflation rate. Since 2012, a randomly selected subset of firms has not received this information. Our baseline sample features only observations for firms that receive the inflation update, but we also assess how much this information treatment matters for our results.

Figure 3.1 presents descriptive statistics. Panel (A) shows how the number of time-series observations varies across firms. More than 530 firms are in the survey for 20 quarters or more, with some firms in the survey for more than 75 quarters. We plot the distribution of forecast errors in Panel (B). It turns out that expectation errors about firms' own prices are generally (for more than 75% of firms) not significantly different from and centered around zero. This is a general pattern in firm surveys (Born et al., 2023). Panel (C) shows a sizeable dispersion of firms' forecast revisions about their own prices, notably in times of macroeconomic turmoil. Finally, Panel (D) illustrates the evolution of average macro news over the sample period. We observe that firms' revisions of inflation expectations are generally moderate, at least during the low and stable inflation environment through 2020. There are more sizeable revisions as inflation takes off in the post-pandemic period.

¹⁶And indeed, we find across a number of specifications that macro news predicts forecast revisions; see the Appendix Table C.2.

¹⁷Appendix Figure C.1 shows that the resulting sample is representative of all firms in the survey based on their geographic location, number of employees, export orientation, and sector.

Figure 3.1 – Basic statistics for SIGE



Notes: Panel (A): distribution of (quarterly) observations per firm. Panel (B): histogram of firm-level average forecast errors for own prices (after winsorizing). Color indicates if estimates significantly differ from zero at the five percent level (light green) or not (dark green). Bins with less than three firms are not shown due to confidentiality constraints. Panel (C): cross-sectional standard deviation of own-price-forecast revision. Panel (D): macro news as average forecast revision for the one-year-ahead inflation rate.

3.3.2 Results for price expectations

We are finally in the position to estimate Specification (3.1). For this purpose, we pool observations across firms, accounting for firm-fixed effects. Table 3.2 reports the results. Consider Column (1) first: Both types of news are statistically significant predictors of forecast errors. Hence, we reject rational expectations for firms, consistent with the result of Bordalo et al. (2020) for professional forecasters. On top of that, we find that the type of news is key for *how* expectations fail to meet the rational expectations benchmark: While positive micro news predicts negative forecast errors, positive macro news predicts

Table 3.2 – Response of price expectations to news

	Firms' forecast errors about their own prices				
	(1)	(2)	(3)	(4)	(5)
Micro news					
Forecast rev. for $\pi_{t+4,t}^i$	−0.423*** (0.018)		−0.422*** (0.031)		−0.426*** (0.018)
Forecast rev. for $\pi_{t+4,t}^i$ net of $\gamma_i \Delta \pi_t$		−0.416*** (0.018)		−0.416*** (0.031)	
Forecast rev. for $\pi_{t+4,t}^i \times$ no treatment					−0.064 (0.046)
Macro news					
Forecast rev. for $\pi_{t+4,t}$	0.275*** (0.050)	0.246*** (0.050)			0.289*** (0.050)
Inflation surprise			3.059*** (0.404)	2.984*** (0.407)	
Forecast rev. for $\pi_{t+4,t} \times$ no treatment					0.305*** (0.111)
Observations	18,284	18,284	6,640	6,640	21,982
R ²	0.071	0.070	0.074	0.073	0.079
Within R ²	0.080	0.079	0.089	0.089	0.086

Notes: Results are based on Specification (3.1), relating firms' forecast errors about their own prices to micro news and macro news. Micro news is measured either by forecast revisions while controlling for macro news, or by forecast revisions purged of changes in aggregate inflation. Macro news is measured by firms' forecast revisions about aggregate inflation or by the surprise component of inflation relative to (mean) professional forecast from Consensus Economics. The top and bottom one percent of firms' forecast revisions for their own prices and aggregate inflation and their forecast errors about their own prices are winsorized. Column (5): results for news interacted with a treatment indicator equal to 1 when the survey question does not feature the recent inflation rate and equal to 0 when the survey question does feature the recent inflation rate (baseline). The treatment indicator is also included in the regression. Sample: 2002 (2013 for inflation surprises) to 2022. Firm-fixed effects are always included, standard errors clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

positive forecast errors. This implies that firms overreact to micro news but underreact to macro news, consistent with our model of island illusion, which restricts the response of expectations to both types of news jointly.

The estimates are quantitatively important. Consider again Column (1). On average, an *upward* revision of price expectations due to micro news by one percentage point (standard deviation) is followed by a forecast error in the other direction: eventually, prices *fall short* of expectations by 0.423 percentage points (0.27 standard deviations of the forecast error). After an *upward* revision of a firm's expectation of aggregate inflation by one percentage point (standard deviation), a firm's own price *exceeds* expectations by 0.275 percentage points (0.05 standard deviations).¹⁸

The measures for news in Column (1) are based entirely on survey responses, reflecting the change in the four-quarter ahead expectations for own prices and aggregate inflation, respectively. In Column (2), we show results obtained once we purge micro news of the

¹⁸The standard deviation of micro (macro) news is 2.456 (0.724); that of the forecast error is 3.791.

most recent change in inflation while allowing for firm-specific loadings, γ_i . In Columns (3) and (4), we vary the measure of macro news, replacing the forecast revision with the inflation surprise relative to the average professional forecast taken from Consensus Economics.¹⁹ Our main result is robust across all specifications: the data never rejects the model’s predictions under island illusion.²⁰

The SIGE is the basis for an influential study on firms’ inflation expectations by Coibion, Gorodnichenko and Ropele (2020). For this purpose, the authors exploit that since 2012, the survey no longer provides information about the most recent inflation rate to a randomly selected group of firms (about one-third of all firms).²¹ Recall that our baseline estimates are based on the sample of firms that receive the information treatment. To assess the extent to which the treatment matters, we interact news in our regression with an indicator variable that assumes a value of one if a firm receives no information about aggregate inflation in a given forecasting period. Column (5) shows the result: Whether firms are treated or not does not change their response to micro news but to macro news. Firms that receive no inflation update tend to underreact more strongly to macro news. This is intuitive: absent treatment, news is less salient and forecasts are revised even less.

We find that our results are robust along a number of further dimensions and provide results in the appendices. Specifically, we limit the sample to observations up until 2019 to check whether results are driven by the pandemic and the following surge of inflation. And indeed, in this case the overreaction to macro news is still significant, but weaker, see Column (1) of Appendix Table C.3. Hence, greater inflation variability after the pandemic likely helps to better identify overreaction to macro news. Results are also robust as we consider an alternative measure of macro news. In the baseline, macro news is the change in the four-quarters-ahead inflation forecast from one quarter to the next and, as such, does not pertain to the same target date. To address this issue, we exploit the fact that

¹⁹Our specification makes sure that the surprise is in a firm’s information set for the relevant survey round. Consider, e.g., the Summer of 2022. On June 13, Consensus Economics polled professional forecasters about their inflation expectations in the second quarter and published results on June 16. The Banca d’Italia published the inflation rate on July 8. The difference between the realized value and the average professional forecast is the measure of macro news that we use for 2022Q3. In this case, the SIGE was conducted between August 25 and September 15. Recall that firms are explicitly informed about the current inflation rate when asked about their inflation expectations (baseline).

²⁰The coefficients on macro news in Columns (3) and (4) are an order of magnitude larger than in Columns (1) and (2). This reflects the fact that inflation surprises are much smaller.

²¹See Table 3.1 for the exact wording. The information treatment generates exogenous variation in inflation expectations and is shown to influence not only pricing decisions but also credit demand, employment, and capital. Ropele, Gorodnichenko and Coibion (2024) utilize the survey to study how the dispersion of beliefs about future inflation influences the misallocation of resources. Rosolia (2021) investigates how the information provision impacts the level of inflation expectations, as well as pricing and labor demand.


the SIGE surveys firms’ annual inflation expectations not only four quarters ahead but also two and eight quarters ahead.²² Hence, to measure news based on a fixed target date, we simply compute the difference between the current short-term forecast and the lagged longer-term forecast. We do so for two alternative macro-news measures by comparing, first, current two-quarters-ahead forecasts of inflation to four-quarters-ahead inflation forecasts reported two quarters ago and, second, by comparing current four-quarters-ahead forecasts to eight-quarters-ahead forecasts reported four quarters ago. For both of these fixed-target measures of macro news, we obtain results that are very similar to the baseline; see again Appendix Table C.3.

We further verify that our results also hold for numerous alternative specifications. Namely, if we no longer winsorize survey variables or if we trim the top and bottom 1% of survey variables, results are very similar to the baseline; see Panels (A) and (B) of Appendix Table C.4. We also consider alternative ways to purge forecast revisions of their macro component. First, we employ time-fixed effects and isolate micro news as the deviation from the cross-sectional mean even though, strictly speaking, this mean is not observable by firms in real time. Second, we purge micro news of recent changes in inflation with homogeneous loading γ (instead of γ_i). Results are similar to the baseline, as Panel (C) shows. Lastly, we show that our results also hold when including only micro and macro news separately instead of including them jointly: Panel (D).

In principle, as discussed in Section 3.2 above, measurement error may induce a negative contemporaneous correlation between the forecast error and the forecast revision—a risk that we reduce by winsorizing observations in the baseline. Now, we go one step further and estimate the dynamic Specification (3.2) and show results in Figure 3.2. The solid line represents the point estimates for various news lags, and the shaded area indicates 95-percent confidence intervals. We find there is overreaction not only to contemporaneous micro news but also to lagged news: Panel (A). We are thus confident that measurement error is not driving our results. We also find that the underreaction to macro news is fairly persistent, albeit to a lesser extent than the overreaction to micro news: Panel (B).

²²Questions about two-quarter-ahead inflation expectations were introduced in 2010. Four-quarters-ahead expectations are available for the entire sample period. Eight-quarters-ahead expectations were introduced in 2009. See Appendix C.2 for details.

Figure 3.2 – Response of price expectations to current and past news



Notes: Results based on Specification (3.2), where we extend the baseline specification to include four lags of micro and macro news, that is, firms' forecast revisions (first differences) for their own prices (micro news) and aggregate inflation (macro news). Panel (A) shows point estimates and confidence intervals for current and lagged micro news. Panel (B) shows point estimates and confidence intervals for macro news. Confidence intervals are at the 95-percent level in both panels. Firm-fixed effects are included and standard errors are clustered at the firm level.

3.4 PRODUCTION EXPECTATIONS

We complement the evidence from the SIGE with evidence from the ifo survey of German firms (IBS-IND, 2023). We do so with a focus on production expectations rather than price expectations. In addition to robustifying our results along this dimension, the ifo survey features long time series for individual firms, thus allowing us to probe into the heterogeneity of the expectation-formation process across firms.

3.4.1 ifo Business Climate Survey

The *ifo Business Climate Survey*, or ifo survey for short, is maintained by the ifo institute in Munich and a mostly qualitative, monthly survey representative of the German economy (Hiersemenzel, Sauer and Wohlrabe, 2022). It has been conducted since 1949 and its design has since been adopted by surveys worldwide (Becker and Wohlrabe, 2008; Born et al., 2023). By way of compensation, participating firms receive sectoral and aggregate survey results, a feature we rely on below. The average monthly response rate is 82 percent, and the sample attrition is moderate (Enders, Hünnekes and Müller, 2022). Our sample runs from April 2004 to December 2019.²³ Given our focus on production expectations, we only use data from the manufacturing sector, and given that we are

²³The ifo survey was launched in 1949, and the underlying micro data is available from 1980 onward. Some aggregate statistics based on it were first used by Theil (1955). As discussed below, our sample is

particularly interested in estimates for individual firms, we restrict our sample to those firms that have been in the survey for at least 30 months.²⁴ In any given month, this leaves us with at least 1,000 responses, and often considerably more.

The ifo survey features questions about expected and realized production, where firms can report either an increase, no change, or a decrease; see Panel (A) of Table 3.3 for details. We follow the approach of Bachmann, Elstner and Sims (2013) to compute forecast errors based on these qualitative responses. We aggregate the monthly changes in realized production over three months as reported ex-post and compare the result with the ex-ante expectation for the same period; see Panel (B) of Table 3.3. As before, our basis for measuring micro news is the forecast revision of firms. Formally, given the qualitative responses in the ifo survey, we define the forecast revision of firm i in month t as the sign of the first difference in production expectations. It is equal to 0 when there is no change in expectations, equal to +1 for an upward revision (for example, from no change in $t - 1$ to an increase in t), and equal to -1 for a downward revision (for example, from no change in $t - 1$ to decrease in t). As mentioned above, we assume the forecast revision reflects news because the overlap in the revision is two months.²⁵ Likewise, we need to control for macro news to isolate the effect of micro news that is reflected in the forecast revision.²⁶

Yet, in contrast to the SIGE, the ifo survey does not solicitate macroeconomic expectations regularly. We, therefore, measure macro news by the surprise component of the ifo index, an aggregate index compiled based on the ifo survey and a widely watched indicator of the German business cycle (Carstensen et al., 2020; Lehmann, 2023). We compute the surprise component by subtracting the consensus forecast for the index, compiled by Bloomberg, from the actual index once it is released.²⁷ Importantly, we require macro news to be in

restricted by data we rely on to construct macro news. The individual filling a firm’s questionnaire is a member of the senior management; 85 percent are CEOs or department heads (Sauer and Wohlrabe, 2019).

²⁴Given the ifo survey runs at monthly frequency, we use a higher threshold for the minimum number of observations than for the SIGE (minimum: 20 quarters). Also, for inclusion, we require firm responses to exhibit some time-series variation in their expectations and expectation errors.

²⁵We verify that the forecast revisions are informative by relating the average forecast revisions over time to German manufacturing production growth, see Figure C.3: The average forecast revision is a leading indicator for changes in manufacturing production. This was especially visible during the Great Recession and in 2018/2019, when the manufacturing sector cooled considerably.

²⁶As with the SIGE, we find that macro news impacts forecast revisions significantly, see Table C.6 for a range of specifications that interact macro news with several indicators.

²⁷Bloomberg surveys professional forecasters who can submit and update their forecasts of macro indicators, for example, GDP and employment, but also the ifo index, up until its release. For the ifo index, the consensus forecast is based on about 40 forecasters; it has been available since April 2004. This constrains the beginning of our sample period.

Table 3.3 – ifo Business Climate Survey

(A) Definition of variables		
Label	Name	Wording
$F_t^i(y_{t+3,t}^i)$	Exp. change in own production	Plans and expectations for the next 3 months: Our production is expected to be: [1] increasing [0] not changing [-1] decreasing
$y_{t,t-1}^i$	Realized change in own production	Review – tendencies in [t-1]: Compared to [t-2], our production: [1] increased [0] stayed about the same [-1] decreased


(B) Construction of micro news, macro news, and forecast errors		
Component	Definition	Alternative/robustness
Micro news	$\text{sign}\{F_t^i(y_{t+3,t}^i) - F_{t-1}^i(y_{t-1+3,t-1}^i)\}$	net of $\gamma_i \Delta \Gamma_t$
Macro news	no macro expectations solicited	ifo index _{t-1} – median forecast _{t-1}
Forecast error	$\begin{cases} 0 & \text{if } \text{sign}(y_{t+3,t}^i) = F_t^i(y_{t+3,t}^i) \\ \frac{1}{3} (y_{t+3,t}^i - F_t^i(y_{t+3,t}^i)) & \text{else} \end{cases}$	–

Notes: Panel (A): definition of variables and questions of Business Climate Survey. All questions are qualitative (increase, decrease or no change); t denotes month of the survey. See Table C.5 for the complete set of questions that we use. Panel (B): definition of variables featured in baseline specification. γ_i indicates firm-specific factor loadings used to purge micro news of their potential macro component Γ_t (see Footnote 29). Definition of forecast error is from Bachmann, Elstner and Sims (2013).

the forecasters’ information set, given Specification (3.1). This is likely to be the case for the surprises of the ifo index for several reasons. As stressed above, participating firms receive the latest release of the ifo index as a token of appreciation. In addition, media attention to the index is high. The ifo index is ranked among Bloomberg’s ‘12 Global Economic Indicators to Watch,’ and news outlets regularly report on the realized value as well as on the professional forecasts. In addition, Link et al. (2023) show that in the ifo survey, firms’ expectations about the macroeconomy align well with those of professional forecasters.

Figure 3.3 presents basic statistics for the ifo survey and is organized in the same way as Figure 3.1 above. Panel (A) shows the distribution of firms sorted according to the number of months a firm is in the sample. Here, the median is around 90 months, and 25 percent of firms are in the survey for over 130 months. These numbers are considerably higher than in the case of the SIGE and allow us to zoom in on the firm-level heterogeneity below. Panel (B) shows the distribution of forecast errors. Just like in the case of the SIGE, they are generally well-behaved: More than 75 percent of firm-level average forecast errors are not significantly different from zero. In Panel (C) we display the cross-sectional dispersion of forecast revisions and how it fluctuates over time. It is largest during the Great Recession, the Euro debt crisis, and towards the end of our sample. Panel (D) shows the time series for macro news.

Figure 3.3 – Basic statistics for the ifo survey



Notes: Panel (A): distribution of monthly firm observations. Panel (B): histogram of firm-level average forecast errors for production. Color indicates if estimates significantly differ from zero at the five percent level (light green) or not (dark green). Panel (C): cross-sectional standard deviation of forecast revisions. The grey line depicts the standard deviation of forecast revision at the monthly level, and the black line depicts the six-month rolling average. Panel (D): macro news over time, defined as the surprise in the ifo index compared to median professional forecasts, see Table 3.3b.

3.4.2 Results for production expectations

Armed with these definitions, we revisit our baseline Specification (3.1) using firms' production expectations. We first pool, as before, observations across time and firms while allowing for firm-fixed effects. Table 3.4 reports the results. Column (1) is based on the specification that simultaneously features forecast revisions and macro news. We find that production expectations, too, overreact to micro news and underreact to macro news. While this pattern is consistent with island illusion, its robustness is remarkable because

Table 3.4 – Response of production expectations to news

	Firms' forecast errors about their own production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for $y_{t+3,t}^i$	-0.191*** (0.001)			
Forecast Revision for $y_{t+3,t}^i$ net of $\gamma_i \Gamma_t$		-0.209*** (0.001)	-0.208*** (0.001)	
Macro News				
Surprise component of the ifo index	0.022*** (0.0007)	0.022*** (0.0007)		0.021*** (0.0007)
Observations	302,737	302,737	302,737	302,737
R ²	0.16260	0.15806	0.15313	0.08967
Within R ²	0.08471	0.07974	0.07435	0.00498

Notes: Results based on Specification (3.1); observations are pooled across firms. The table shows results for the qualitative production expectations (3-month horizon) for German firms surveyed monthly in the ifo survey. Macro news is the surprise component of the ifo index. Column (1): micro news measured by forecast revisions (while controlling for macro news). Columns (2) and (3): micro news measured by forecast revisions net of real-time observable macro indicators, Γ_t , with firm-specific loading γ_i (see Table 3.3). Column (4): macro news only. All specifications include firm-fixed effects and standard errors clustered at the firm level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

it emerges as we estimate Specification (3.1) on data sets which differ not only—among other things—in terms of country, data frequency, and survey design but, importantly, also in terms of firm-level variables: price expectations vs. production expectations.²⁸

The remaining columns of the table further validate the robustness of our results: The micro coefficient remains negative and highly significant when we purge the forecast revision of the impact of real-time macro indicators, Γ_t , with firm-specific loadings γ_i (second column).²⁹ In fact, estimates hardly differ from the baseline. In what follows, we, therefore, always measure micro news by the forecast revision net of the macro factors. Note further that when we drop macro news from the regression, the result for the impact of micro news also remains virtually unchanged (third column). The macro coefficient remains positive and significant when including only macro news in the regression (fourth column).

²⁸The ifo survey also features price expectations. We verify that in this case, too, our main result obtains: see Appendix Table C.9 in the appendix.


²⁹Since the ifo survey does not ask firms about their aggregate inflation expectations, we rely on a set of macroeconomic indicators to proxy the relevant information set of firms: real-time monthly changes in German industrial production, the CPI, manufacturing orders, the stock market index DAX, as well as month-fixed effects to control for potential seasonality.

Note that the magnitude of the coefficients in Table 3.4 is again quantitatively relevant—although not as straightforward to assess due to the qualitative nature of the forecast revisions. Consider, for example, Column (2). A one-standard-deviation increase in micro news (0.46) predicts a forecast error of -0.1 (-0.209×0.46 ; that is, 0.27 standard deviations of the forecast error). A one-standard-deviation increase in macro news (1.2) leads to a forecast error of 0.03 (0.022×1.2 ; 0.08 standard deviations of the forecast error). Hence, the effects on forecast errors are not negligible, and the micro-news effect is about three times stronger than that of macro news. In terms of standard deviations, the magnitude is very close to our findings for the SIGE.

We further verify that results are robust across a range of alternative specifications. We show first that forecast errors are also significantly affected by past news in Appendix Figure C.5 and report results of additional robustness tests in Appendix Table C.7. This includes, first, estimates based on a subsample of observations restricted to firms that expect no production change. In this way, we ensure that the results are not mechanically biased by the qualitative revision scale. Next, in addition, we also set small forecast errors (i.e., $\pm 1/3$) to zero. Second, we estimate a specification in which we again set small forecast errors—potentially driven by measurement error—to zero and, in a further addition, do so only for firms that expect ‘no change’ in production. Third, we use an ordered logit rather than OLS for the estimation. Fourth, we consider alternative ways to measure macro news. Specifically, we purge firms’ forecast revision through time-fixed and time-sector-fixed effects. Lastly, we vary the measure for macro news and replace the surprise component in the ifo index, in turn, with the surprise component in manufacturing orders, the change in the ifo index, the average forecast revision, the average forecast revision per sector, and the change in the stock market index. We find robust overreaction to micro news and underreaction to macro news across all specifications.

The ifo also elicits a quantitative measure of firm expectations, which pertains to a firm’s expected business situation over the next six months, and answers are provided in a range from 0 (rather less favorable) to 100 (rather favorable). Correspondingly, the survey also asks about the current business situation, with possible answers ranging from 0 (bad) to 100 (good); see Appendix Table C.5. We use these answers to compile a quantitative measure of forecast errors but note that in this case, data is only available for a subset of firms and since September 2005. Micro and macro news are measured in the exact same way as above, except that micro news is measured quantitatively in terms of revisions in business expectations instead of production expectations. Here, we also find overreaction to micro and underreaction to macro news; see Appendix Table C.8.

Figure 3.4 – Response to news over time




Notes: Estimates based on five-year rolling windows. Black lines represent point estimates; grey areas correspond to 95 percent confidence intervals.

Finally, because the ifo panel is sufficiently large, we can also investigate whether the patterns in the data are robust over time. For this purpose, we follow Coibion and Gorodnichenko (2015a) and estimate the baseline specification on five-year rolling windows. Figure 3.4 shows how the estimated response coefficients evolve over time both for micro news (left) and for macro news (right). We observe, first, that firms overreact to micro news and underreact to macro news over the entire sample. Second, there is a co-movement of overreaction and underreaction. Such a pattern is consistent with Proposition 1: If the extent of island illusion varies over time, overreaction and underreaction will co-move in sync. Third, for macro news, the variation over time appears to be substantial in economic terms: the underreaction is about three times as large during the Great Recession compared to non-recession periods. Taken at face value, this may appear surprising because one might expect firms to pay more attention to the aggregate economy in times of crisis (see also, Flynn and Sastry, 2024). Yet, as we document below, an increased underreaction may reflect a stronger impact of macro variables on firm outcomes without a sufficiently large increase in attention. Consistent with this notion, we find a weaker underreaction to macro news after the Great Recession: The crisis may have increased the attention firms pay to the macroeconomy, resulting in less underreaction once the macro environment stabilized.

3.4.3 Firm-level variation

The ifo survey features sufficiently long time series for individual firms so that we can confidently estimate the reaction to news firm by firm. Specifically, we re-estimate Spec-

Figure 3.5 – Firm-by-firm estimates



Notes: Distribution of estimates for individual firms on Specification (3.1) for production expectations of German firms in the ifo survey (three-month horizon, qualitative data). Grey area represents insignificant estimates, light green area represents estimates significant at the 10% level, dark green area indicates significance at the 5% level.


ification (3.1) for each of the 3,215 firms in our sample and show the distribution of the estimates for β_1 and β_2 in Figure 3.5. There is a clear pattern: Panel (A) shows that the mass of the estimates for β_1 is concentrated to the left of zero, with most estimates being significantly smaller than zero (dark green bars). Specifically, for the subset of significant estimates, β_1 is negative for all firms. Hence, overreaction to micro news is a pertinent feature in firms' expectation formation. Instead, the estimates for β_2 are centered to the right of zero; see Panel (B). In this case, estimates are not always significantly different from zero (grey bars), but when we consider significant estimates only, the macro coefficient is positive for 92 percent of firms.³⁰

Given the firm-level variation, we investigate whether a systematic relationship exists between the estimated coefficient in the cross-section of firms. To this end, we show a binned scatterplot of the micro and macro news coefficients in Figure 3.6: There is a negative association, in particular for the set of firms with significant overreaction to micro news and underreaction to macro news, indicated by the dark triangles. Recall from Proposition 1 that such a negative association is precisely what we would expect if island illusion varies across firms.

We conclude this section with a brief analysis of potential drivers of the variation of firm-level estimates apparent in Figure 3.5. Specifically, we re-run the estimation on pooled samples but add interaction terms that control for specific dimensions of heterogeneity. Table 3.5 reports results. The micro coefficient is robustly negative in the cross-section

³⁰Repeating this exercise for the SIGE, where the number of firms is substantially smaller and the time series per firm are shorter, yields similar results, see Appendix Figure C.2.

Figure 3.6 – Relation between macro and micro coefficients at the firm level



Notes: The figure displays two binned scatter plots (15 bins) between firm-level micro-news coefficients and macro-news coefficients. The grey points display the binned scatter for the set of firms with negative micro-news coefficients and positive macro-news coefficients ($\rho = -0.09$). The green triangles display the binned scatter based on the set of firms with significantly negative micro-news coefficients and significantly positive macro-news coefficients ($\rho = -0.35$).

and does not differ systematically along many dimensions. We run a Wald test and find that across different levels of firm age (Panel 3), time in the survey (Panel 4), and the self-reported importance of the business cycle (Panel 4), there are no significant differences. The overreaction significantly decreases with firm size (Panel 2), but the differences in terms of magnitude are small. This is consistent with the evidence in Panel (a) of Figure 3.5, which shows that the estimates for β_1 cluster in a fairly tight range. Turning to the response to macro news, we find the underreaction to macro news is strictly and statistically significantly increasing across employee quartiles (Panel 2). The underreaction of the largest firms is twice as strong as that of the smallest firms. This result may reflect a stronger impact of the macroeconomy on the production—and hence the forecast errors—of larger firms. Regarding firm age, there is no statistical difference in the response to macro news between young and old firms. When comparing the underreaction of firms that recently joined the survey (within six months) to firms with longer tenure, we find evidence for “learning through survey” (Kim and Binder, 2023). The underreaction among longer-tenured firms is about one-third smaller than for firms that

Table 3.5 – Accounting for firm-level heterogeneity

Interaction	N	Micro News			Macro News		
		$\hat{\beta}_i$	$SE(\hat{\beta}_i)$	W	$\hat{\beta}_i$	$SE(\hat{\beta}_i)$	W
(1) News Overall (see Table 3.4)	302,737	−0.209***	0.001		0.022***	0.001	
(2) News	302,737			0.001			0.000
× 1. Quartile by employees		−0.216***	0.003		0.013***	0.002	
× 2. Quartile by employees		−0.211***	0.002		0.019***	0.001	
× 3. Quartile by employees		−0.210***	0.002		0.022***	0.001	
× 4. Quartile by employees		−0.203***	0.002		0.026***	0.001	
(3) News	162,776			0.554			0.408
× Firm age < 20 years		−0.205***	0.005		0.019***	0.003	
× Firm age ≥ 20 years		−0.208***	0.002		0.021***	0.001	
(4) News	302,737			0.919			0.045
× Time in survey < half a year		−0.210***	0.010		0.033***	0.006	
× Time in survey ≥ half a year		−0.209***	0.001		0.021***	0.001	
(5) News	129,053			0.25			0.038
× Low business-cycle exposure		−0.203***	0.003		0.016***	0.002	
× Med. business-cycle exposure		−0.209***	0.002		0.021***	0.001	
× High business-cycle exposure		−0.208***	0.003		0.022***	0.002	

Notes: All regressions include micro and macro news with interaction terms and firm-fixed effects. Standard errors are clustered at the firm level. N is the number of observations, $\hat{\beta}_i$ is the point estimate and $SE(\hat{\beta}_i)$ its standard error. Column W reports the p-value for the null that the news coefficients are jointly the same. We run the Wald test separately for each type of news. For (quartiles of) the number of employees, we rely on annual questions in the ifo survey. For firm age, we rely on a one-time question about the year the firm was founded. To compute the firm age, we subtract the year of response from the year of foundation. For business-cycle exposure, we rely on a one-time question, where firms rank the importance of general economic developments in Germany for their business on a five-point scale from very important [1] to unimportant [5]. Business-cycle exposure is high when the response is very important [1], medium when the response is important [2], and low otherwise [3-5]. *** p<0.01, ** p<0.05, * p<0.1.

recently joined the survey, and the difference is statistically significant. For the exposure to the business cycle, we distinguish between firms that rank the business cycle as very important, important, or less important to them. Here, in line with the effect of firm size, a high business-cycle exposure comes with a significantly larger underreaction.

3.4.4 Real effects

Expectations matter for firm decisions and outcomes, as a recent study confirms for the ifo survey (Enders, Hünnekes and Müller, 2022). Against this background, we investigate whether overreaction and underreaction to news are associated with measures of firm performance. Specifically, in the cross-section of firms, we relate the estimated response coefficients to profits, production volatility, and forecast error volatility. We restrict the sample to firms that overreact to micro and underreact to macro news.

Table 3.6 – Over- and underreaction to news and real activity

	mean _i (profits _{it})		sd _i (production _{it})		sd _i (error _{it})	
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	0.224 (0.177)		0.383*** (0.011)		0.226*** (0.007)	
Reaction to micro news ($\beta_1 < 0$)	1.70** (0.782)	1.79** (0.756)	-0.371*** (0.046)	-0.360*** (0.046)	-0.318*** (0.028)	-0.312*** (0.028)
Reaction to macro news ($\beta_2 > 0$)	-0.673 (1.79)	-1.10 (1.78)	1.63*** (0.097)	1.61*** (0.097)	1.31*** (0.062)	1.30*** (0.062)
Observations	1,691	1,691	2,227	2,227	2,227	2,227
R ²	0.003	0.051	0.146	0.162	0.230	0.252
Within R ²		0.004		0.143		0.228
Sector FE		✓		✓		✓
Size FE		✓		✓		✓

Notes: Estimates from linear regressions of average profits, Columns (1)–(2), production dispersion of firms, Columns (3)–(4), and forecast-error dispersion, Columns (5)–(6), on the firm-by-firm estimates of the micro and macro news coefficients. The sample is restricted to firms that overreact to micro news and underreact to macro news. Size-fixed effects refer to firm-size quartiles based on the number of employees. Standard errors are clustered at the firm level. *** p<0.01, ** p<0.05, * p<0.1.

Since 2009, the ifo survey asks firms twice a year about profits in the current year, in May and September.³¹ For each firm, we calculate the average profit and regress them on the micro and macro news coefficients estimated in Section 3.4.2. In addition, we absorb sector- and size-fixed effects. Table 3.6, Columns (1) and (2), shows that a stronger overreaction to micro news is associated with significantly lower average profits, while a stronger underreaction to macro news is not significantly related to the average profits. Quantitatively, a one standard deviation increase in the overreaction to micro news leads to an average reduction in profits by about 0.14 percentage points.

Second, we calculate the standard deviation of realized production changes as a proxy for firm-level production volatility. We then follow the procedure above and regress it on the estimated response coefficients to micro and macro news, obtained in Section 3.4.2. Columns (3) and (4) in Table 3.6 show a tight relation between production volatility and the over- and underreaction to news at the firm level. An increase in the overreaction to micro news is associated with higher volatility. While the point estimate is larger for micro news than for macro news, a one standard deviation increase in the estimated coefficient is associated with a somewhat stronger increase of output volatility in the case of macro news. Projecting these cross-sectional estimates to the macro level implies higher micro-

³¹We build our measure of firm profits based on the September wave, subtracting from the response the annual cross-sectional average of firm profits to ensure that the results are not confounded by heterogeneity over time (in a recession, profits are lower and underreaction stronger, see Section 3.4.3).

level volatility in the presence of over- and underreaction, potentially explaining the high observed idiosyncratic volatility of firm outcome variables (Bachmann, Elstner and Sims, 2013; Bloom, 2009; Bloom et al., 2018).

Lastly, we do the same for the standard deviation of qualitative forecast errors as a proxy for the accuracy of firm expectations. Columns (5) and (6) in Table 3.6 display the results. Again, the estimates indicate a tight (negative) relation between the accuracy of forecasts and the over- and underreaction to news at the firm level.

3.5 AGGREGATE IMPLICATIONS

The firm-level evidence points to a significant departure from rational expectations, consistent with the notion of island illusion. In what follows, we perform a first assessment as to whether island illusion matters for macroeconomic dynamics.³² For this purpose, we embed it in the workhorse model of noisy information by Lorenzoni (2009). The original version of the model relaxes the assumption of full information, assuming that firms and households do not observe shocks directly. Yet it still maintains rational expectations because the estimates of the shocks are updated rationally given the signals that households and firms receive. This setup is particularly suited to introduce island illusion: We may simply assume, just like in the stylized version of the model introduced in Section 3.2.2 above, that firms' subjective weights in the Kalman filter are no longer equal to the true shock variances but are distorted by island illusion.³³

In contrast to Section 3.2.2, we now employ the full model and, to tie our hands, stick to the original calibration of Lorenzoni (2009). Allowing for island illusion then amounts to introducing one additional model parameter only: Υ . To pin down this parameter, we rely on an indirect inference procedure: We simulate the model for alternative values of Υ and keep track of individual firms and their expectations. Based on the simulated time

³²Recent work has explored the aggregate implications of *diagnostic expectations*. They may give rise to boom-bust cycles and can amplify shocks within business cycle models (Bianchi, Ilut and Saijo, 2024a; L'Huillier, Singh and Yoo, 2024). Yet, as discussed above, they do not account for the simultaneous under- and overreaction to news that we find to be a pervasive feature of firm-level data.

³³As discussed in Section 3.2.2, the information flow in the Lorenzoni model spans multiple periods. In particular, the last period's forecasts are revised today as news materializes, changing, in turn, expectations about future prices.

Table 3.7 – Overreaction and underreaction to news: Monte Carlo evidence

	Firms' forecast errors about their own prices	
	$\Upsilon = 1$	$\Upsilon = 1.003$
Micro News		
Forecast Revision for $\pi_{t+1,t}^i$	0.002 (0.002)	-0.416*** (0.008)
Macro News		
Public signal s_t	-0.008 (0.006)	1.67*** (0.052)
Repetitions	1,000	1,000

Notes: Results based on Specification (3.1). The table reports results based on simulated data of expectations of firms' own prices. Macro news is the public signal s_t . Column (1) without island illusion, Column (2) with island illusion. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

series, we estimate Specification (3.1) using firms' forecast errors and forecast revisions for their own prices.³⁴ We pin down Υ as we target our estimate for overreaction to micro news, β_1 , reported in Table 3.2 above.


We report results in Table 3.7. The left column shows estimates for the rational expectations version of the model (no island illusion). In this case, there is neither significant over- nor underreaction to micro or macro news, in line with Proposition 1. The right column shows the result as we target the actual estimate for β_1 . In this case, Υ takes a value of 1.003.³⁵ The table also reports the implied estimate for β_2 . Even as we do not target it, we obtain a significantly positive estimate. This implies underreaction to macro news, just like in the data and predicted by Proposition 1.

Finally, to get a first idea of how island illusion might impinge on macroeconomic dynamics, we revisit the transmission of an aggregate TFP shock in the Lorenzoni model. Figure 3.7 shows the results, contrasting the impulse responses to a permanent TFP shock under FIRE (black solid lines) with the noisy information case of the original version of the model (blue dashed line) and with island illusion (red dotted lines). We focus on the responses of output (left panel) and inflation (right panel). Under full information rational expectations (FIRE), output jumps immediately to its new level while inflation remains flat since the shock is (known to be) permanent. In case information is noisy but expectations are rational ($\Upsilon = 1$, as in Lorenzoni, 2009), households and firms realize only over time that the technology improvement is permanent. The output adjustment is,

³⁴To be precise, we simulate the model for 1,000 periods with an additional 2,000 burn-in periods. We define micro and macro news as in Section 3.2.2 above, that is, as the forecast revision and the public signal. We repeat this procedure 1,000 times for each draw of Υ and compute the average estimate and confidence intervals based on the distribution of β_1 and β_2 .

³⁵A key factor that determines Υ is the ratio of the variances of idiosyncratic and aggregate shocks. Again, here we tie our hands by sticking with Lorenzoni's original calibration.

Figure 3.7 – How island illusion alters macro dynamics: Model simulation



Notes: Impulse response functions (IRFs) to one-standard-deviation TFP shock in model of Lorenzoni (2009) under alternative assumptions on expectation formation. Black solid lines: IRFs under FIRE; blue dashed lines: IRFs w/ noisy information and rational expectations ($\Upsilon = 1$), as in Lorenzoni (2009); red dotted lines: IRFs w/ noisy information and island illusion ($\Upsilon = 1.003$). Vertical axis measures percentage deviations from steady state, horizontal axis measures time in quarters.

hence, gradual and expected demand is relatively low compared to productivity, causing inflation to drop. As prices converge towards the long-run price level, inflation slightly overshoots.

Under island illusion, adjustment dynamics differ considerably. Compared to the rational expectation case, we observe that the adjustment of output is more drawn out; later output overshoots its new long-run level. The inflation response is generally much more persistent. Because of island illusion, as discussed in Section 3.2.2 above, firms attribute a larger share of the improvement in technology to idiosyncratic developments—they fail to see the permanent and common nature of the shock. Consequently, and more so than under rational expectations, they overestimate their competitiveness, expect too high demand, and set prices too high. As they slowly learn about the true nature of the shock, they reduce prices gradually towards the long-run equilibrium. Monetary policy responds by keeping interest rates lower. This, in turn, explains why output overshoots after some time. Note that island illusion alters the adjustment dynamics in a way that makes the model prediction more closely aligned with the time-series evidence in Galí (1999), who finds, in particular, that inflation rebounds only gradually after a technology shock.

3.6 CONCLUSION

How do firms adjust their expectations to news? As we address this question, we make a material and a conceptual contribution. First, our material contribution concerns a

potentially important aspect of firm behavior. We find a simultaneous overreaction and underreaction of firm expectations to news, which we interpret—against the background of our model—as an instance of salience: Firms are excessively preoccupied with one portion of the environment at the expense of another. This is likely to matter for firm outcomes, and we present some evidence for island illusion to have real effects. Investigating this issue further and in more detail appears to be a fruitful avenue for future research. The same holds for macroeconomic dynamics: once we embed island illusion in the canonical model of noisy information, its predictions become better aligned with time-series evidence regarding the aggregate effects of TFP shocks.

Second, conceptually, by emphasizing the distinction between different types of news, we move beyond existing work on expectation formation that links forecast errors to news. The distinction turns out to be essential for understanding firm expectations, and we conjecture that it matters beyond the context of firms. For instance, the distinction between micro and macro news may also be relevant for investor expectations or for workers where related patterns have been identified (e.g., Daniel, Hirshleifer and Subrahmanyam, 1998; Jäger et al., 2024). Hence, we expect—an instance of overreaction, perhaps—this conjecture to be the subject of further research, too.

Chapter 4

The Devil is in the Tail: Macroeconomic Tail Risk Expectations of Firms

This chapter is single authored. Regarding the use of generative models, I am fully responsible for the selection, adoption and all results of the generated output. I used DeepL and ChatGPT for spell checking and linguistic refinement of the text.

4.1 INTRODUCTION

The global economy has recently faced significant macroeconomic tail events—such as the Great Recession, the COVID-19 pandemic, and the war in Ukraine. Yet, little is known about firms’ beliefs about occurrence of these events in the future. Since realized macro tail events are rare, firms have limited information not only about the likelihood of such events, but also about the impact of such events on the firms’ business, adding an additional layer of uncertainty.

At the same time, expectations matter for firms’ investment. Indeed, firms’ subjective expectations and uncertainty—corresponding to the first and second moments of their forecast distributions—are important drivers of investment decisions; see, e.g., Gennaioli, Ma and Shleifer (2016) and Kumar, Gorodnichenko and Coibion (2023), respectively. While macro tail risk is a crucial determinant of asset prices and the business cycle in many quantitative models (e.g., Gabaix, 2012; Wachter, 2013; Gourio, 2012), it remains an open question to what extent the beliefs of firm decision-makers regarding macro tail events influence their investment behavior.

This paper studies the nature, role, and consequences of macroeconomic tail risk beliefs of firms. I elicit the subjective probability of four macro tail events multiple times in a large panel of firms: a drop in GDP of at least 5%, a financial crisis similar in magnitude to the Great Recession, a pandemic, and a geopolitical escalation. On average, firms believe that these macro tail events are fairly likely, but beliefs vary widely across firms and fluctuate significantly over time. Changes in these subjective probabilities co-move strongly with changes in firms’ expectations and uncertainty about their own business. Turning to exposure, both structured and open-ended survey questions reveal that managers are highly uncertain about the exposure of their firms to these macro tail events.

Using a threefold empirical approach based on novel data that exploits (i) firm-level observational variation, (ii) experimental variation, and (iii) time-series variation, I examine how changes in macro tail risk beliefs affect firm decisions, focusing on investment as the key outcome variable. Across these different methodologies, a consistent pattern emerges: a higher subjective probability of a macro tail event leads to substantially lower investment. Crucially, at least half of this effect persists after conditioning on subjective first and second moments. The effect is particularly driven by firms that expect to be heavily exposed to macro tail events in worst-case scenarios.

To rationalize the findings and analyze stabilization policies, I explore the role of macro tail risk beliefs in a quantitative heterogeneous firm model. The key mechanism driving

investment decisions beyond first and second moments is ambiguity regarding firm-specific exposure to macro tail events. Both fiscal and monetary policy can stabilize investment after an increase in macro tail risk. I find that fiscal policy in the form of generous loss carrybacks is particularly effective in a world with ambiguous exposure to macro tail events.

I shed light on the nature of macro tail risk expectations of firms using the representative German ifo Business Survey, a well-established large firm panel with a high response rate, typically answered by the firm’s CEO or owner (Born et al., 2023). In the survey, I introduce a set of questions regarding the subjective probability of macro tail events—the first layer of macro tail risk. I focus on four specific macro tail events, as described above, asked at three different points in time. Firms consider these events to be fairly likely. For example, in October 2022, the median firm believes that a large decline in GDP ($\geq 5\%$) occurs with a 50% probability within the next five years, corresponding to a 3.4% quarterly arrival rate. This number aligns with a frequentist approach that considers the development of German GDP growth over the last 20 years. During this period, two significant GDP drops occurred within four five-year windows: the Great Recession in 2008/2009 and the COVID-19 pandemic in 2020.

Macro tail risk beliefs are highly dispersed across firms. Continuing with the example, the interquartile range (IQR) spans from a subjective probability of 20% to a subjective probability of 80%. This substantial cross-sectional dispersion contrasts with firms’ point forecasts of GDP growth and the unemployment rate, which are more similar and closely aligned with those of professional forecasters (Link et al., 2023). For the financial crisis scenario, the median subjective probability is slightly lower and the IQR more narrow. Both subjective probabilities of the median firm decreased by 20 and 30 percentage points, respectively, by October 2023. The pandemic and geopolitical escalation scenarios are perceived to be less likely, but still show substantial dispersion. An increase in the subjective probability of macro tail events is closely associated with lower business expectations and higher business uncertainty, both in the cross-section and within a given firm over time. This is consistent with the notion that macro tail risk shifts firms’ subjective expectations and uncertainty.

A second layer of macro tail risk for firms is how severely they will be affected if a macro tail event materializes. Therefore, I also elicit firms’ beliefs about their exposure to macro tail events with both structured and open-ended survey questions. In the structured survey question, I elicit firms’ beliefs about the change in their real sales in the event of a 5% drop in GDP. I ask them about the change in their real sales in the best-case, on average, and in the worst-case. Beliefs about the average exposure scenario range

from a 30% decrease in sales at the 90th percentile to no change in sales at the 10th percentile. Even conditional on average exposure, there is large heterogeneity in expected best- and worst-case exposures. I find that stronger past exposure to realized macro tail events, such as the Great Recession and COVID-19, significantly predicts a more negative expected worst-case exposure to future macro tail events, even after conditioning on expected average exposure. This finding is in line with the hypothesis by Kozlowski, Veldkamp and Venkateswaran (2020) of scarring effects after strong economic downturns. For the median firm, there is a 20 percentage point range between the best- and worst-case scenarios, indicating high uncertainty about exposure to macro tail events. This range increases with the level of expected average-case exposure.

I also design an open-ended survey question that allows firms to freely express their thoughts on exposure to macro tail events. The survey responses are manually coded into categories according to a codebook. Three main narratives emerge that corroborate the substantial uncertainty in exposure identified in the structured question: (i) firms directly express significant difficulties in forecasting exposure, (ii) when discussing beliefs about exposure in terms of sales, they tend to express it in ranges or max/min scenarios rather than average exposure, (iii) firms often refer to past macro tail events, underscoring the limited set of experiences available to predict exposure.

To study the role of macro tail risk beliefs for investment, I link these beliefs to firm decisions. To fix ideas, I present a stylized firm choice model, which later forms the core of the quantitative model. In the model, consistent with the design of the survey questions, macro tail risk is characterized by two layers of uncertainty for firms. The first layer is the likelihood of occurrence of a macro tail event. The second layer is firm-specific exposure to macro tail events. Motivated by the empirical evidence, I assume incomplete knowledge about the firm-specific exposure to a macro tail event: firms have scenarios in mind regarding the extent to which they might be exposed in this case. However, they are not able to assign probabilities to the respective scenarios, making the exposure ambiguous—also known as Knightian uncertainty (Knight, 1921; Keynes, 1921).¹ If firms are ambiguity averse, i.e. they prefer to know the likelihood of events, they act *as if* they are relatively pessimistic about their exposure to macro tail events. Specifically, they choose the capital stock for the next period *as if* their worst-case exposure will be realized in the macro tail event, see Ilut and Schneider (2022) for a review on ambiguity. Hence, an increase in the subjective probability of a macro tail event results in a significant decline in investment beyond the first-moment effect, even when firms are risk neutral.

¹Bachmann et al. (2020) document that Knightian uncertainty is pervasive among firms.

I employ three complementary empirical approaches to examine the role of macro tail risk expectations for firm decisions, focusing on investment. The first approach (“micro approach”) exploits firm-level variation in macro tail risk beliefs to estimate the effect on firm investment. As outcome variable, I construct an indicator for a planned decrease in investment in the next year, elicited one month after the macro tail risk beliefs.² I find that a one standard deviation increase in the subjective probability of a drop in GDP of at least 5% predicts a 7.5 percentage points higher probability of reducing investment in the next year. When accounting for firms’ subjective first and second moments, specifically adjustments in forecasts and uncertainty about their own business, at least half of this relationship persists. To rule out the influence of any fixed firm characteristics on the results, I next use only within-firm variation. While this naturally reduces available variation and increases noise, I still find a sizable and significant negative effect of macro tail risk expectations on investment. The negative coefficient on investment is particularly pronounced for firms with substantial expected worst-case exposure to a macro tail event, consistent with the theoretical predictions outlined above. The results hold for a battery of robustness checks, including different measures of subjective first and second moments.

In the second approach (“experimental approach”), I use experimental variation to establish causality and to test the proposed mechanism on ambiguity about firm exposure, as explained above. I conduct a survey experiment, where I confront firms with different hypothetical macroeconomic outlooks in a vignette using the ifo Management Survey. The outlooks have the same mean and variance, but varying higher-order moments such as skewness, resulting in different macro tail risk. Firms decide in which outlook they would invest more. Consistent with the results of the first approach, firms are significantly less likely to invest when there is substantial macro tail risk. A main advantage of the experimental approach is that the first layer of macro tail risk expectations—the probability of a future macro tail event realization—is fixed by the outlooks and, therefore, can be separated from the role of firms’ beliefs about exposure, which constitutes the second layer of macro tail risk expectations. I relate the firms’ investment decisions in the vignette to their expected exposure scenarios in the case of a macro tail event. As predicted by the model, strongly negative worst-case exposure is highly associated with investing less in the outlook incorporating substantial macro tail risk, controlling for average and best-case exposure.

The third approach (“macro approach”) exploits time-series variation in macro tail risk to increase external validity. It combines the idea that firms’ macro tail risk beliefs are formed

²To study how the survey responses translate into actual decisions, I relate the survey responses on investment to changes in the capital stock in balance sheet data. A planned decrease in investment reported in the survey implies a 5% decline in the capital stock in the following year.

through news consumption (Carroll, 2003; Chahrour, Nimark and Pitschner, 2021) with the fact that large language models have demonstrated an increasing ability to behave like humans (Aher, Arriaga and Kalai, 2023; Bybee, 2023; Horton, 2023). Specifically, I feed articles from Germany’s largest business newspaper, *Handelsblatt*, between 1986 and 2023 into OpenAI’s GPT-3.5 and ask the model to forecast macro tail risk based on each article. To disentangle macro tail risk—defined as a sharp decline in GDP ($\geq 5\%$)—from other moments of macro news, I also construct indices for average macro forecasts and upside risk macro expectations. The upside risk index isolates changes in macro tail risk from simultaneous movements in upside risk, thus capturing the second moment. Using local projections, I find that an interquartile range increase in the macro tail risk index, conditional on average and upside macro risk, foreshadows a persistent decline in aggregate investment of about 7%.

To rationalize the empirically documented negative effect of macro tail risk beliefs on investment, explore the aggregate consequences, and draw policy implications, I propose a quantitative heterogeneous firm model. In the model, firms choose their optimal capital stock exposed to stochastic aggregate and idiosyncratic productivity as well as capital adjustment costs. As introduced in the stylized model, firms differ in their expected exposure scenarios in the presence of a macro tail event, which corresponds to low aggregate productivity. I calibrate the model using standard parameters and my empirical estimates. I then simulate an increase in macro tail risk expectations without an actual realization of the macro tail event: news of higher macro tail risk strongly dampens investment. Aggregating across firms informs about macroeconomic dynamics: higher macro tail risk generates a recession, where the distribution of firm investment growth is left skewed with a fat left tail due to the reaction to different worst-case exposures embedded in the policy functions.

Next, I study the extent to which monetary and fiscal policy can stimulate investment after an increase in macro tail risk, comparing a world with and without ambiguous exposure to the macro tail event. I conduct counterfactual analyses involving credible policy rule changes in response to a realized macro tail event, which act as automatic stabilizers (Dolls, Fuest and Peichl, 2012; Edgerton, 2010). Counteracting monetary policy in the form of a interest rate cut in a macro tail event, motivated by the Taylor rule, significantly stimulates investment in both worlds, with and without ambiguous macro tail risk exposure. Regarding fiscal policy, I consider a rule-based switch to symmetric corporate taxation—similar to generous loss carrybacks that allow firms to offset current losses against past profits to obtain tax refunds—in macro tail events. The counterfactual analysis indicates that this policy mitigates the negative investment response to an increase in macro tail risk beliefs by 20–25% in a world with ambiguous exposure, while having

little effect in a world without ambiguity. Intuitively, this counter-cyclical fiscal policy is particularly effective, as firms expecting significant losses in their worst-case exposure scenario to the macro tail event benefit most from the switch to symmetric taxation.

Related literature I contribute to five strands of the literature. First, the paper adds to the literature examining macro tail risk beliefs. The approach of directly eliciting macro tail risk beliefs was pioneered by Shiller (2000), who elicits stock market crash expectations. Subsequent surveys have also captured other macro tail event beliefs of investors (Giglio et al., 2021; Goetzmann, Kim and Shiller, 2024) or households (Dietrich, Müller and Schoenle, 2023; Beutel and Stockerl, 2023).³ This study is the first to ask firms about their beliefs regarding a range of macro tail events and their expected exposure in a large panel. To complement this, I also generate a time series of macro tail risk beliefs based on newspaper articles following Bybee (2023). This connects to the extensive literature that creates various types of macro indices based on newspaper articles, such as the Economic Policy Uncertainty Index by Baker, Bloom and Davis (2016), the “recession” topic by Bybee et al. (2024), or the Geopolitical Risk Index by Caldara and Iacoviello (2022). I extend this avenue by disentangling different moments of macro forecasts using large language models.

Second, I contribute to the literature on the importance of firm expectations for decision-making, focusing on investment. Most of the evidence is based on firm-level expectations and uncertainty about the business outlook, often measured by firms’ sales growth forecasts. For instance, Gennaioli, Ma and Shleifer (2016) document a dominant role for sales growth expectations in explaining changes in investment. Guiso and Parigi (1999) and Fiori and Scoccianti (2023) show negative effects of higher subjective sales growth uncertainty on investment. Alfaro, Bloom and Lin (2024) measure uncertainty using realized and option-implied volatility and employ an IV strategy to estimate the causal effect on firms’ decision-making. They also find a negative impact on investment. Bachmann, Elstner and Sims (2013) exploit time-series variation in firm-level uncertainty, measured in the ifo Business Survey, and document a negative effect on economic activity. Evidence is more limited regarding the influence of firms’ macro expectations on decision-making.⁴ An exception is Kumar, Gorodnichenko and Coibion (2023), who provide causal evidence for the negative effects of aggregate uncertainty on firm decisions. In this paper, I show

³Roth and Wohlfart (2020) elicit household expectations on the likelihood of a fall in GDP. They find elevated recession expectations and show that people extrapolate to their personal unemployment prospects.

⁴For inflation expectations, there is more causal evidence regarding the effects on firms’ decisions (Coibion, Gorodnichenko and Kumar, 2018; Coibion, Gorodnichenko and Ropele, 2020; Ropele, Gorodnichenko and Coibion, 2022).

that firms' subjective beliefs about macro tail events are also an important driver of investment decisions. This effect persists after conditioning on subjective first and second moments.

Third, there is a rapidly growing literature highlighting the importance of higher-order moments for economic dynamics. A consistent finding is the procyclical skewness of economic outcomes, which drives the business-cycle (Iseringhausen, Petrella and Theodoridis, 2023). Most closely connected is the work by Salgado, Guvenen and Bloom (2023), who focus on firm growth rates and also document procyclical skewness.⁵ Complementary to this, I provide direct firm-level evidence that subjective higher-order macro moments significantly influence decision-making. Based on my empirical results, I propose a mechanism in a quantitative model that generates a fat left tail in recessions after a macro tail risk news shock.

Fourth, motivated to explain the equity premium puzzle, there is a large body of work introducing rare disasters into asset pricing models (e.g., Rietz, 1988; Longstaff and Piazzesi, 2004; Barro, 2006; Gabaix, 2012; Wachter, 2013) and business cycle models (e.g., Gourio, 2012; Kozlowski, Veldkamp and Venkateswaran, 2020). Backus, Chernov and Martin (2011) and Kelly and Jiang (2014) both explore how macro tail risk can be inferred from financial markets, with the former using option prices to estimate the probability of macro disasters and the latter proposing a firm-level approach based on the cross-section of returns.⁶ I add to this literature by providing direct empirical evidence on the effects of macro tail risk on investment by exploiting variation in macro tail risk beliefs of firm decision makers.

Fifth, the heterogeneous firm model builds on Cooper and Haltiwanger (2006) and is a stripped-down version of more general models that carefully analyze the effects of uncertainty shocks at the firm level on investment and economic dynamics (e.g., Bloom, 2009; Bloom et al., 2018; Alfaro, Bloom and Lin, 2024; Bachmann and Bayer, 2014). I introduce heterogeneous exposure to aggregate downside risk into this model class. The mechanism in the model is based on ambiguity aversion, which has already been shown to be an important driver of the business cycle and portfolio decisions (e.g., Ilut and Schneider, 2014; Dimmock et al., 2016; Bianchi, Ilut and Schneider, 2018). In the literature, ambiguity

⁵Procyclical skewness is also documented for income risk (Busch et al., 2022; Guvenen, Ozkan and Song, 2014).

⁶This also touches on the empirical asset pricing literature, which shows that investors care more about downside than upside risk sensitivity (Ang, Chen and Xing, 2006; Lettau, Maggiori and Weber, 2014). Relatedly, Fiore, Martin and Nagler (2024) show that corporate investment decisions are affected through a risk-based disaster-vulnerability channel.

is typically about aggregate technology, while I introduce ambiguity about idiosyncratic exposure to macro tail events. Hence, the effects of ambiguity scale with increases in macro tail risk, with implications for, among other things, the effectiveness of policy.

The remainder of this paper is organized as follows. In Section 4.2, I introduce the data and characterize the nature of macro tail risk beliefs. After fixing ideas theoretically in Section 4.3.1, I relate macro tail risk beliefs to firm decisions based on a micro approach (Section 4.3.2), an experimental approach (Section 4.3.3), and a macro approach (Section 4.3.4). In Section 4.4, I present a quantitative heterogeneous firm model and analyze the effects of macro tail risk news shocks in interaction with policy. A final section concludes.

4.2 NATURE OF MACROECONOMIC TAIL RISK EXPECTATIONS

How do firms assess the subjective probability of a macro tail event occurring? In this section, I introduce the firm survey and document the nature of firms' macro tail risk expectations over time. In addition, I analyze how changes in macro tail risk expectations are associated with changes in firms' expectations about their own business. Lastly, I examine firms' beliefs about their exposure to macro tail events.

4.2.1 ifo Business Survey

To study the nature of firms' macro tail risk expectations, I add new questions to the ifo Business Survey (IBS-IND, 2023; IBS-SERV, 2023; IBS-TRA, 2023). Since 1949, the ifo Institute has conducted the IBS, a monthly survey of about 9,000 German firms that was first used for research by Theil (1955).⁷ The survey is used to provide business sentiment indicators in Germany, such as the renowned ifo Business Climate Index, which is among the top 20 macro news events moving global asset prices (Kerssenfischer and Schmeling, 2024) and one of Bloomberg's "12 Global Economic Indicators to Watch." Answers to the regular questions of the survey are reported for the EU Harmonized Business Survey. The ifo Institute incentivizes participation in the survey by automatically providing participants with the survey results free of charge as a thank you. To further encourage participation in the IBS, this report contains more detailed information, such as more

⁷Other papers that use the IBS include, for example, Bachmann, Elstner and Sims (2013); Bachmann et al. (2024b); Link et al. (2023); Buchheim et al. (2022); Buchheim, Link and Möhrle (2024); Enders, Hünnekes and Müller (2022); Bachmann and Elstner (2015).

disaggregated sectoral data, than the publicly available results. The panel dimension of the survey has a low attrition rate, making it well-suited for studying the expectations and behavior of firms over time. For more details and a comparison to other firm surveys, see Born et al. (2023). The IBS is representative of the German economy (Hiersemenzel, Sauer and Wohlrabe, 2022), enabling statements about macroeconomic effects. Crucially for analyzing the effects of expectations on decision-making, the survey is typically answered by the CEO or owner of the firm (Hennrich, Sauer and Wohlrabe, 2023).

Subjective expectations and uncertainty The survey contains a wide range of both qualitative and quantitative questions on firms’ expectations and uncertainty about their own business. In this paper, I focus on the quantitative questions.


In the main part of the survey, firms report subjective expectations regarding their own business over the next six months with a slider on a visual analog scale from 0 to 100, where 50 means no change. Similarly, firms report subjective uncertainty about their own business over the same period also with a slider on a visual analog scale from 0 to 100, where 50 means average uncertainty.

A subset of firms participates in the “uncertainty module” initiated by Bachmann et al. (2024b). Here, firms report their sales growth expectations in % for different horizons—one quarter ahead, one year ahead, and average annual growth two to four years ahead—in the worst, average, and best case. I use the average-case sales growth expectations as an alternative measure for business expectations and the range between best- and worst-case sales growth expectations as an alternative measure of business uncertainty. Appendix D.1 provides the wording of all survey questions used, translated into English.

Firms’ business decisions Importantly, the IBS also elicits information on business decisions. Most relevant for my subsequent analysis is the information on investment, which is collected twice a year in the survey. The questions are qualitative in nature and are framed in terms of changes compared to the previous year, see Appendix D.1 for the wording of the questions. My focus is on the November survey waves, in which firms report their investment plans for the next year as well as their investment in the current year.

Do the qualitative survey responses on investment decisions map into actual changes in the capital stock? For a subset of firms, I can link the survey data to annual balance sheet data provided by Orbis. While I cannot use the balance sheet data directly in my analysis due to the smaller sample size and missing data for 2021 to 2024, I can examine

Figure 4.1 – Survey responses on investment and change in tangible fixed assets



Notes: The figure shows how survey responses on investment (indicators) predict changes in the tangible fixed assets (balance sheet data). The bars represent the point estimates of three separate regressions based on Equation (4.1). The blue bars show the estimate of the constant α . The orange bars show the estimate of $\alpha + \beta$. Sample: 2016–2019.

how a survey response about changes in investment generally translates into changes in the stock of tangible fixed assets (a major component of the capital stock) on the balance sheet between 2016 and 2019. To do this, I run the following regression:

$$\Delta \log(Capital_{i,t}) = \alpha + \beta \mathbb{1}\{\Delta Investment_{i,t} < 0\} + \gamma_{t,s} + \delta_i + \epsilon_{i,t}. \quad (4.1)$$

For the survey responses on investment, I use (i) an indicator for a planned decrease in investment in the next year, (ii) an indicator for a realized decrease in investment in the current year, and (iii) an indicator for a downward revision of investment compared to previous plans. I exploit only the variation within firms at a given point in time within industries by including firm fixed effects δ_i and year $\times$ two-digit NACE industry fixed effects $\gamma_{t,s}$. Figure 4.1 shows the estimated average year-over-year change in tangible fixed assets $\Delta \log(Capital_{i,t})$ when the respective indicator is zero (in blue) or one (in orange).

The results show that the survey responses on investment predict large and significant changes in the capital stock. For example, when firms in the November wave of the survey report that they plan to reduce investment in the next year, tangible fixed assets actually fall by almost 5% on average in that following year. When firms do not plan to reduce investment, the capital stock increases by more than 2% on average. Hence, the qualitative survey responses predict very large movements in investment.

4.2.2 Firms' macroeconomic tail risk expectations

In October 2022, I supplemented the regular IBS with a set of questions eliciting firms' subjective probability of the occurrence of four specific macro tail events, which reads as follows:

What do you think:

What is the probability that the following events will occur in the next 5 years?


[Your answers can each be in a range between 0 and 100%, where 0 means absolutely unlikely and 100% means absolutely certain.]

- 1. A drop in annual GDP in the magnitude of at least 5%: ____*
- 2. A financial crisis (with an overall magnitude similar to the 2008/2009 financial crisis): ____*
- 3. A pandemic (with an overall magnitude similar to the COVID-19 pandemic): ____*
- 4. An escalation of geopolitical conflicts (direct war involvement of NATO countries):
____*

I chose these four scenarios because a large drop in GDP serves as a reduced-form measure of macro tail events, while the other three scenarios capture the typical reasons why the economy ends up in a macro tail event (Barro and Ursúa, 2012; Nakamura et al., 2013). I have added a short explanation to the events, where the references to recent crises should help participants to get an idea of the magnitude. The geopolitical scenario is framed as general as possible in order to capture overall geopolitical tensions. The range of five years is chosen to overcome potential issues with stating small probabilities.

It is important to note that there is no 'correct answer' to these questions, as such events are notoriously difficult to predict. Even when relying solely on the frequency of past occurrences, the small number of events can result in drastically different conclusions. Over the past 100 years, there have been four macroeconomic tail events in Germany: (i) the COVID-19 pandemic, (ii) the Great Recession, (iii) World War II, and (iv) the Great Depression. Thus, a reference to the last 100 years—20 five-year windows—implies a frequency of 20% in a five-year window. Considering only the post-WWII period leads to a lower probability, while looking at the last 20 years implies a probability of 50%. I repeated the questions about macro tail beliefs in April 2023 and October 2023, enabling me to study changes in beliefs within firms over time.

Figure 4.2 – Distribution of firms’ macro tail risk expectations over time



Notes: The figure shows the cross-sectional distribution of firms’ subjective probability of four macro tail events over the next five years for October 2022, April 2023, and October 2023 (balanced panel). The triangles indicate means, and the crosses indicate medians. The boxes correspond to the interquartile ranges. The lines correspond to the p10-p90 ranges. The question on the respective macro tail event is shown in the text.

Figure 4.2 shows the cross-sectional distribution of firms’ subjective probability of the four macro tail events over the three survey waves. Observations from October 2022 reveal notably elevated expectations on average: the median firm, marked with crosses, expects a large decline in GDP ($\geq 5\%$) with 50% probability, a financial crisis with 40%, a pandemic with 10%, and a geopolitical escalation with 30% probability over the next five years. In April and October 2023, the subjective probability of a significant decline in GDP decreased sharply by 30 percentage points (p.p.), those for a financial crisis by 20 p.p., while the median subjective probability of pandemic and geopolitical escalation scenarios remained unchanged.⁸ The particularly high subjective probabilities in October 2022 are consistent with the fears discussed in the media at that time of a major recession caused by a cut in energy supplies from Russia (Bachmann et al., 2022a; Moll, Schularick and Zachmann, 2023).

The elevated macro tail risk expectations might question firms’ general knowledge about macroeconomic developments. Appendix Figure D.1 shows the relation between the subjective probability of a large decline in GDP in the next five years (elicited in October 2022) and point forecasts of GDP growth in 2020 and 2021 (elicited in August 2020).⁹

⁸The strong fluctuations in macro tail risk expectations are consistent with Orlik and Veldkamp (2022), who provide a framework in which new data can lead to large revisions in tail probabilities.

⁹Dovern, Müller and Wohlrabe (2023) analyze these expectations of GDP growth in the IBS in more detail.

While there is a slightly negative relationship, even firms that are very pessimistic about a large decline in GDP still have very reasonable point forecasts of GDP growth that are close to those of professional forecasters during this period.

We also observe a large cross-sectional dispersion across firms. For example, the interquartile range spans from 20 to 80 percent for the GDP drop scenario, 20 to 60 percent for the financial crisis, 5 to 25 percent for the pandemic, and 20 to 50 percent for the geopolitical escalation scenario in October 2022. A variance decomposition reveals that macro tail risk expectations are persistent at the firm level: 52% to 62% of the variation can be explained by firm fixed effects, while common time variation accounts for between 1% and 13%, depending on the specific macro tail risk scenario. In Appendix D.4, I analyze the drivers of the level differences across firms. The results suggest an important role for experience in shaping firms' macro tail risk expectations.

Relation to subjective expectations and uncertainty Are macro tail risk expectations associated with firms' business outlooks? It is conceivable that firms might not incorporate macro tail risk expectations into their own business outlooks, particularly if they exhibit overconfidence (Malmendier and Tate, 2015) or focus on business or industry-specific news when forming expectations (Born et al., 2024; Andrade et al., 2022). I test the link between firms' macro tail risk expectations and business outlooks by regressing firms' business expectations (first moment) and business uncertainty (second moment) on their subjective probability of a macro tail event. To mitigate potential reverse causality—specifically, using local information to forecast aggregate growth—I control for year×industry fixed effects and the current business situations of firms. I add firm fixed effects to use only the variation in macro tail risk expectations within firms over time.

Table 4.1 presents the regression results for the four different macro tail events. An increase in macro tail risk expectations is significantly associated with a decline in business expectations (Columns 1–2) and an increase in business uncertainty (Columns 3–4); a relationship that persists even after controlling for business expectations (Column 5). In terms of magnitude, a one standard deviation increase in firms' subjective probability of a drop in GDP of at least 5% decreases their business expectations by 0.064–0.13 standard deviations and increases their business uncertainty by 0.053–0.097 standard deviations. The results are robust for the other three more specific macro tail events, where subjective expectations regarding the occurrence of a pandemic and a geopolitical escalation are more strongly associated with business uncertainty than business expectations. Appendix Table D.1 replicates the analysis using firms' quantitative sales growth expectations and uncertainty. Although the results become noisier as the sample size is halved, we once

Table 4.1 – Relation of macro tail risk expectations and business outlooks

	Business Expectations		Business Uncertainty		
	(1)	(2)	(3)	(4)	(5)
Subj. Prob. GDP Drop $\geq 5\%$	-0.13*** (0.011)	-0.064*** (0.013)	0.097*** (0.013)	0.061*** (0.012)	0.053*** (0.012)
Business Expectations					-0.12*** (0.013)
Subj. Prob. Financial Crisis	-0.13*** (0.012)	-0.056*** (0.014)	0.10*** (0.013)	0.049*** (0.013)	0.042*** (0.013)
Business Expectations					-0.12*** (0.013)
Subj. Prob. Pandemic	-0.033*** (0.011)	-0.018 (0.012)	0.054*** (0.013)	0.035*** (0.011)	0.032*** (0.011)
Business Expectations					-0.12*** (0.013)
Subj. Prob. Geop. Escalation	-0.075*** (0.010)	-0.024* (0.014)	0.064*** (0.012)	0.042*** (0.011)	0.039*** (0.011)
Business Expectations					-0.12*** (0.013)
Observations	10715	10715	10715	10715	10715
Business Situation	Yes	Yes	Yes	Yes	Yes
Firm FE	No	Yes	No	Yes	Yes
Year X Industry FE	Yes	Yes	Yes	Yes	Yes

Notes: The table presents the estimation results based on linear regressions of firms' subjective probability of a GDP drop of at least 5% / financial crisis / pandemic / geopolitical escalation in the next five years on firms' business expectations for the next six months (Columns 1-2) and firms' business uncertainty for the next six months (Columns 3-5). Both are measured on a visual analog scale from 0 (negative expectations / low uncertainty) to 100 (positive expectations / high uncertainty). Appendix D.1 provides the wording of the questions. All variables are normalized by their cross-sectional standard deviation. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

again observe a negative association with sales growth expectations and a positive association with sales growth uncertainty for the coming year. These results are consistent with the notion that macro tail risk expectations affect firms' business outlooks. Therefore, it seems natural that they may impact their decision making. In Section 4.3, I test whether macro tail risk expectations affect firms' investment decisions, even after conditioning on their first and second moment expectations.

Macro tail risk expectations also shift the mean of macro forecast distributions. Thus, it could be that macro tail risk expectations are mostly reflected in firms' macroeconomic point forecasts. Macro point forecasts are collected irregularly for a subset of firms in the IBS in the form of one-year-ahead unemployment rate expectations.¹⁰ Figures D.3 and D.4 in the Appendix show binscatter plots of the subjective probability of a large decline in GDP in the next five years and unemployment rate expectations as well as uncertainty about it one year ahead.¹¹ Unemployment rate expectations are elicited in the last month of the quarter, so I can relate the macro tail risk expectations reported in October to the unemployment rate expectations elicited in the respective previous September and subsequent December wave. While there is indeed a positive relationship—higher subjective probabilities of a large decline in GDP are associated with higher unemployment rate expectations—unemployment rate expectations are surprisingly similar across firms that strongly disagree on the probability of a macro tail event occurring. In fact, the relationship between the subjective probability of a large decline in GDP and the point forecast of the unemployment rate disappears when I use only within-firm variation. These results are consistent with Link et al. (2023), who find comparatively little variation in macro point forecasts across firms and close alignment with professional forecasters.¹² Firms' uncertainty about the unemployment rate one year ahead is also only weakly related to the subjective probability of a large decline in GDP and disappears when only within-firm variation is used. In sum, macro tail risk expectations represent a new dimension of heterogeneity that may influence firms' decision-making.

4.2.3 Firms' expected exposure to macro tail events

For firms, there is a second layer of macro tail risk expectations, which is their belief about what happens to the firm in the case of a macro tail event. To better understand these beliefs, I elicit the expected firm exposure to a hypothetical 5% decline in German GDP using both structured and open-ended survey questions.

Structured survey question To obtain quantitative exposure beliefs of firms, I elicit expected exposure to a 5% decline in German GDP in terms of real sales changes. Since

¹⁰GDP growth expectations are collected even less regularly, as mentioned in the text above.

¹¹This information comes from a quarterly voluntary module in the IBS, see Link et al. (2023) and Link et al. (2024a).

¹²Note that point forecasts in surveys provide a central tendency of beliefs, although it is not clear whether firms report means, medians, or modes, see Manski (2018). The weak link between the subjective probability of a macro tail event and firms' macro point forecasts suggests that firms may be more likely to report the median or mode.

firm exposure is notoriously difficult for managers to predict, I ask for potential exposure *scenarios* in the worst case, on average, and in the best case. Specifically, I added the following question to the IBS in October 2023:

Imagine a macroeconomic crisis in Germany (GDP collapses by $\approx 5\%$ in a year). To what extent do you expect that your firm would be affected? Please give a rough estimate of the potential impact on your firm's real sales in percent.


- *in the best case: — %*
- *on average: — %*
- *in the worst case: — %*

Firms' expected average exposure to a large drop in GDP is strongly heterogeneous. The median firm expects a 10% decline in real sales in the average scenario, while firms at the 10th percentile expect no change in sales and firms at the 90th percentile expect a 30% decline, see Appendix Table D.2. Figure 4.3 shows, in a jittered scatter plot, how firms' expected average exposure (horizontal axis) relates to their expected best- and worst-case scenarios (vertical axis; shown in blue and orange, respectively). We observe that even conditional on expected average exposure, firms report very different best- and worst-case exposures. On average, the difference between the worst-case exposure and average exposure is larger than the difference between the best-case exposure and average exposure, indicating a left-skewed exposure distribution, consistent with the findings of Salgado, Guvenen and Bloom (2023).¹³

There is substantial exposure uncertainty, measured by the distance between best- and worst-case exposure scenarios. For the median firm, there is a 20 percentage points difference between the expected best- and worst-case exposure. This uncertainty is systematically related to expected average exposure: while the best- and worst-case exposures are close to the dashed 45-degree line for expected average exposures around zero, the gap widens as the average exposure becomes more negative. Since the expected worst-case exposure is at the center of the proposed mechanism in the next section, I analyze it in more detail.

¹³Appendix Figure D.2 shows that higher subjective exposure beliefs are slightly positively correlated with a higher subjective probability of the occurrence of a macro tail event, but not with changes in those subjective probabilities.

Figure 4.3 – Expected firm exposure scenarios in case of a 5% decline in GDP



Notes: The figure shows how firms’ expected average exposure scenario relate to their worst- and best-case scenarios in terms of real sales changes in case of a 5% decline in German GDP. The scatter plot is jittered—random noise is added to the scatter—to align with data-protection rules and avoid overplotting. The blue (orange) line shows the fitted linear regression of average exposure on best-case (worst-case) exposure. The black dashed line is the 45-degree line.

Motivated by Kozlowski, Veldkamp and Venkateswaran (2020), one hypothesis is that firms’ expected exposure scenarios are shaped by experienced exposure to realized macro tail events in the past. The idea is that firms have imperfect knowledge about exposure and view the realized exposure in a macro tail event as a draw from their firm-specific distribution. I test this potential mechanism by relating firms’ expected worst-case exposure, elicited in 2023, to their realized exposure in the Great Recession and the COVID-19 pandemic, leveraging the panel dimension of the survey. I proxy firm-specific exposure to the Great Recession with the share of months in which firms report negative business conditions in 2008 and 2009. I utilize a more direct proxy for exposure to the COVID-19 pandemic, as firms were asked in the survey about the impact of COVID-19 on their business on a seven-point Likert scale from -3 (negative) to $+3$ (positive). I construct an indicator of a strong negative impact from COVID-19, defined as firms reporting an average COVID-19 impact score of less than -1 , following Balleer et al. (2024). Table 4.2 presents the regression results. Indeed, higher exposure to realized macro tail events in the past is significantly associated with a more negative expected worst-case exposure in 2023. Although the magnitude of this relationship is diminished, it remains significant when controlling for average and best case exposure and is also robust to the inclusion of industry fixed effects. First, these results act as a validation exercise, confirming that

Table 4.2 – Expected worst-case exposure and realized exposure to the Great Recession and the COVID-19 Pandemic

	Expected worst-case firm sales change if GDP drops by 5%					
	(1)	(2)	(3)	(4)	(5)	(6)
Neg. Business State (08/09)	-13.3*** (2.26)	-2.71** (1.19)	-2.86** (1.23)			
Strong COVID-19 impact				-6.41*** (0.77)	-0.89** (0.42)	-1.19*** (0.45)
Average exposure		1.73*** (0.069)	1.72*** (0.070)		1.83*** (0.039)	1.78*** (0.038)
Best-case exposure		-0.63*** (0.10)	-0.61*** (0.10)		-0.75*** (0.056)	-0.68*** (0.054)
Constant	-24.6*** (0.68)	-5.85*** (0.56)	-5.88*** (0.58)	-27.1*** (0.52)	-6.33*** (0.38)	-6.63*** (0.41)
Observations	1252	1252	1252	3167	3167	3167
Industry FE	-	-	✓	-	-	✓

Notes: The table shows how firms’ expected worst-case exposure in terms of real sales changes in case of a 5% decline in German GDP relate to firms’ past realized exposure in macroeconomic crises. Columns 1-3 are based on the share of months in which a firm reports a negative business situation in 2008 and 2009 as a proxy for the exposure to the Great Recession 2008/2009. Columns 4-6 are based on the self-reported impact of COVID-19 on the business as a proxy for the exposure to the COVID-19 pandemic. “Strong COVID-19 impact” is an indicator based on a reported COVID-19 impact in the IBS on a 7-point Likert scale from -3 (negative) to +3 (positive). The indicator is one when firms report an average COVID-19 impact score of less than -1. Robust standard errors in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

firms’ subjective exposure beliefs are meaningfully correlated with realized exposure during past macro tail events. Second, the results align with the idea of scarring effects at the firm level, where strong negative experiences in a macro tail event persistently shape firms’ expectations.

Open-ended survey question To gain a deeper understanding of how managers think about what happens to their firm in the case of a macro tail event, I incorporated an open-ended question to the survey. The main advantage of the open-ended survey question format is that it prevents respondents from being influenced by predefined response categories, see Haaland et al. (2024) for a review on open-ended questions. In my setup, this approach allows firms to freely express their thoughts on exposure beliefs. Specifically, I posed the following question in the IBS in June 2024:

To what extent can you generally forecast the effects of such a potential macroeconomic crisis (GDP collapses by around 5% within a year) on your real sales? Please let us know your thoughts on this macroeconomic scenario here:

The responses are hand-coded into categories by a research assistant following a codebook provided in the Appendix Table D.10. I received 2220 responses with an average length of 161 characters. The main takeaway from analyzing the exposure narratives of firms is that there is substantial uncertainty about potential exposure to macro tail events. In the following, I list the three main findings that lead to this conclusion. Appendix D.5 provides a more detailed discussion of the open-ended text responses.

First, the most common response from firms directly expresses major difficulties in forecasting exposure to macro tail events. Among those who comment on assessability, 94% report that it is difficult or even (almost) impossible to assess their potential exposure to such events. Second, when firms share exposure in terms of quantitative sales changes, they tend to report scenarios that deviate from central tendencies. In fact, they are more likely to share exposure in terms of ranges or max/min scenarios rather than averages. Third, firms refer to their—sometimes heterogeneous—exposure to past macro tail events. This highlights that firms have to rely on a small set of experiences—typically mentioning the Great Recession and COVID-19—to gauge their exposure to future macro tail events.

4.3 MACROECONOMIC TAIL RISK EXPECTATIONS AND FIRM DECISIONS

In this section, I analyze how macro tail risk expectations influence firms’ decisions, focusing on investment. After presenting a stylized model, I take a threefold approach to empirically assess the role of macro tail risk expectations. First, I employ a “micro approach” that utilizes observed variation in firms’ macro tail risk expectations in the cross-section and over time. Second, I implement an “experimental approach,” where I confront firms in vignettes with macroeconomic outlooks, varying only the skewness, and elicit the corresponding investment decisions. Third, in a “macro approach,” I construct a time series of macro tail risk expectations and other subjective macro moments based on newspaper articles, and then analyze the relation with future aggregate investment dynamics.

4.3.1 Stylized model framework

Before diving into the empirical analysis, I illustrate with a stylized model, how a firms’ macroeconomic tail risk expectations may affect the investment decision of the firm.

Model setup Consider a firm choice model with heterogeneous exposure to macro tail events. Firms, indexed by i , use capital, which fully depreciates each period, to produce output with a decreasing return to scale production function $A_i k^\alpha$, where A_i denotes the productivity of the firm.

Differences in productivity across firms arise only due to firm-specific exposure to aggregate productivity. Aggregate productivity A is stochastic and follows a two-state Markov process, switching between a high (A^h) and low (A^l) productivity state. Here, A^l represents periods of severe economic downturns, i.e. macro tail events. For simplicity, all firms have the same exposure to the normal productivity state A^h . However, there is heterogeneous exposure β_i to the macro tail event A^l . Firms contemplate three different scenarios of β_i : a worst-case, average-case, and best-case scenario, such that $\beta_i \in \{\beta_{i,worst}, \beta_{i,average}, \beta_{i,best}\}$. Therefore, the expectation about future productivity for a given firm in a current high productivity state can be written as follows:

$$\mathbb{E}[A_{i,t+1}|A^h] = p_{h \rightarrow h}A^h + p_{h \rightarrow l}[A^h - \mathbb{E}[\beta_i](A^h - A^l)] \quad (4.2)$$

where $\mathbb{E}[A_{i,t+1}]$ denotes expected productivity of the firm and $\mathbb{E}[\beta_i]$ is the expected exposure.¹⁴ For now, I assume that firms can confidently assess probabilities to the different exposure scenarios. This assumption will be subsequently relaxed by introducing ambiguous exposure. Intuitively, the exposure of the firm β_i scales the aggregate productivity change ($A^h - A^l$), where for firms with $\mathbb{E}[\beta_i] = 1$ own productivity equals aggregate productivity in the macro tail event.

The decision to introduce these different exposure scenarios in the model is motivated by the fact that firms rarely experience macro tail events as well as evidence on scenario planning of firms by Graham (2022) and Bachmann et al. (2024b). The imperfect knowledge about firms' exposure adds a second layer of uncertainty: There is uncertainty not only about future macroeconomic conditions, but also about firm-specific exposure in severe recessions, a situation of compound risk.

Investment decision Firms' objective is to maximize expected profits. Firms are risk neutral and discount the future with $1/(1+r)$, where r is the cost of capital. In this stylized setup, the optimal capital stock is:

¹⁴For simplicity, I assume that the average-case exposure scenario $\beta_{average}$ equals the firm's expected exposure.

$$k_i^* = \left(\frac{\alpha \mathbb{E}[A_{i,t+1}|A^h]}{1+r} \right)^{\frac{1}{1-\alpha}} \quad (4.3)$$

Higher expected productivity $\mathbb{E}[A_{i,t+1}|A^h]$ leads to a larger optimal capital stock, which is equal to investment in this simple model. For a given expected exposure $\mathbb{E}[\beta_i]$, we can state that an increase in the subjective probability of a macro tail event $p_{h \rightarrow l}$ has a first-moment effect on the optimal level of capital, as visualized in Figure 4.4 with the blue line. When a firm has a higher expected exposure $\mathbb{E}[\beta_i]$, this would imply a larger first-moment effect on the optimal capital stock, operating through the expectation of productivity. Note that only the expectation of the next period's production matters for the choice of capital. Beliefs about higher order moments do not play a role in this setup. For instance, an increase in uncertainty through a mean-preserving spread of the aggregate productivity process or changes in worst-case and best-case firm-specific exposure scenarios have no effect on the optimal capital stock.


Ambiguity aversion How can macro tail risk expectations affect investment decisions beyond the first moment? In general, there are many potential sources of non-linearity in firms' policy functions that imply effects of changes in macro tail risk expectations on investment decisions beyond the first-moment effect. In the following, I focus on a novel mechanism that takes into account the empirically documented substantial uncertainty of firms about their exposure to macro tail events.

Specifically, I make the arguably reasonable assumption that firms are not confident enough to assign probabilities to the exposure scenarios. This is an instance of ambiguity, which refers to uncertainty when—distinct from risk—the odds are not known, see Ilut and Schneider (2022) for a review. In general, uncertainty about probability assessments can be described by the size of the belief set governing the distribution of the next period state, as conceptualized by Gilboa and Schmeidler (1989). In the context of ambiguous exposure to macro tail events, I parameterize the set of beliefs by a set of event realizations consisting of the firm-level exposure scenarios $\beta_i \in \{\beta_{i,worst}, \beta_{i,average}, \beta_{i,best}\}$.¹⁵

Agents prefer to know the odds of an event—they are ambiguity averse. Indeed, a large body of experimental work shows that ambiguity aversion is a widespread phenomenon,

¹⁵Specifying the uncertainty about probability assessments by a set of event realizations is a standard approach in the literature. For example, Ilut and Schneider (2014) parameterize ambiguity about TFP growth by an unknown TFP component with an interval of means centered around zero.

Figure 4.4 – Optimal capital stock and macro tail risk expectations



Notes: This figure shows the percentage deviation from the high productivity state ($A^h = 1$) optimal capital stock due to changes in the expected probability of transitioning to the low productivity state ($A^l = 0.9$) in the next period. The first-moment effect is visualized in blue and arises from the fact that the subjective probability of a macro tail event influences expected productivity. The ambiguity aversion effect is visualized in orange and is determined by firms' expected worst-case exposure to the low productivity state. The interest rate is set to $r = 0.02$ and output elasticity $\alpha = 0.33$. The calibration of the exposure to macro tail events is based on the median firm in the survey, see Table D.2. Average exposure is $\beta_{average} = 1$ and worst-case exposure is $\beta_{worst} = 2.5$.

see Trautmann and Van De Kuilen (2015) for a review.¹⁶ Following Gilboa and Schmeidler (1989), agents deal with ambiguity by choosing actions *as if* they were using the most pessimistic plan from the belief set. Through the lens of the model, firms act *as if* their worst-case exposure to the macro tail event, that minimizes expected future output, will be realized when the economy is in A^l . Thereby, the optimal capital stock decreases beyond the first-moment effect as the subjective probability of a macro tail event increases, as visualized in Figure 4.4 with the orange line. One key advantage of ambiguity is its source dependence, meaning that it does not influence decisions made under risk, where probabilities are known. Thus, when the subjective probability of a macro tail event is very low, the model converges to a standard firm choice problem.

Note that while for actions—such as investment—there is an observationally equivalent model with pessimistic beliefs, this equivalence breaks down when we ask people about their forecast, such as firms' sales forecasts. To further illustrate this point, consider the classic Ellsberg paradox involving a choice between two urns, urn A with a known

¹⁶Interestingly, there are studies showing that ambiguity aversion especially arises in the context of compound risk (Abdellaoui, Klibanoff and Placido, 2015; Halevy, 2007).

distribution of red and black balls (50/50 ratio) and urn B with an unknown distribution (Ellsberg, 1961). Although participants are indifferent between betting on a particular color from urn B with an unknown distribution—implying an expected 50/50 ratio—they show a clear preference for urn A with a known distribution over urn B. This behavior indicates that while their forecasts (e.g., the expected ratio of red to black balls) remain neutral, their actions—preferring the known risk of urn A over the ambiguity of the unknown in urn B—reveal an underlying aversion to ambiguity. This analogy underscores that firms’ cautious behavior in the face of uncertainty does not imply pessimistic forecasts in surveys.

In Appendix D.6, I explore two other mechanisms that can also generate negative effects on investment beyond the first-moment effect. First, I consider shareholder maximization with risk averse investors in a Bayesian world, where probabilities can be assigned to each state. Second, I introduce bankruptcy risk. I show that these alternatives generate quantitatively rather small effects beyond the first-moment effect. It would require an extremely high degree of risk aversion (CRRA parameter $\eta = 20$) or bankruptcy risk to obtain similar magnitudes as in the model with ambiguity aversion in this stylized setup. In addition, since idiosyncratic risk diversifies away—in contrast to ambiguity (see Appendix D.6 for details)—the investment decisions are not directly associated with firms’ worst-case exposure beliefs.¹⁷ I also empirically test both alternative mechanisms in the next section, and the findings suggest that they are not the main drivers of the results.


4.3.2 Micro approach

4.3.2.1 Empirical strategy

The goal is to establish a causal link between macro tail risk expectations and firms’ investment decisions. In an ideal world, one would compare two completely identical economies, shock one of them with higher macro tail risk expectations and then compare the investment dynamics. I employ this approach in a model simulation exercise in Section 4.4. Since this is not empirically feasible, the main idea of the micro approach

¹⁷In an entrepreneur-type model, where the entrepreneur’s only source of income is the output of her own firm, there is a direct relationship with firm’s worst-case exposure beliefs. Again, however, the risk aversion parameter must be extremely high ($\eta \approx 10$) to obtain similar magnitudes as in the model with ambiguity aversion, see Appendix D.6 for details.

is to exploit the substantial variation in the subjective probability of a macro tail event occurring at the firm level. Macro tail risk expectations are elicited in October 2022 and 2023, while investment plans for the next year are elicited one month later, see Figure 4.5.



I regress an indicator for a planned decrease in investment $\mathbb{1}\{\Delta Investment_{i,t+1} < 0\}$ in the next year on the subjective probability of a macro tail event occurring within the next five years $\mathbb{P}_{i,t}(MacroTailEvent)$ of firm i :

$$\mathbb{1}\{\Delta Investment_{i,t+1} < 0\} = \alpha_t + \beta \mathbb{P}_{i,t}(MacroTailEvent) + \gamma X_{i,t} + \delta_i + \epsilon_{i,t}. \quad (4.4)$$

Here, time fixed effects α_t flexibly absorb macroeconomic circumstances. In addition, I control for subjective first and second moments at the firm level by including business expectations and business uncertainty in the regression. I also control for current business conditions. These three variables are based on a visual analog scale ranging from 0 to 100, as introduced in the previous section. I conduct several robustness checks by varying the set of control variables $X_{i,t}$. By including firm fixed effects δ_i , I use only within-firm variation.

While this first analysis is rather descriptive, four aspects speak against major bias. First, firms are forecasting the probability of a single macro outcome at the same point in time, namely a macro tail event. Therefore, the estimation is unlikely to be prone to reverse causality in the sense that firms' investment plans will mechanically influence managers to be more pessimistic about macro tail events, see Gennaioli, Ma and Shleifer (2016) for a similar argument. Second, investment plans are elicited in the subsequent wave after the macro tail risk expectations. This minimizes potential response bias due to general sentiment at the time of the survey. Third, by controlling for business conditions, business expectations, and business uncertainty, I account for many other potentially confounding

factors. For example, if a firm is simply more pessimistic in general or extrapolates from their own business situation. Fourth, I additionally estimate specifications that rely only on within-firm variation between October/November 2022 and 2023 by controlling for firm fixed effects. This addresses the concern that certain firm characteristics may be correlated with both, macro tail risk beliefs and investment decisions.


4.3.2.2 Results

I start by mapping macro tail risk expectations to the decision to reduce investment in the following year using the pooled data. Figure 4.6 visualizes the relation in a binned scatterplot. The raw relation in Panel (A), which includes only time fixed effects, reveals a strong and significant positive association between the subjective probability of a large decline in GDP and the decision to reduce investment in the next year. In terms of magnitude, a standard deviation increase in macro tail risk expectations is associated with a 7.5 percentage point increase in the share of lower investment plans. Importantly, Panel (B) demonstrates that the positive association remains intact when we control for the current business situation as well as for the subjective first and second moments, namely business expectations and uncertainty over the next six months. Consistent with the model predictions in the previous section, the point estimate shrinks as the subjective probability of a macro tail event also moves business expectations and uncertainty. However, more than half of the total effect size remains even after accounting for the subjective first and second moments. Appendix Figure D.5 also displays the regression results based on the expectations of the occurrence of the other three more specific macro tail events: financial crisis, pandemic, and geopolitical escalation. In all cases, I estimate a strongly negative effect on investment.

I now add firm fixed effects and thereby exploit only variation within firms over time. Naturally, I lose a large amount of variation in macro tail risk beliefs (up to 62%, see Section 4.2) and attenuation bias due to measurement error is potentially larger. While the point estimate is indeed somewhat smaller and estimated less precisely, the significant positive relation is still there, see Panel (C) in Figure 4.6. The effects also remain significant for the financial crisis and pandemic scenarios, but become insignificant for the geopolitical escalation scenario, see Appendix Figure D.5.

In the next step, I analyze heterogeneity with respect to firms' expected worst-case exposure to a large GDP drop. The theoretical framework in Section 4.3.1 predicts particularly strong effects beyond the first-moment effect when firms are highly exposed in the expected worst case. Indeed, the results in Panel (D) of Figure 4.6 show particularly

Figure 4.6 – Firms’ macro tail risk expectations and investment plans



Notes: This figure shows binned scatterplot based on Equation (4.4). Table D.3 displays the corresponding point estimates. The dependent variable is an indicator of a planned decrease in investment in the next year compared to the previous year, elicited in the November waves 2022 and 2023. The macro tail risk variable is the subjective probability of the occurrence of a large decline in annual GDP ($\geq 5\%$) in the next five years, elicited in the October 2022 and 2023 waves. In Panel (D), all explanatory variables are interacted with an indicator for above median worst-case exposure to a macro tail event. Standard errors are clustered at the firm level.

strong effects for firms with above-median worst-case exposure, which are significantly different from those of firms with below-median worst-case exposure.¹⁸ The results are similar when flexibly controlling for industry-specific time trends with year $\times$ industry fixed effects instead of year fixed effects, see Appendix Table D.3. For a subset of firms, the survey includes information on quantitative sales growth expectations and uncertainty

¹⁸Further, as explained in Section 4.3.1, ambiguity aversion does not imply that firms make their forecasts based on their worst-case exposure. Indeed, I find that the change in business expectations associated with changes in macro tail risk expectations does not significantly differ between firms with above- and below-median worst-case exposure, see Appendix Figure D.7.


introduced in the previous section. In Panel (B) of Appendix Table D.3, I repeat the analysis using quantitative year-on-year sales growth expectations as controls and obtain very similar results.

The main takeaways from the analysis so far are threefold. First, higher macro tail risk expectations have a strong negative effect on investment. Second, the effect remains significant and sizable even after controlling for proxies of first- and second-moment expectations. Third, the negative effect is particularly pronounced for firms that expect to be heavily exposed in worst-case scenarios. This provides suggestive evidence that firms act on their worst-case exposure as predicted by the stylized model in Section 4.3.1. When using average exposure instead of worst-case exposure, the observed differences are not as strong, see Appendix Figure D.6. In the experimental approach in Section 4.3.3, I test this mechanism more formally.

Additional robustness One potential caveat of the baseline analysis is that the controls for the business expectations and uncertainty cover only the next six months, while the macro tail risk expectations are elicited over a five-year horizon to circumvent problems with eliciting small probabilities. The alternative specification in Appendix Table D.3 with quantitative year-on-year sales growth expectations discussed above, already partly addresses this issue. Moreover, the quantitative sales growth expectations and uncertainty question is available over various horizons, up to four years into the future (although answered by fewer firms). I re-run the main specification for short-, medium-, and long-term sales growth expectations and find very similar effect sizes, although the estimates are less precise due to the smaller sample size, see Appendix Figure D.8. For an even smaller subset of firms, I also observe unemployment rate expectations and the associated uncertainty in the last month of a given quarter. Using the September and December unemployment rate expectations as controls to absorb macro expectations and uncertainty leads to very similar results, see Appendix Figure D.9.

Other outcome variables Instead of planned investment, I also examine investment revisions by constructing an indicator of whether firms invested less than planned in 2023 and relating this to the change in macro tail risk expectations from 2022 to 2023. For all four subjective probabilities of macro tail events, I find a strong negative effect on investment revisions, consistent with the baseline results, see Appendix Table D.4. In addition, I use information on quantitative realized investment, collected in the survey for 2021, 2022, and 2023. Although the timing is not as closely aligned with the question of the macro tail risk expectations, I still find that a higher subjective probability of a macro tail event in October 2022 predicts a strong decline in investment from 2021 to

Figure 4.7 – Heterogeneity across firm characteristics



Notes: This figure shows the estimates of a one percentage point increase in macro tail risk expectations on the likelihood of reducing investment in the next year based on Equation (4.4), where all explanatory variables are interacted with indicators that split the sample by firm characteristics. The dependent variable is an indicator of a planned decrease in investment in the next year compared to the previous year, elicited in the November waves 2022 and 2023. The macro tail risk variable is the subjective probability of the occurrence of a large decline in annual GDP ($\geq 5\%$) in the next five years, elicited in the October 2022 and 2023 waves. The sample is restricted to firms with above-median worst-case exposure. “Firm size” is an indicator that equals one if the number of employees is ≥ 50 . “Family firm” is an indicator of whether the firm considers itself to be a family firm. “Risk taker” is an indicator that equals one if the manager selects a value > 6 (median) on a Likert scale from 0 (not at all willing to take risks) to 10 (very willing to take risks) when asked whether she considers herself a “risk taker” in general. “Equity ratio” is an indicator that is one if the equity ratio (self-reported in 2022) is above the median. 90% confidence intervals are depicted. Standard errors are clustered at the firm level.

2022, while April and October 2023 macro tail risk expectations show no relationship. This changes when analyzing the dynamics of investment from 2021 to 2023, where a higher subjective probability of a macro tail event in April 2023 now predicts a strong decline, see Appendix Table D.5. In terms of magnitude, a one standard deviation increase in the subjective probability of a macro tail event leads to a decline in investment of about 8.4% to 9.6% (based on Columns 6 and 7).

Heterogeneity I now study heterogeneity in the effect sizes across a number of firm characteristics. This allows me to test for further mechanisms. I use the baseline regression setup with and without firm fixed effects and interact firm characteristic indicators with

all explanatory variables. To ensure that the results are not driven by the exposure to macro tail events, I restrict the analysis to firms with above-median worst-case exposure (corresponding to the orange group in Figure 4.6 Panel (D)).

Figure 4.7 presents the results. First, the effect sizes are similar for small and large firms. If anything, the effects are greater for larger firms, implying that changes in macro tail risk also have strong aggregate effects on investment dynamics. Second, we observe larger effect sizes for family-owned firms, which could be due to several factors. One is that there are typically fewer decision-makers in family firms, see Bertrand and Schoar (2006) for an overview of the characteristics of family firms. Thus, idiosyncratic expectations of the owner are likely to have a stronger impact on the firm’s decision making. Moreover, larger effect sizes for family-owned firms support the idea that the owners’ primary goal is the long-term success of the firm. In contrast, managers in non-family firms are more likely to ignore tail risks due to limited liability, as formalized by Maćkowiak and Wiederholt (2018). Third, I analyze heterogeneity with respect to the degree of risk-taking of the managers. Theoretically, the mechanism driving the strong effects beyond the first moment could be due to risk-averse managers. However, if anything, the heterogeneity results indicate rather stronger effects for firms with *more* risk-taking managers, labeled as “risk taker” in Figure 4.7. One reason for this could be that the firms with more risk-averse managers are better prepared and diversified for a macro tail event. Lastly, I examine the role of firms’ leverage. The effect sizes are similar for firms with high and low equity ratios, suggesting that financial frictions, including bankruptcy risk, do not seem to be a dominant factor.

4.3.3 Experimental approach

4.3.3.1 Hypothetical vignettes of macroeconomic outlooks

In the second approach, I employ experimental methods to identify the causal effect of macroeconomic tail risk expectations on investment.¹⁹ To isolate the effect of macroeconomic tail risk expectations on investment decisions beyond first- and second-moment effects, I design hypothetical vignettes on macroeconomic outlooks, in which I vary only the higher-order moments. Specifically, I present managers two simple macroeconomic outlooks that have different skewness and kurtosis but identical mean and variance. They

¹⁹See Haaland, Roth and Wohlfart (2023) for a review of hypothetical vignettes. Recent examples of studies employing vignettes in a macroeconomic context are Roth, Wiederholt and Wohlfart (2023), who improve the understanding of how monetary policy affects household consumption and Gödl-Hanisch and Menkhoff (2023), who study price pass-through dynamics at the firm level.

are then asked to rank the outlooks based on which one they are more likely to make large investments in or whether they are indifferent. The baseline vignette reads as follows:

Outlook A:

*The German economy (measured by the **GDP**) will either **decline by -1% with 50% probability or grow by 5% with 50% probability** in the next year.*

Outlook B:

*The German economy (measured by the **GDP**) will either **decline by -7% with 10% probability or grow by 3% with 90% probability** in the next year.*

(In both outlooks, the GDP level returns back to the current level in the subsequent year.)

Under which outlook would your firm be more likely to make major investments?

Outlook A	<i>O</i>
No difference	<i>O</i>
Outlook B	<i>O</i>

The hypothesis is that firms are less likely to invest in Outlook B, where considerable macro tail risk is present due to the negative skewness and higher kurtosis.²⁰ To circumvent different expectations of subsequent economic dynamics that could affect the response, I add the information that the GDP level returns to the previous level in the subsequent year in both outlooks. To convey the outlooks as clearly as possible, I also provide a visualization as shown in Figure D.10 and chose outlooks where the growth rates and probabilities have no decimal places.

A potential limitation of the approach is that the results may depend too much on the specific values chosen in the simplified outlooks. A specific concern is potential loss aversion or recession aversion, which implies a central role for the sign of GDP growth rates. To address this concern, both outlooks include a scenario with positive and negative growth. Under loss aversion, firms would invest less in Outlook A, because there is a higher probability of a recession in the next year. Thus, the baseline vignette can be interpreted as a rather conservative approach and a lower bound for testing the hypothesis. Therefore, for those firms that actually chose to invest more in Outlook B with negative skewness, I present a second vignette, in which all growth rates are shifted upward by two percentage

²⁰In Outlook A and B, the mean (1. moment) is 2 and the variance (2. moment) is 9. The skewness (3. moment) is 0 in Outlook A and -2.6 in Outlook B. The kurtosis (4. moment) is 1 in Outlook A and 8.1 in Outlook B.

points, see Figure D.11. Now, the mean is at 4% and the other moments are the same as before. Firms again choose which outlook they would invest more in. Since Outlook A now contains only positive scenarios, the results can be interpreted more as an upper bound to test the hypothesis.

I conducted the experiment in the ifo Management Survey IMS (2023) during July and August 2023. The survey asks decision-makers in German firms about current economic policy issues and changes in the macroeconomic environment (Sauer, Schasching and Wohlrabe 2023). The survey covers about 350 firms across all sectors of the economy. The majority of respondents are the owner or CEO of their respective firm. I also elicit firms' expected exposure scenarios in the worst, average, and best cases in terms of real sales changes conditional on a GDP drop of 5% in exactly the same way as in the ifo Business Survey to formally analyze whether firms' exposure scenarios predict their investment decisions in the vignettes.


4.3.3.2 Results

In the baseline scenario, significantly more firms would invest more in symmetric Outlook A than in Outlook B, which has a 10% probability of exposing firms to a macro tail event in the next year (a 7% drop in GDP), as visualized in Figure 4.8. This is despite the fact that the symmetric Outlook A has a higher probability of a recession. When we also consider the shifted scenarios, we observe an even stronger preference for the symmetric scenario. About 40% of firms would invest more in Outlook A, more than twice as many as in Outlook B.²¹ The results are similar when considering only firms where the survey respondent is known to own at least part of the firm, see Panel (B) of Figure 4.8.

A crucial advantage of the hypothetical vignette is that it fixes the probability of a macro tail event occurring—the first layer of macro tail risk for firms. We can now analyze the extent to which the decisions of the firms in the vignette are related to their exposure beliefs—the second layer of macro tail risk for firms—in the case of a 5% drop in GDP. Specifically, I run a regression with the firm's chosen outlook as dependent variable and the firms' expected best-case, average, and worst-case exposure scenarios as explanatory variables. Through the lens of the mechanism proposed in Section 4.3.1, one would expect that firms with a highly negative worst-case scenario would be less likely to invest in Outlook B, which contains negative skewness and therefore high macro tail risk.

²¹Note that this result includes the responses to the baseline scenario of firms that would invest more in the symmetric Outlook A or that were indifferent between the outlooks.

Figure 4.8 – Firms’ investment choice in the hypothetical macro outlooks



Notes: The figure shows the share of firms choosing the symmetric outlook A (blue) and skewed outlook B (green) in the baseline scenario and shifted scenario. Panel (A) is based on the full sample, while Panel (B) focuses solely on firms where the respondent is known to be at least a partial owner. The whiskers reflect the standard errors of the means. The asterisks indicate the significance levels testing the rejection of the equivalence between the proportion of firms choosing the symmetric outlook A and those choosing the skewed outlook B. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.3 – Investment choice and firm exposure to a macro tail event

	Base scenarios		Shifted scenarios	
	Invest more in ...			
	(1)	(2)	(3)	(4)
	1: sym ;	1: sym	1: sym	1: sym
	0: no diff	0: no diff	0: no diff	0: no diff
	-1: skew	0: skew	-1: skew	0: skew
$\beta_{i,worst}$ (exposure to GDP↓ 5%)	-0.018**	-0.010**	-0.021***	-0.014***
	(0.008)	(0.005)	(0.006)	(0.004)
Observations	327	327	327	327
Average & best-case exposure ctrls	Yes	Yes	Yes	Yes

Notes: The table shows the estimation results of a linear regression. The dependent variable is the firms’ decision in the hypothetical vignette, coded differently across columns. Columns 1 and 2 are based on firms’ responses in the baseline scenario, while Columns 3 and 4 also consider the shifted scenarios. The explanatory variables are firms’ beliefs on exposure scenarios (worst-case, average, and best-case) in terms of changes in real sales, when GDP drops by 5%, elicited after the vignette in the ifo Management Survey. In Section 4.2, I introduce the question. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.3 shows the regression results for different codings of the choice in the vignette. Indeed, a more negative expected worst-case exposure significantly predicts the choice to

invest more in the symmetric outlook, while controlling for average and best-case exposure. This holds for both, the baseline and the shifted scenarios. The slightly stronger effects in the shifted scenarios in Columns (3) and (4) may be due to the fact that some firms interpret the -1% growth scenario already as a macro tail event. The magnitude is economically meaningful: expecting a 10 p.p. stronger decline in real sales in the worst-case exposure scenario increases the probability of choosing the symmetric scenario by 10 to 14 p.p. based on Columns (2) and (4).

4.3.4 Macro approach

How do firms' macro tail risk expectations evolve over time and how do they relate to future aggregate investment dynamics?

In the third approach, I exploit recent advances in artificial general intelligence research, where state-of-the-art models exhibit increasingly human-like behavior. I generate beliefs about macro tail risk with a large language model (LLM) based on articles of the largest business newspaper in Germany. The underlying premise is that managers' macroeconomic outlooks are shaped by the information provided by daily newspapers. First, I explain the belief generation process using the GPT API. Second, I present and examine the macro tail risk expectations index as well as an index of average macro forecasts and macro upside risk expectations to control for other moments of macro news. Then, I run local projections to analyze the relation between changes in macro tail risk beliefs and future aggregate investment dynamics.

4.3.4.1 Generating beliefs with a large language model

How can macroeconomic beliefs be generated using LLMs, and why is this approach feasible? There has been significant progress in the field of LLMs. A milestone was the invention of the "transformer architecture", which is also used in the GPT family. The GPT models generate output by predicting the next word in a sequence based on the words that have come before. They accomplish this by focusing on the importance of specific words in the input and using that information to make predictions. This process is repeated for each word, resulting in the generation of coherent and contextually appropriate text (Vaswani et al., 2017). Horton (2023) demonstrates that these models can be used as "simulated economic agents" by showing similar behavior as humans in classic experiments from the behavioral economics literature. Bybee (2023) is the first using LLMs to generate beliefs about macroeconomic outlooks and finds similar biases as in the previous literature based on surveys. I also refer to Bybee (2023) for a more

detailed methodological discussion. To date, the economic literature has primarily relied on more traditional natural language processing tools, such as dictionary approaches or topic modeling. While the clear advantage of these approaches is the simplicity and transparency of the index construction, they lack the capacity to generate human-like beliefs, especially to distinguish between different moments of a forecast distribution.

As newspaper data source, I use the *Handelsblatt*, which is the largest business newspaper in Germany. I have access to all articles from 1986 to 2023. The *Handelsblatt* has been a standard daily newspaper source for managers for decades and is therefore a suitable newspaper for generating firm beliefs. I include all front-page articles in my analysis, along with 300 articles randomly sampled each month from the remaining articles. To generate beliefs, I use the API of OpenAI’s GPT-3.5 model (Brown et al., 2020), a state-of-the-art large language model that also powered the well-known chatbot “ChatGPT”. I feed the *Handelsblatt* articles as part of a prompt to the API, which has the following structure: First, I tell the model to take on the role of a manager predicting the future of the German economy based on the information in the article. Then, I provide the title and textbody of the article. Afterwards, I ask for three forecasts, given the information of the article:

1. *Assessment of strong recession:*

Does the likelihood of a strong recession (Gross Domestic Product declines of at least 5%) in Germany over the next 12 months increase or decrease? [increase, decrease, neutral]

2. *Assessment of strong upturn:*

Does the likelihood of a strong upturn (Gross Domestic Product increases by at least 3%) in Germany over the next 12 months increase or decrease? [increase, decrease, neutral]

3. *Assessment on average:*

In the most likely scenario, how will the economic development in Germany evolve in the next 12 months? [negative, positive, neutral]

Then, I provide a structured format in which way GPT should generate the response. I classify decrease/negative as -1, increase/positive as +1, and neutral as 0; implying that the article contains no clear signal for the forecast at hand. Additionally, I ask for the confidence in the answer and a short explanation for each forecast to reduce the likelihood of hallucination. Appendix D.9 provides the exact wording of the prompt.

The design of the macro tail risk elicitation (“assessment of strong recession”) aligns with the question I formulated for the ifo Business Survey, focusing on the probability of a substantial recession characterized by a decline in GDP of at least 5% . The upturn scenario (“assessment of strong upturn”) is designed in a similar fashion, capturing potential strong growth opportunities. However, I selected 3% instead of 5% growth to reflect a more realistic upside scenario, given that 5% GDP growth has only been achieved during the reunification in 1990/1991. The third macro scenario (“assessment on average”) is formulated in the same style as qualitative expectations are elicited in surveys, e.g. business expectations in the ifo Business Survey. To obtain a robust measure, the mode forecast (“most likely scenario”) is asked for a given article. After averaging over many articles, the resulting index provides a qualitative measure of the average macro forecast.


Finally, I derive three indices by averaging the article-level classifications at the monthly level for each of the three macro forecasts. Following Bybee (2023), I weight each article equally. To ensure replicability, I re-ran the classifications a second time. The correlation between the first and second run at the monthly level exceeds 0.97 for all three indices. For the subsequent analysis, I use the average of the two runs.

4.3.4.2 Macro tail risk expectations index

Figure 4.9 shows the macro tail risk expectations index (*MacroTail* index), the average macro forecast index (*MacroAvg* index), and the macro upside risk expectations index (*MacroUp* index) on a monthly basis over time. The shaded areas indicate recession periods defined by the German Council of Economic Experts. We observe a strong co-movement of the three indices over time, see Appendix Table D.6 for the correlation coefficients. In recessions, macro tail risk expectations are high, while average macro forecasts and macro upside risk expectations are low. This highlights the importance of disentangling macro tail risk expectations from other moments of the macro forecast distribution. The *MacroTail* index shows a slight upward trend: following the Great Recession, the index did not return to pre-2008 levels and has remained particularly elevated in recent years. The dashed lines indicate the three months during which I elicit macro tail risk expectations from firms in the survey. Consistent with the survey results, macro tail risk expectations are higher in October 2022 than in April and October 2023.

How do the indices correlate with conventional uncertainty measures? Appendix Table D.6 shows a positive correlation with the VIX for Germany (0.29) and the European Economic Policy Uncertainty Index (0.56). There is also positive correlation with the US VIX (0.36)

Figure 4.9 – Macroeconomic beliefs over time based on newspaper data



Notes: The figure shows macroeconomic belief indices based on articles from the *Handelsblatt* over the time period 1986-2023. The beliefs are generated by prompting forecasts with the large language model GPT 3.5, see Appendix D.9 for the wording of the prompt. The light blue lines indicate the respective raw monthly index and the bold dark blue lines a moving average over the last three months. The shaded areas show recessions as defined by the German Council of Economic Experts. The dashed lines indicate the three months during which I elicit the subjective probability of a macro tail event from firms in the survey. Panel (A) presents the macro tail risk expectations index (*MacroTail* index), panel (B) presents the average macro forecast index (*MacroAvg* index), and panel (C) presents the macro upside risk expectations index (*MacroUp* index).

and a US skewness index (0.11), both based on S&P 500 returns. I conclude that while the *MacroTail* index shows—reassuringly—correlations in the expected direction with other uncertainty proxies, the relatively low correlation coefficients suggest that the index captures unique variation in macro tail risk.

To gain an intuition about GPT’s classifications, I analyze which type of articles have the greatest impact on changes in macro expectations. Using a dictionary-based approach,

I group articles into the following topics: news about the aggregate economy, idiosyncratic news about firms, financial news, news about national politics, and news about international politics. News that affect macro tail risk expectations are mostly about the aggregate economy and financial topics, while idiosyncratic news are relatively more important for average macro forecasts, see Appendix Table D.7.


A potential concern is that GPT’s expectations may reflect future event knowledge that could bias the responses due to the strong overlap between GPT’s training period, which lasts until 2021, and my sample. However, two aspects speak against major biases. First, Bybee (2023) evaluate out-of-sample macro expectations and finds no significant difference in predictive power compared to in-sample results. Second, I explicitly include in the prompt that GPT should only consider the information from the article. An example demonstrating that this approach seems effective is an article from 2001 about the relationship between Putin and George W. Bush, including Putin saying, "The differences weighed far less than what unites us", see for more details in Appendix D.9. GPT classifies this article as reducing macro tail risk expectations while having no impact on average forecasts and upside risk expectations, despite NATO’s eastward expansion in later years (this possibility is even mentioned in the article) and Russia’s invasion and annexation of Crimea during the training period.

4.3.4.3 Dynamics of macro tail risk expectations and aggregate investment

The literature measuring various types of macro indices based on newspaper articles typically employs parsimonious time series models to analyze how changes in these indices predict macroeconomic dynamics (e.g., Baker, Bloom and Davis, 2016; Bybee et al., 2024; Caldara and Iacoviello, 2022). I use local projections (Jordà, 2005) to estimate the link between changes in macro tail risk expectations and subsequent macroeconomic dynamics, focusing on investment. To disentangle macro tail risk expectations from other macro moment expectations, I explicitly include the two additional indices for average macro forecasts and upside risk macro expectations. I estimate the following equation on a monthly level over a horizon of up to four years:

$$y_{t+h} = \alpha_h + \beta_{1,h} \text{MacroTail}_t + \beta_{2,h} \text{MacroAvg}_t + \beta_{3,h} \text{MacroUp}_t + \Gamma X_t + \epsilon_t, \quad (4.5)$$

Figure 4.10 – Investment dynamics after an increase in macro tail risk expectations (*MacroTail*)



Notes: This figure shows local projections of investment good production in response to an interquartile-range increase in the macro tail risk expectations index (*MacroTail* index) based on Equation 4.5. In Panel (A), it is controlled for the average macro forecast index (*MacroAvg* index) and two lags of the outcome variable, *MacroTail* index, and *MacroAvg* index. In Panel (B), it is additionally controlled for the macro upside risk expectations index (*MacroUp* index) and two lags of it. Confidence bands are shown at the 68% and 90% level based on lag augmented standard errors (Montiel Olea and Plagborg-Møller, 2021).


where X_t contains two lags of the outcome variable, the *MacroTail* index, *MacroAvg* index and *MacroUp* index. The main outcome variable is the monthly log of the seasonally adjusted production of investment goods index in Germany, which serves as a high-frequency proxy for aggregate investment developments.²²

Figure 4.10 shows the estimation results for an interquartile-range increase in the macro tail risk expectations index. The left panel presents the results from local projections controlling for the average macro forecast index, while the right panel includes additional controls for the macro upside risk expectations index. In both cases, an increase in macro tail risk expectations foreshadows a large decline in investment. The impact accumulates gradually over time and persistently depresses investment by up to 7%. An increase in the average macro forecast index predicts a positive short-lived increase in investment, while changes in upside risk macro expectations show no significant relation to subsequent investment dynamics, see Appendix Figure D.12.

To test the robustness of the results, I estimate several alternative specifications, as summarized in Appendix Figure D.13. I find that the results are robust to changing the

²²The results remain robust when controlling for 12 or 24 lags of the outcome variable.

Figure 4.11 – Impact of macro tail risk expectations (*MacroTail*) on investment: Split by change in dispersion



Notes: In the left panel, the change in macro tail risk index $MacroTail_t - MacroTail_{t-1}$ is related to the change in the macro upside risk index $MacroUp_t - MacroUp_{t-1}$ in a scatter plot. The right panel show the local projection results following an increase in macro tail risk, split by co-movement with changes in the macro upside risk index based on Equation (4.6). Confidence bands are shown at the 68% and 90% level based on lag augmented standard errors (Montiel Olea and Plagborg-Møller, 2021).

number of lags, taking out the Great Recession, including the European Economic Policy Uncertainty Index of Baker, Bloom and Davis (2016), and when controlling for the DAX and the VIX (proxied by the realized monthly volatility of daily stock market returns) as additional measures for first and second moment expectations.

Another way to isolate the impact of changes in macro tail risk from changes in the second moment (dispersion) is to differentiate changes in macro tail risk based on whether the macro upside risk index moves in the same direction, while still controlling for average macroeconomic forecasts. When an increase in macro tail risk is accompanied by an increase in the macro upside risk, both tails of the macroeconomic outlook widen, leading to an overall increase in the second moment. Conversely, if only the macro tail risk index rises without a corresponding increase in the macro upside risk, the impact is not confounded by concurrent changes in second moment expectations.

Panel (A) of Figure 4.11 shows a scatter plot of how month-to-month changes in macro tail risk are related to changes in macro upside risk. In most months, an increase in macro tail risk expectations does not coincide with an increase in macro upside risk expectations (orange shaded area, $SameDirection = 0$). However, there are also months in which both indices co-move (blue shaded area, $SameDirection = 1$). Panel (B) presents the

investment changes following an increase in macro tail risk, differentiated by whether there is co-movement with changes in the macro upside risk index:

$$y_{t+h} = \alpha_h + \beta_{1,h} \text{MacroTail}_t \times \text{SameDirection}_t + \beta_{2,h} \text{MacroAvg}_t + \text{SameDirection}_t + \Gamma X_t + \epsilon_t \quad (4.6)$$

First, we observe that an increase in macro tail risk expectations, even without a corresponding increase in macro upside risk expectations, has a negative effect on investment beyond the first moment (orange line). Second, when the increase in macro tail risk expectations is accompanied by an increase in macro upside risk expectations, the negative investment effect tends to be even larger. This is consistent with additional real option effects documented in the literature, e.g., Bloom (2009).

4.4 QUANTITATIVE HETEROGENEOUS FIRM MODEL

In this section, I extend the stylized model from Section 4.3.1 into a quantitative heterogeneous firm model, simulate macro tail risk news shocks, and analyze policy counterfactuals. The model builds on a standard heterogeneous firm model (Cooper and Haltiwanger, 2006), which also forms the core of more elaborate RBC models that contain a heterogeneous firm model block (e.g., Bloom et al. 2018). I abstract from labor as a factor of production and general equilibrium effects, while introducing ambiguous firm-specific exposure to the low aggregate productivity state—the macro tail event. I, thereby, propose a new framework that characterizes the stochastic environment for firms.

4.4.1 Model setup

The model is an augmented version of the stylized model in Section 4.3.1 with more realistic dynamics. The main additional features are (i) idiosyncratic productivity, (ii) capital only partially depreciates in a given period, and (iii) there are three types of capital adjustment costs.

Technology and productivity Firms, indexed by i , produce output with a decreasing returns to scale production function. Their productivity is the product of idiosyncratic productivity z and aggregate productivity A . The log of idiosyncratic productivity follows an AR(1) process. Aggregate productivity varies over time according to a three-state Markov chain: (i) a high productivity state A^h with low macro tail risk, (ii) a high productivity state A^h with high macro tail risk, and (iii) a low productivity state A^l

corresponding to the macro tail event. The stochastic process of aggregate productivity is not symmetric, but rather characterized by infrequent significant plunges in productivity, consistent with the view of economic fluctuations as described by the plucking model (Friedman, 1964, 1993).

Importantly, firms are differentially exposed to the low aggregate productivity state. The decline in aggregate productivity from the high to the low productivity state is multiplied by a firm-specific exposure parameter β_i , resulting in productivity

$$A_i^l = A^h - \beta_i(A^h - A^l). \quad (4.7)$$

Due to the parameter β_i , there is an increase in cross-sectional dispersion of productivity in a macro tail event. The increase in dispersion in a recession is a prominent feature of business cycle models with time-varying volatility (Bloom et al., 2018) or skewness (Salgado, Guvenen and Bloom, 2023) of idiosyncratic risk. In contrast to these papers, the increase in dispersion in my model is tied to firm exposure to the macro tail event. Importantly, firms have ex-ante heterogeneous beliefs about this exposure, which has significant policy implications, as discussed below.

Firms consider a set of exposure scenarios $\beta_i \in \{\beta_{worst}, \beta_{average}, \beta_{best}\}$, but they are unable to assign probabilities to these scenarios, making the firm-specific exposure ambiguous, which is an instance of Knightian uncertainty (Keynes, 1921; Knight, 1921). There are three key advantages of this modelling assumption. First, it provides a parsimonious approach to generate large effects of macro tail risk on firms' investment decisions operating beyond the first-moment effect, as discussed in Section 4.3.1. Second, it accounts for the empirical evidence of substantial uncertainty about firm exposure, documented in Section 4.2.3. Third, the Knightian uncertainty is state dependent. Hence, when macro tail risk is very low, such as in the aggregate state (i), we converge back to a standard firm model.

Capital adjustment costs Firms face convex and non-convex capital adjustment costs. Investment is given by

$$I_{i,t} = k_{i,t+1} - (1 - \delta)k_{i,t} \quad (4.8)$$

Capital depreciates with δ every period. Capital adjustment costs consist of (i) fixed disruption cost ϕ_1 for non-zero investment, (ii) quadratic adjustment cost ϕ_2 , and (iii) resale cost ϕ_3 for net disinvestment, as capital is partially irreversible.

$$\psi(k_{t+1}, k_t) = \phi_1 \mathbb{I}_{I_{i,t} \neq 0} + \frac{\phi_2}{2} \left(\frac{I_{i,t}}{k_{i,t}} \right)^2 + (1 - \phi_3) |I_{i,t}| \mathbb{I}_{I_{i,t} < 0} \quad (4.9)$$

Corporate income tax and additional costs To study fiscal policy implications, I introduce a corporate income tax τ on firms' profits. As it is typically the case in advanced economies, I assume that the tax rate is asymmetric: firms pay taxes on their profit but do not receive a transfer if they are making losses, see Fuest and Neumeier (2023) for a review on corporate income taxation. To achieve realistic profit margins in the cross-section and over time, I add fixed production costs γ_0 to capture overhead costs as well as production costs that are linear in capital γ_1 as a reduced form way to add labor costs.

Firm problem Firms are risk neutral and ambiguity averse, i.e. they prefer to know the probability of future events. Their dynamic problem consists of choosing the capital stock for the next period to maximize the present discounted value of future profits.

When evaluating different investment plans, firms draw the exposure to macro tail events from the set of firm-specific exposure scenarios that minimizes expected continuation utility.²³ Hence, the dynamic problem of the firm is based on the expected firm-specific worst-case exposure $\beta_{i,worst}$ to the macro tail event, reflecting firms' cautious behavior in the presence of Knightian uncertainty:

$$V(k, z, \beta, A) = \max_{k'} \left\{ \begin{array}{l} (1 - \tau)(Azk^\alpha - \delta k - \psi(k', k) - \gamma_0 - \gamma_1 k) \\ -k' + k + \frac{1}{1+r} \min_{\beta \in \{\beta_{worst}, \beta_{average}, \beta_{best}\}} \mathbb{E}[V(k', z', \beta, A')] \end{array} \right\} \quad (4.10)$$

4.4.2 Model results

Calibration I solve the firm problem numerically using value function iteration. I discipline the degree of ambiguity aversion with the elicited worst-case exposures from the Ifo Business Survey and solve the model for firms at the 10th, 25th, 50th, 75th, and 90th percentile of the cross-sectional exposure distribution, as well as for a firm with average exposure ($\beta = 1$). I parameterize the transition probabilities to the low aggregate productivity state guided by firms' subjective macro tail risk expectations from the survey. In the high productivity state with low macro tail risk, the quarterly probability of transitioning


²³Epstein and Schneider (2003b) axiomatize the recursive version of behavior under ambiguity.

to the low productivity state is 1%, and in the one with the high macro tail risk it is 4%, equivalent to an 18% and 56% probability, respectively, over a five-year window. Hence, the change in the likelihood from the low to high macro tail risk state is slightly larger than the change in the average subjective probability reported in the Ifo Business Survey between October 2022 and April 2023, see Section 4.2. Productivity of the median firm decreases by 7.5% in the realized macro tail risk state, which lasts for five quarters on average (length of the Great Recession in Germany). The rest of the calibration assumes mostly standard parameters from the literature on a quarterly level, see Appendix D.7, where I also analyze the sensitivity of the results to parameter values.

Simulation of macro tail risk news shock Next, I simulate a macro tail risk news shock. For 60 quarters the economies are in the high productivity, low macro tail risk state. Then, there is a macro tail risk news shock: the economies are still in the high productivity state, but now face a higher probability of transitioning to the low productivity state in the future. From the next period on, economies switch back to the low macro tail risk state with a certain probability, but simulations where the macro tail event is realized are excluded.

Figure 4.12 shows how investment and the capital stock evolve over time after the macro tail risk news shock at event time zero. I plot the investment responses depending on firms' expected worst-case exposure. As macro tail risk increases, firms reduce investment sharply—higher macro tail risk expectations induce a recession. There is significant dispersion in the investment response across varying worst-case exposure beliefs. Given the left-skewness of the worst-case exposure distribution, the resulting cross-sectional changes following the macro tail risk news shock exhibit a fat left tail. Firms with below-median expected worst-case exposure (25th percentile of the exposure distribution) reduce investment by approximately 8% on impact, while those anticipating disproportionate impacts in the worst-case scenario (75th percentile) cut investment by over 25% on impact. The average decline in investment over the first four quarters after the shock amounts to around 10% for the median firm. This is slightly larger but in the ballpark of the effect sizes estimated using the first empirical approach. Based on a back-of-the-envelope calculation—see Appendix D.8 for details—my empirical estimates predict a decline between 6%–9% in investment.

Figure 4.12 – Baseline simulation of a macro tail risk news shock



Notes: The figure shows the average simulation results over 1000 economies, each populated with 4000 firms, for a given worst-case exposure to the low productivity state. For 60 quarters, the economies are in the high productivity, low macro tail risk state. At time 0, there is a macro tail risk news shock: the economies are now in the high productivity state (same productivity as before) with a higher probability of transitioning to the low productivity state in the future. Thereafter, the aggregate productivity process evolves naturally, but excludes the simulations in which the macro tail event is actually realized. Deviations of investment and capital stock from the high productivity state with low macro tail risk are shown in percentages. The different lines correspond to different expected worst-case exposures to the macro tail event state. For example “p75 worst: 4 × avg” represents the 75th percentile of the exposure distribution, where the expected decline in sales in the worst-case exposure scenario is 4 times larger than in the average exposure scenario of the median firm.


4.4.3 Policy implications

To what extent can counteracting monetary and fiscal policy stimulate investment following an increase in macro tail risk? Is there a difference between a world with and without ambiguous exposure to macro tail events? To study these questions, I conduct counterfactual analyses in which I introduce credible rule-based changes in monetary and fiscal policy in the case of a realized macro tail event. I repeat the simulations of macro tail risk news shock to analyze the effectiveness of the policies in stimulating aggregate investment. Note that the respective policy does not change in this case, since only the probability of transitioning to the macro tail event increases without a realization of it. Importantly, I assume that the policy changes are credible, since only then firms’ worst-case scenario in a macro tail event is improved.²⁴


²⁴Both policies considered, the lower interest rate and the switch to symmetric corporate income taxation, also lead to level shifts of the capital stock in the low macro tail risk state. I abstract from these level shifts and focus only on the relative change in the capital stock following an increase in macro tail risk.

Figure 4.13 – Macro tail risk news shocks and rule-based policy counterfactuals: Dynamics of the aggregate capital stock

Monetary policy counterfactual



Fiscal policy counterfactual



Notes: The figures show the aggregate change in capital after a macro tail risk news shock. Note that only macro tail risk increases without a realization of the macro tail event. The left (right) panels are based on a world with (without) ambiguous exposure to macro tail events. The upper panels consider a monetary policy counterfactual (dashed blue lines) in which the interest rate decreases in a realized macro tail event from 2% to 1%. The bottom panels consider a fiscal policy counterfactual (dashed blue lines) in which corporate taxation becomes symmetric in a realized macro tail event.

Monetary policy counterfactual For studying monetary policy, I consider a counterfactual in which the interest rate is no longer fixed. Specifically, the interest rate is rule-based and decreases in the case of a realized macro tail event. This counterfactual is motivated by the Taylor rule, which predicts that the central bank cuts the nominal interest rate after a decline in output. Since I abstract from nominal rigidities, I assume

that central banks would react disproportionately strongly to a potential decline in inflation, ensuring that the real interest rate also falls. Firms incorporate this interest rate cut in the case of a macro tail event into their optimization problem.

To what extent does the anticipated decline in real interest rates change firms' investment behavior after an increase in macro tail risk? I simulate a monetary policy counterfactual in which the annual interest rate decreases from 2% to 1% in the case of a macro tail event. The upper left panel in Figure 4.13 shows that the decline in the aggregate capital stock is mitigated by about 9–11% due to the credible interest rate cut by one percentage point. The stimulus works through expected lower user cost of capital in case of a realization of the macro tail event, leading to a higher optimal capital stock. The effect turns out to be similar in absolute terms in the world with and without ambiguous exposure. Importantly, this depends on the assumption that the user cost of capital is independent of firms' expected worst-case exposure to the macro tail event.

Fiscal policy counterfactual To study fiscal policy, I analyze a counterfactual in which there is a credible switch from asymmetric to symmetric corporate income taxation in case of a macro tail event. Intuitively, this means that firms receive money from the government when they incur losses. Importantly, the benefits of symmetric corporate taxes scale linearly with the size of the losses, thereby partially insuring firms. In practice, this is similar to generous loss carrybacks, where losses can be deducted from previously paid taxes on profits.²⁵

How does the introduction of symmetric corporate income taxation in the macro tail event affect firms investment behavior after an increase in macro tail risk? Again, as with monetary policy, I simulate a fiscal policy counterfactual in which there is symmetric corporate income taxation in case of a macro tail event. In a world with ambiguous exposure (lower left panel in Figure 4.13), the decline in the aggregate capital stock is mitigated by about 20–25%, whereas there is no large effect of the fiscal policy counterfactual in the world without ambiguous exposure (lower right panel in Figure 4.13). Why is there such a large difference? In contrast to monetary policy, the benefit of symmetric corporate taxation is directly linked to firms' worst-case exposure: the larger the expected losses are in the worst-case exposure scenario in the macro tail event, the more funds do the firms receive from the government. In the world without ambiguous exposure, only the average exposure scenario matters, with significantly better firm performance in the macro tail event compared to the worst-case scenario.

²⁵Note that this comparison only holds if firms were profitable and paid taxes before a macro tail event.

Note that the exposure of firms in a realized macro tail event would ultimately be the average exposure scenario rather than the worst-case scenario. Thus, the announced switch to symmetric taxation in a macro tail event is less costly to the government yet still significantly stimulates investment, thereby enhancing the efficiency of the policy compared to a scenario without ambiguity aversion. Why not implement symmetric corporate taxation universally? Especially in normal times, firms with low idiosyncratic productivity would benefit most, implying lower average productivity and higher costs of the policy.

4.5 CONCLUSION

This paper provides a comprehensive analysis of firms' macroeconomic tail risk expectations. I show that firms perceive macro tail events as likely to occur and that these expectations are highly dispersed across firms and strongly co-moving with firms' expectations about their own business. Based on three complementary empirical approaches, I consistently find that higher macro tail risk expectations have a strong negative impact on investment beyond first- and second-moment expectations. Ambiguity aversion regarding firm-specific exposure to macro tail events can rationalize the findings: the effects are strongest for firms that expect to be highly exposed to the macro tail event in the worst-case scenario.

I conduct policy counterfactuals in a quantitative heterogeneous firm model and show that fiscal policy is particularly effective in stimulating investment during periods of elevated macro tail risk. Crucially, policies must be credibly announced such that firms' expected worst-case scenario in a macro tail event is mitigated. Therefore, automatic stabilizers such as loss carrybacks are suitable policy measures, while uncertain payouts of subsidies in a macro tail event would not stimulate investment.

Further investigation into the origins of macro tail risk expectations and their role in various economic decisions made by different economic agents presents fertile ground for future research. Additionally, examining the efficiency of policies that can mitigate the negative effects of macro tail risk expectations, along with the evaluation of the importance of rule-based announcements, seems promising for future analysis.

A Appendix to Chapter 1

A.1 BUSINESS TAXATION IN GERMANY

In Germany, business profits are subject to two different taxes. At the national level, profits are either taxed under the personal income tax or under the corporate income tax, depending on the legal form of the firm. In addition, both corporate and non-corporate firms are subject to the local business tax (LBT) at the municipality level.

Corporate income tax Profits of incorporated firms are subject to the national corporate income tax (*Körperschaftsteuer*). The rate of the corporate income tax is currently 15 percent. Until 2000, a split rate imputation system existed in Germany, where retained profits were subject to a tax rate of 40-45 percent, whereas distributed profits were taxed at a rate of 30 percent. From 2001 to 2007, all profits were equally taxed at 25 percent. In all years since 1991, a so-called solidarity surcharge (*Solidaritätszuschlag*) of 5.5 percent of the corporate tax rate was added, dedicated to financing the costs of the German reunification.

Personal income tax Profits of non-corporated firms are subject to the progressive income tax (*Einkommensteuer*). The top marginal tax rate of the personal income tax is currently 45 percent but has been higher in the past, with a maximum of 56 percent in the 1980s. Since 2001, sole proprietors and partners in a partnership have been able to partially offset LBT payments tax against their income tax. This regulation, limiting the bite of the LBT, is however not relevant in our setting, as it only applies to unincorporated businesses, whereas we focus exclusively on the corporate sector.


Local business tax In addition, both corporate and non-corporate firms are subject to the LBT (*Gewerbesteuer*). As the corporate tax and the personal income tax, the LBT is a federal tax. For this reason, tax base and liability criteria of the LBT are set at the federal level. The tax rate, in turn, falls under the discretion of the municipalities. More

precisely, municipalities decide autonomously on a scaling factor that is then multiplied with a uniform basic tax rate. This results in the following formula:

$$\text{Local Business Tax Rate} = \text{Basic Federal Tax Rate} \times \text{Municipal Scaling Factor}$$

The basic rate, which is fixed at the national level, has been constant with exception to a change in 2008, when it was decreased from 5.0 to 3.5 percent. This means that for the median municipal scaling factor of 3.2, the resulting LBT rate was 16 percent before 2008. After 2008, the tax rate for the median scaling factor of 3.5 was 12.25 percent.

Figure A.1 – Timing of tax hike news



Notes: This figure provides evidence on the point in time when firms typically learn about a tax hike by displaying the number of monthly newspaper articles covering increases in the LBT, obtained from the German press database Genios. Under the broad definition, we counted search matches for “gewerbesteuer erhöh*”, under the narrow definition for “gewerbesteuer (erhöht* || angehob* || erhöhung) (beschl* || entschei*)”.

Each year, the municipal council has to vote on next year’s municipal scaling factor, even if it remains unchanged. The decision on next year’s local scaling factor is taken jointly with the adoption of the budget in the year’s last meeting of the municipal council. For this reason, tax hikes are typically announced in December. In Figure A.1, we substantiate this empirically, showing that newspaper coverage of municipal tax hikes in a given year indeed peaks in December. This holds for both a narrower definition (in red) and a

broader definition (in blue) of newspaper coverage of a hike in the LBT. As documented in Appendix A.2.1, a decision to increase the LBT sends no clear signal about the likelihood of future tax changes.

Around three quarters of the revenues of the LBT accrue directly to the municipalities, whereas one quarter is transferred to the federal government. Taxable profits of firms with establishments in more than one municipality are divided between municipalities according to formula apportionment based on the payroll share. As a consequence, profit shifting between municipalities requires the actual re-allocation of the employees (or wages) of a firm, and is thus associated with relatively high costs. The revenues from the LBT are of key importance for municipal budgets, as the LBT constitutes the most important original source of revenue for municipalities in Germany. Besides own tax revenues, municipal budgets are strongly dependent on fiscal transfers from the federal government or the federal states. As the municipalities cannot directly influence these fiscal transfers, the rate of the LBT is the central budget parameter under their control.

A.2 DATA APPENDIX

This appendix provides comprehensive information on the data sets used in the empirical analysis (including the translated wording of the relevant survey questions from the ifo Investment Survey), explains how we obtain our analysis sample, and reports summary statistics and aggregate time series of our final sample.

A.2.1 Administrative data at the municipality level

The administrative data on tax rates and municipality revenues and expenditures used in this paper cover the period from 1980 to 2018. The data largely correspond to the municipality data underlying the analysis in Fuest, Peichl and Siegloch (2018), comprising the period 1993 to 2018. Data for the period from 1980 to 1992 were obtained by filing individual requests to the respective Statistical Offices of the German Federal States. For the state of Schleswig-Holstein, data were not available in the year 1980. For Bremen and Saarland, data are only available since 1990. As these are the two smallest states of Germany in terms of GDP and population, jointly comprising less than 2% of the German population, this does not substantially change the composition of our sample. For all years, the data contain information on scaling factors of the LBT. In addition, we know the full municipality budget, that is all categories of expenditures and revenues, for most years. For a more detailed description of the data, we refer to Fuest, Peichl and Siegloch (2018) and Isphording et al. (2021).

There is substantial variation in LBT rates across municipalities and over time. To document this variation, we use the subset of municipalities, where we observe at least one firm during our sample period in the ifo Investment Survey. Figure A.2 plots the raw data of the local scaling factors for each municipality in Western Germany (excl. Berlin) over time, demonstrating that there is a lot of variation in local business taxes in any given year. Municipalities tend to increase the LBT approximately ten times as frequently as decreasing it. In consequence, the statistical power of this variation is too low to investigate the effect of tax drops in our data, and the analysis is thus restricted to tax hikes.¹ Accordingly, Figure A.3 shows that the share of municipalities that increased the LBT in a given year is relatively stable over time and does not differ between recessions and expansions. Moreover, Panel (A) of Figure A.4 plots the fraction of municipalities that

¹The number of tax decreases that could in principle be used in the analysis is very low. If we followed the protocol in Appendix A.2.3 to combine the municipality-level data on LBT rates and the firm-level data from the IVS, our analysis could only exploit 236 firm observations (0.7% of all firm-year observations) that face a tax drop in a given year despite spanning a time frame of almost four decades.


underwent a given number of tax hikes in the period between 1980 and 2018. The median municipality experienced three tax hikes, while taxes were never increased in only 7% of municipalities. The average duration between two tax hikes in our sample is 14.6 years, the median duration 13 years. Panel (B) displays the mean and various percentiles of the size of tax hikes over time. The distribution of tax hikes is rather stable over time in terms of average size and dispersion. If anything, tax hikes were slightly larger in the early 1980s and slightly lower in the 2010s.

To shed light on the dynamic aspects of tax hikes, Figure A.5 documents how a tax hike in year t_0 influences the probability for future tax hikes in the same municipality. Specifically, the figure displays the coefficients of separate regressions of the following form

$$TaxHike_{m,t+x} = \beta TaxHike_{m,t} + \mu_m + \epsilon_{m,t} \quad \forall x = \{1, 20\},$$


where $TaxHike_{m,t+x}$ is an indicator for a tax hike occurring x years after a tax hike in the same municipality m in year t that is estimated separately for each year in the future $x \in \{1, 20\}$. In the right panel, we include municipality fixed effects. The results show that tax hikes contain little predictive power for future tax hikes. While the unconditional probability for future tax hikes is slightly elevated if a tax hike has recently been enacted, the association is very weak and completely vanishes when including municipality fixed effects, which corresponds to the tax rate variation exploited in our main analyses (that applies firm fixed effects which are themselves nested within municipalities).

Figure A.2 – Time series of local scaling factors by municipality



Notes: This figure shows the local scaling factors underlying the LBT for each municipality in West Germany (excl. West-Berlin) over the period between 1980 and 2018. “Northern Germany” summarizes the states of Schleswig-Holstein, Hamburg, Bremen, and Lower Saxony.


Figure A.3 – Share of municipalities increasing the LBT over time




Notes: This figure shows the share of municipalities that increased the LBT in a given year. Gray shaded areas indicate recessions as defined by the German Council of Economic Experts.

Figure A.4 – Number of tax hikes and distribution of tax changes

(A) Number of tax hikes per municipality




(B) Distribution of tax hikes over time



Notes: Panel (A) plots the fraction of municipalities that underwent a given number of tax hikes in the period between 1980 and 2018. Panel (B) displays the average size of tax hikes (in percentage points) along with various distributional parameters, i.e., the median, the interquartile range, and the range between the 10th and 90th percentile of tax hikes in a given year.

Figure A.5 – Predictability of tax hikes as a function of past tax hikes in the same municipality



Notes: This figure reports how a tax hike in year t_0 influences the probability for future tax hikes in the same municipality, by showing the estimates of separate regressions with tax hike indicators x years in the future as dependent variable and a tax hike indicator for the current year as explanatory variable: $TaxHike_{m,t+x} = \beta TaxHike_{m,t} + \mu_m + \epsilon_{m,t} \forall x = \{1, 20\}$. In the right panel, we include municipality fixed effects, so that the graph shows the probability of future tax hikes conditional on knowing the institutional and political economy patterns of the own municipality.

A.2.2 ifo Investment Survey

General information The ifo Investment Survey (IVS, 2019) is a firm-level survey of the German manufacturing sector. Since its inception in 1955, it is conducted biannually by the ifo Institute, with survey waves in spring and fall of each year. The aim of the IVS is to supplement investment data collected by the German Statistical Office, which is only available with a time lag of two years, with more recent data by means of extrapolations at the industry level. The survey is part of the European Commission’s sponsored investment surveys in its member countries and participation in the survey is supported by several industry associations. All of this background information is contained in the cover letter of each survey. The aggregated investment volume of the participants of the IVS represents approximately 56% of overall investment in the manufacturing sector (see Sauer, Schasching and Wohlrabe (2023), p. 145).

The repeated panel structure of the ifo Investment Survey allows tracking approximately 1,500 firms over time. As outlined in greater detail below, the questionnaire elicits three types of questions, covering (i) the planned volume of investment, (ii) the realized volume of investment, and (iii) investment objectives. Realized investment is always reported for the previous year. Next to these investment-related variables, firms also report annual revenues and the number of employees. For all model specifications which include year fixed effects at the industry level, we rely on the ifo industry classification that maps firms into 34 industries over the entire sample period. The ifo industry classification is slightly more granular than, but largely comparable to two-digit NACE industries. All items of the questionnaire refer to the firms’ plants located in Germany. Sauer, Schasching and Wohlrabe (2023) provide a comprehensive overview and detailed description of this data source. After a protection period of one and a half years, the anonymized data can be accessed via the LMU-ifo Economics & Business Data Center under strict non-disclosure regulations (<https://www.ifo.de/en/ebdc>).

The survey is usually completed by high-level management personnel at the firms’ controlling departments (Sauer, Schasching and Wohlrabe, 2023). The ifo Institute incentivizes the participation to the survey by automatically providing the participants with the survey results free of charge as a thank-you for their cooperation. In order to create an additional incentive for participation in the investment survey, this reporting includes more detailed information, e.g., at more disaggregate sectoral levels, compared to the results that are reported publicly.

Table A.1 – Distribution of firms in the IVS by industry and size

WZ08	Industry	ifo Investment Survey				Actual Germany by			
		Small	Medium	Large	Total	Count	Employees	GVA	Payroll
10-12	Food, beverages, and tobacco	1.1	3.6	3.6	8.2	14.0	12.4	7.8	7.0
13-15	Textiles, apparel, and leather	1.2	1.8	1.0	4.1	4.2	1.8	1.1	1.1
16-18	Wood/paper products and printing	3.0	5.7	3.5	12.2	11.8	5.5	4.3	4.0
19	Coke and refined petroleum	-	-	0.3	0.6	0.0	0.3	0.8	0.5
20	Chemicals	-	1.1	3.4	4.7	1.5	4.6	6.9	6.0
21	Pharmaceuticals, medicinal chemical, and botanical	-	0.4	1.0	1.4	0.3	1.9	3.1	3.0
22+23	Rubber/plastic products, and other non-metallic	1.4	6.4	6.4	14.2	8.1	9.0	7.6	7.7
24+25	Basic and fabricated metal products	2.1	6.8	8.4	17.3	21.9	15.7	13.2	13.5
26	Computers, electronics, and optical products	-	1.0	2.4	3.6	3.7	4.8	5.4	5.5
27	Electrical equipment	-	1.3	3.8	5.3	2.9	6.4	7.0	7.5
28	Machinery and equipment	0.5	5.3	11.1	17.0	7.7	15.7	16.7	18.0
29+30	Transport equipment	-	0.8	4.0	4.9	1.9	13.2	19.0	19.1
31-33	Other, and installation of machinery and equipment	1.2	2.1	3.2	6.4	21.9	8.6	7.0	7.0
Total		11.4	36.5	52.2	100	100	100	100	100
Actual GER by Count		89.7	7.7	2.6	100				
Actual GER by Employees		19.1	18.6	62.3	100				
Actual GER by Gross Value Added (GVA)		10.6	13.2	76.1	100				
Actual GER by Payroll		10.0	13.9	76.1	100				

Notes: This table compares the distribution of firms in the ifo Investment Survey to the distribution of firms in administrative data by industry and firm size. The ifo Investment Survey data is based on the year 2018. The administrative data is based on the 2018 Statistics on Small and Medium-sized Enterprises (*“Statistik für kleine und mittlere Unternehmen”*) provided by the Federal Statistical Office (EVAS Code 48121). Definition of size classes: small: 0-49 employees; medium: 50-249 employees; large: 250+ employees. Cells are empty if there are less than 4 observations due to data protection.

Representativeness and accuracy In Table A.1, we demonstrate the representativeness of the ifo Investment Survey by comparing it to the distribution of firms in administrative data by industry and firm size. The numbers depicted in the table display the percentage share of firms in the respective cells. For instance, 17.3% of firms in the 2018 ifo Investment Survey are in the basic metals and fabricated metal products industry (2-digit WZ08: 24 and 25). This is in between the share of firms by count (21.9%) and weighted by employees (15.7%) in the administrative data. The share of firms by gross value added and payroll in this industry is around 13% in population. Overall, the industry-composition of the ifo Investment Survey is very close to the distribution in administrative data. Regarding the distribution across firm size, the ifo Investment Survey covers a substantial share in each size category. Around a third of firms have between 50 and 249 employees. Thereby, the survey slightly oversamples medium-sized firms while still being representative for small and large firms, since the share of firms is in between the population share of firms by count on the one hand, and by employees, gross value added, or payroll on the other hand.

In general, the accuracy of the IVS data appears to be quite high, as the average deviations of the survey results from the data of the Federal Statistical Office for the manufacturing sector as a whole are only relatively minor. For instance, Bachmann and Zorn (2020) show

that aggregate investment growth calculated from the microdata of the ifo Investment Survey is highly correlated with manufacturing investment growth reported by the Federal Statistical Office. Similarly, benchmarking the investment growth rates calculated from the survey against official statistics from the German Statistical Office for the period 1980 to 2016, Sauer, Schasching and Wohlrabe (2023) report an average absolute estimation error of less than two percentage points. Sauer, Schasching and Wohlrabe (2023) stress that it should be borne in mind that, at the time investments were recorded in the survey, the balance sheets of some of the companies may not yet be final, while the official results, on the other hand, are based on the final balance sheet figures.

Lastly, and in line with evidence presented in Appendix A.2.3, Bachmann, Elstner and Hristov (2017) present a series of stylized facts on the cross-sectional and time-series properties of revisions of investment plans, i.e., the difference between ex ante planned and ex post realized investment volumes, showing that these deviations are meaningful along many dimensions. For example, they document that the overall distribution of revisions is not systematically skewed, while their cross-sectional average is procyclical. This indicates that participants provide accurate investment plans given their current level of knowledge at the time of the survey.

Wording of questions in the IVS used in the paper In the following, we present the translated wording of the questions of the IVS that we use in the paper.

Fall questionnaire

1. General company information on the current financial year

Employees (as of Sept. 30th): _____ Total revenue (TEUR): _____

2. Gross fixed capital formation (equipment and buildings) in TEUR

	last year	this year	next year
Total (equipment + buildings):	_____	_____	_____

3. Investment targets this year and next year

Our domestic investment activity is influenced positively/negatively by the following factors:

		inducement		no	hampering	
		strong	little	influence	little	strong
This year:	a) Financing situation	☐	☐	☐	☐	☐
	b) ...					
Next year:	a) Financing situation	☐	☐	☐	☐	☐
	b) ...					

Spring questionnaire

1. General company information on the last financial year

Financial year from: _____ to: _____ Focus of production: _____
 Employees (as of Sept. 30th): _____ Total revenue (TEUR): _____

2. Gross fixed capital formation (equipment and buildings) in TEUR

	two years ago	last year	this year
Buildings:	_____	_____	_____
Equipment:	_____	_____	_____
Total (buildings + equipment):	_____	_____	_____

A.2.3 Construction and descriptive statistics of the merged dataset

Protocol for construction of merged dataset In constructing the final sample used for analysis, we have aimed at establishing a valid control group to analyze corporate tax hikes over time, and at cleaning the data to ensure that the results are not driven by outliers. To obtain our final sample, we follow the protocol outlined below:

- We restrict our sample to West Germany and, as Fuest, Peichl and Siegloch (2018), drop all municipalities which underwent municipal mergers in the observation period. As most of these municipalities were located in East Germany anyway, this does not substantially restrict our sample further (less than 1% of municipalities affected).
- We drop observations in a window of two years before and after a tax hike, if another tax hike occurred in that window.
- We drop all observations for which a tax decrease was enacted, as well as the two years before and after the tax decrease. Fuest, Peichl and Siegloch (2018) find that while tax hikes are arguably exogenous to shocks to economic variables, a potential endogeneity to economic conditions cannot be ruled out for tax decreases. In addition, only 13.5% of tax changes in the sample are tax decreases. In our setting, we do not have enough statistical power to separately analyze tax decreases.
- In total, the outlined sample selection above reduces the sample size from 8,522 municipalities and 326,274 municipality $\times$ year observations to 8,266 municipalities and 283,846 municipality $\times$ year observations.
- In the firm survey, for variables that are elicited both in the spring and the fall (last year's number of employees, revenues, and total investment volume), we follow Bachmann, Elstner and Hristov (2017) and compute a yearly value by taking the average. We drop the observation if both values deviate more than 20% from the mean.
- As Fuest, Peichl and Siegloch (2018), we drop firms with legal forms which are exempt from paying the LBT (this affects only 6.2% of the observations).
- We drop firms for which we observe revisions in investment plans in less than 5 years.

- To construct the Log Revision Ratio, we calculate the ratio of realized investments over planned investments, take the natural logarithm, and drop outliers (all values smaller/larger than $p1/p99$ in each year).
- Matching the municipal and firm-level samples, the final sample consists of 35,310 observations that are spread across 1,192 municipalities.
- We express all nominal variables, i.e., the amounts of revenues and investments, in real terms of constant 2015 Euro by converting German Mark to Euro and adjusting for inflation using the German Consumer Price Index (CPI).

Firms in the merged dataset: Descriptive statistics Table A.2 displays summary statistics for the firms in our sample. For each firm, we can rely on information on reported planned and realized investment volumes in 17 years on average. The median firm is a typical representative of the “German Mittelstand” employing 264 workers, generating annual revenues of 45 million Euro (CPI inflation-adjusted and—if denominated in German Mark—converted to 2015 Euros), and investing 1.4 million Euro each year. As described in Appendix A.2.2, the IVS covers firms of all sizes. While slightly oversampling medium-sized firms, it is still representative for small and large firms. Accordingly, 10% of firms in our sample have at most 38 employees, annual revenues of 5.2 million Euro and invest as little as 88,000 Euro per annum. In contrast, the 10% largest firms employ at least 1,950 workers and have annual revenues of almost half a billion Euro and total annual investment of at least 19 million Euro. As shown in Figure A.6, the firm size is consequently highly skewed according to the number of employees (Panel A), while the distribution of its logarithm displays a bell-shape (Panel B). As demonstrated in Panel A of Table A.3, firms experiencing a tax hike in year t are largely comparable to firms in the control group according to key firm characteristics measured in year $t - 1$. Further, the pre-treatment averages of both main outcome variables (Log Revision Ratio and Downward Revision Indicator measured in year t_{-1}) are not statistically different for firms that eventually are affected by a tax hike in year y_0 and firms ending up in the control group, see Panel B.

Documenting variation in investment over time, Figure A.7 displays a calendar time graph of the investment plans and investment realizations. Relatedly, Figure A.8 presents the share of downward revisions of investment (blue, solid) and the average log revision ratio (red, dashed) over time. The gray shaded areas indicate recession periods. During recessions, the share of downward revisions increases and the log revision ratio decreases. In addition, there might be a slight time trend towards a higher share of firms that revise

their investments downwards. Note, however, that this potential trend does not affect our analysis since we include year fixed effects in the regressions and thus rely on differences between firms in a given year for identification.

Figure A.9 shows the share of firms that report a decline in revenues by more than 10% compared to the previous year. In normal times, we observe that around 10% of firms experience such a revenue drop. In recessions, this share spikes up to 60%. This variable is used in Section 1.4.3, where we discuss potential channels of state-dependence in the effect of tax hikes on investment revisions.

Table A.2 – Summary statistics of firms in the sample

	p10	p50	p90	Mean
Employees	38	264	1,950	1,361
Revenues	5,194	44,901	451,899	418,842
Investment	88	1,435	19,163	17,751
Obs. per Firm	7	16	29	17


Notes: This table shows the 10th, 50th, and 90th percentile, and the mean of employees, revenues, and realized investment for the firms in our sample. “Obs. per Firm” refers to the number of years a firm is observed in our sample, i.e, the number of years for which firms report both ex ante planned and ex post realized volumes of investment. Revenues and investment are displayed in thousands of Euro.

Table A.3 – Balance statistics of firms in the treatment and control group

	Treated	Control	p-value
	(1)	(2)	(3)
	Mean _{<i>t</i>-1}	Mean _{<i>t</i>-1}	(1)=(2)
<i>Panel (A): Firm Characteristics</i>			
Employees	1,283	1,376	0.54
Revenues	337,895	434,071	0.14
Realized Investment	14,401	18,372	0.15
<i>Panel (B): Main Outcome Variables</i>			
Downward Revision of Investment Plans (Share)	0.54	0.53	0.34
Log Revision Ratio	-0.04	-0.03	0.59


Notes: This table presents balance statistics of firms experiencing a tax hike in year t_0 and untreated firms, respectively (both measured in the pre-event year t_{-1}). Panel (A) covers firm characteristics including the number of employees, annual revenues (in thousands of Euro), and realized investment (in thousands of Euro). Panel (B) refers to the main outcome variables, i.e., the Downward Revision Indicator and the Log Revision Ratio. Column (3) presents the p-values of a t-test on the equality of the means depicted in Columns (1) and (2).

Figure A.6 – Distribution of firms by number of employees




Notes: Panel (A) shows a histogram of the number of employees for the firms in our sample. The distribution is winsorized at a value of 4,000 employees. The vertical line denotes the median number of employees, which is 264. Panel (B) shows a histogram of the natural logarithm of the number of employees for the firms in our sample. The labels on the x-axis refer to the absolute number of employees.

Figure A.7 – Time series of investment plans and realizations




Notes: This figure shows time trends of log planned investment and log realized investment in the period 1980 to 2018, for all firms with a non-missing log revision ratio. The shaded areas indicate 95% point-wise confidence intervals. Gray shaded areas indicate recessions as defined by the German Council of Economic Experts.

Figure A.8 – Time series of investment revisions



Notes: This figure shows time series of the Log Revision Ratio (right axis), defined as the logarithm of the ratio between realized and planned investment, and the downward revision dummy (left axis), indicating whether a firm has invested less than planned, for the period 1980 to 2018 in our sample. Blue and red shaded areas indicate 95% point-wise confidence intervals. Gray shaded areas indicate recessions as defined by the German Council of Economic Experts.

Figure A.9 – Time series of share of large revenue drops




Notes: This figure depicts the time series of the share of firms with large revenue drops, defined as a year-to-year decline in revenues of more than 10%, over the period 1980 to 2018. Blue shaded areas indicate 95% point-wise confidence intervals, while gray shaded areas indicate recessions as defined by the German Council of Economic Experts.

Relationship between planned and realized investment Our identification approach relies on the investment plans of firms. In the following, we display the distribution of the investment revision ratio and illustrate the strong explanatory power of investment plans for actual investments.

Figure A.10 shows the distribution of the log revision ratio, trimmed at the first and 99th percentile. The log revision ratio is centered around zero, which means that on average, firms invest as much as they have previously planned. Overall, the approximately normal distribution in Figure A.10 indicates that firms revise investments frequently and similarly upwards and downwards.

Next, we provide further evidence that investment plans are highly informative for subsequently realized investment volumes. As shown in Figure 1.2 in the main part of the paper, the relationship between planned and realized investment volumes is highly linear and virtually corresponding to the 45 degree line. According to the corresponding regression output presented in Column (1) of Table A.4, 84% of the unconditional variation in (log) realized investment is explained by the investment plans for the respective year ($R^2 = 0.84$). The estimated slope is 0.91 and thus close to one. Moreover, the horse-race regression depicted in Column (2) demonstrates that planned investment regarding year t is much more strongly correlated with the ex post realizations in t than with realized levels in the previous year. As shown in Columns (3) and (4), these patterns even hold when controlling for firm fixed effects and investment plans are still strongly positively associated with ex post realized investment. Taken together, investment plans appear to contain accurate information on subsequent year's investment that goes beyond the extrapolation of the level of investment that was realized in the year these plans are reported to the IVS.

Figure A.10 – Distribution of the log revision ratio



Notes: This figure shows a histogram of the Log Revision Ratio in our sample. The Log Revision Ratio is defined as the logarithm of the ratio between realized and planned investment and constitutes one of the two main variables used in the analysis. For exhibitional reasons, the outliers below p1 and above p99 are not depicted here.


Table A.4 – Information content of investment plans for realized investment

	Log(Realized Investment)			
	(1)	(2)	(3)	(4)
Log(Planned Investment)	0.908 (0.005)	0.552 (0.011)	0.574 (0.012)	0.462 (0.012)
L.Log(Realized Investment)		0.395 (0.011)		0.195 (0.011)
Constant	1.276 (0.067)	0.731 (0.047)	6.064 (0.165)	4.886 (0.164)
Observations	25282	25282	25282	25282
R^2	0.84	0.87	0.89	0.89
R^2 (within)	-	-	0.27	0.30
Firm FE	-	-	✓	✓

Notes: This table reports estimates from linear regressions of log realized investment in year t_0 (I_{t0}) on log planned investment ($E_{t-1}(I_{t0})$) and log realized investment in the previous year (I_{t-1}). Columns (3) and (4) in addition purge for fixed effects at the firm-level. The sample is restricted to observations in years without changes in the LBT. Robust standard errors in parentheses.


A.3 SUPPLEMENTARY FIGURES AND TABLES

Figure A.11 – Long event study: Effect of tax hike on investment plans, realizations, and revisions




Notes: This figure shows event-study estimates of the downward revision (blue, solid line) and the log revision ratio (red, solid lines) in the left panel and log planned investment (green, short dashed lines) and log realized investment (orange, long dashed lines) in the right panel on the tax hike indicator and fixed effects at the levels of firm identifiers and years with longer event windows. The reference period is t_{-1} . “Log Revision Ratio” is the natural logarithm of the ratio between the ex post realized and ex ante planned volume of investment. In the right panel, the sample is trimmed outside the event window. Inspired by Dube et al. (2023), when estimating the effects with respect to log planned and realized investment, firms are assigned to another firm identifier after the year that is in the middle between two tax hikes in order to ensure that there is only one treatment for each unit and to allow for different long-run trends. In addition, end-periods $t-4$ and $t+4$ are binned in the right panel. The confidence intervals refer to the significance levels of 90% (thick lines) and 95% (thin lines).

**Figure A.12 – Investment revision effect after a tax hike:
Alternative estimators**




Notes: This figure shows the estimates of the imputation estimator introduced by Borusyak, Jaravel and Spiess (2024) (solid lines) and the interaction-weighted estimator by Sun and Abraham (2021) (dashed lines). The dependent variable is based on the ratio of realized investments over planned investments (elicited in fall the year before). “Downward Revision” is an indicator that is one if the ratio is below one (blue/circle). “Log Revision Ratio” is the log of the ratio (red/square). The treatment “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the previous year. Time fixed effects and firm fixed effects are absorbed in the estimation. Confidence bands refer to the 95% level.

Figure A.13 – Event study: Expenditures and revenues of municipalities




Notes: This figure shows the estimates of the following event-study regression: $Y_{m,t} = \sum_{j=-2}^2 \gamma_j \text{TaxHike}_{m,t}^j + \mu_i + \phi_{l,t} + \psi_{s,t} + \varepsilon_{i,t}$, where μ_i are firm fixed effects, $\psi_{s,t}$ year fixed effects at the industry level, and $\phi_{l,t}$ state-year fixed effects. In the left panel, $Y_{m,t}$ represents an indicator that is one when municipal revenues/spending increases compared to the previous year. In the right panel, $Y_{m,t}$ represents log municipal revenues/spending. The reference period is $t - 1$. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. Standard errors are clustered at the municipality level. The thick and thin confidence bands refer to the levels of 90% and 95%.

Figure A.14 – Investment revisions after a tax hike: Permutation test



Notes: This figure reports the empirical cumulative distribution functions of estimates from 2000 placebo tests. In a Monte Carlo exercise, tax hikes ($\mathbb{1}(\Delta tax_{m,t} > 0)$) are randomly allocated to municipalities by holding the share of treated municipalities constant. Then, Model (1) is estimated with the full set of fixed effects. In Panel (A), the dependent variable is $\mathbb{1}\left(\frac{I_{i,t}}{E_{i,t-1}(I_{i,t})} < 1\right)$, i.e., an indicator that is one if the fraction of realized investment over planned investment is below one. In Panel (B), the dependent variable is $\ln\left(\frac{I_{i,t}}{E_{i,t-1}(I_{i,t})}\right)$, i.e., the natural logarithm of the investment revision ratio. The vertical lines correspond to the baseline estimates from Column 5 in Panels A1 and B1 of Table 1.2. In Panel (A), 0.05% of the estimates are equal or larger than the baseline estimate. In Panel (B), 1.15% of the estimates are equal or smaller than the baseline estimate.


Figure A.15 – Collectively bargained wage growth in manufacturing




Notes: This figure shows year-on-year changes of the index of hourly earnings in the manufacturing sector without special payments obtained from the German Statistical Office. Grey shaded areas indicate recessions as defined by the German Council of Economic Experts.

Figure A.16 – Obstacles to investment by firm size

(A) Investment hampered by financing conditions



(B) Investment hampered by earnings situation



Notes: This binscatter plot depicts the share of firms reporting that their investment activity is strongly negatively affected by adverse financing conditions (Panel A) and the earnings situation (Panel B) by firm size, separately for recession and non-recession years. Recession years are defined as 1980-1982, 1992-1993, 2001-2003, and 2008-2009, as classified by the German Council of Economic Experts. Panel A uses the same survey question as Table A.11 in which firms report how financial constraints influence their current year's investment activity on a scale between 1 (strongly induced) and 5 (strongly hampered), see Appendix A.2 for the exact wording. Panel B uses the answer category "role of earnings situation" of the same survey question.

Table A.5 – Robustness: Baseline estimates excl. reunification period

	(1)	(2)	(3)	(4)	(5)
<i>Panel (A): Downward Revision</i>					
A1: Tax Hike Indicator: $\mathbb{1}(\Delta tax_{m,t} > 0)$					
	0.026	0.031	0.033	0.039	0.049
	(0.013)	(0.013)	(0.014)	(0.014)	(0.014)
Constant	0.540	0.539	0.539	0.539	0.538
	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
A2: Tax Hike in Percentage Points: $\Delta tax_{m,t}$					
	0.008	0.019	0.022	0.030	0.038
	(0.012)	(0.012)	(0.012)	(0.012)	(0.013)
Constant	0.541	0.540	0.540	0.540	0.539
	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
Observations	25960	25960	25911	25911	25911
<i>Panel (B): Log Revision Ratio</i>					
B1: Tax Hike Indicator: $\mathbb{1}(\Delta tax_{m,t} > 0)$					
	-0.039	-0.049	-0.035	-0.046	-0.062
	(0.020)	(0.019)	(0.020)	(0.020)	(0.022)
Constant	-0.039	-0.038	-0.039	-0.038	-0.037
	(0.007)	(0.007)	(0.001)	(0.001)	(0.001)
B2: Tax Hike in Percentage Points: $\Delta tax_{m,t}$					
	-0.031	-0.051	-0.047	-0.062	-0.073
	(0.017)	(0.016)	(0.018)	(0.018)	(0.021)
Constant	-0.039	-0.038	-0.038	-0.037	-0.037
	(0.007)	(0.007)	(0.001)	(0.001)	(0.001)
Observations	25310	25310	25255	25255	25255
Firm FE	-	-	✓	✓	✓
Year FE	-	✓	-	✓	-
Year × State FE	-	-	-	-	✓
Year × Industry FE	-	-	-	-	✓

Notes: This table re-estimates our baseline results from Table 1.2, excluding the years between the reunification of Germany in 1990 and the end of the government of Helmut Kohl in 1998, i.e., a period when many subsidy programs for investment, especially in East Germany, were in place that might have influenced investment decisions of West German firms. “Downward Revision” is an indicator that is one if the fraction of realized investment over planned investment is below one. “Log Revision Ratio” is the natural logarithm of this ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. Industry fixed effects refer to the ifo industry classification, comparable to two-digit NACE industries. Standard errors in parentheses are clustered at the municipality level.

Table A.6 – Treatment effect heterogeneity: State dependence

	Downward Revision				Log Revision Ratio			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel (A): Baseline Recession Definition by the German Council of Economic Experts</i>								
Tax Hike Indicator ×								
No Recession	0.018 (0.012)	0.021 (0.013)			-0.011 (0.018)	-0.019 (0.020)		
Recession	0.062 (0.022)	0.069 (0.024)			-0.084 (0.034)	-0.086 (0.036)		
Tax Hike ×								
No Recession			0.013 (0.011)	0.015 (0.012)			-0.019 (0.016)	-0.024 (0.019)
Recession			0.037 (0.016)	0.043 (0.018)			-0.064 (0.027)	-0.063 (0.028)
Constant	0.536 (0.001)	0.535 (0.001)	0.536 (0.001)	0.536 (0.001)	-0.033 (0.001)	-0.032 (0.001)	-0.033 (0.001)	-0.032 (0.001)
H0: Coefficients Equal	0.069	0.074	0.201	0.195	0.059	0.105	0.16	0.266
<i>Panel (B): Alternative Recession Definition by Negative Year-on-Year Real GDP Growth</i>								
Tax Hike Indicator ×								
No Recession	0.019 (0.011)	0.020 (0.012)			-0.021 (0.017)	-0.024 (0.018)		
Recession	0.089 (0.028)	0.112 (0.030)			-0.079 (0.047)	-0.107 (0.051)		
Tax Hike ×								
No Recession			0.013 (0.010)	0.013 (0.011)			-0.024 (0.015)	-0.024 (0.017)
Recession			0.066 (0.024)	0.084 (0.026)			-0.089 (0.039)	-0.107 (0.042)
Constant	0.536 (0.001)	0.535 (0.001)	0.536 (0.001)	0.536 (0.001)	-0.033 (0.001)	-0.032 (0.001)	-0.033 (0.001)	-0.032 (0.001)
H0: Coefficients Equal	0.017	0.004	0.038	0.012	0.243	0.126	0.12	0.065
Observations	35310	35310	35310	35310	34421	34421	34421	34421
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	-	✓	-	✓	-	✓	-
Year × State FE	-	✓	-	✓	-	✓	-	✓
Year × Industry FE	-	✓	-	✓	-	✓	-	✓

Notes: This table reports estimates from linear regressions based on Equation (1.1), where the tax hike treatment is split into recession and non-recession years. In Panel (A), 1980-1982, 1992-1993, 2001-2003, and 2008-2009 are classified as recession years as defined by the German Council of Economic Experts. In Panel (B), 1982, 1993, 2002, 2003, and 2009 are classified as recession years as these years showed negative real GDP growth according to World Bank data: <https://data.worldbank.org/indicator/NY.GDP.MKTP.KD.ZG?locations=DE>). The dependent variable is based on the ratio of realized investments over planned investments (elicited in fall the year before). “Downward Revision” is an indicator that is one if the ratio is below one. “Log Revision Ratio” is the log of the ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. The p-values at the bottom of each panel indicate whether the coefficients are statistically different from each other. Standard errors in parentheses are clustered at the municipality level.

Table A.7 – Difference-in-differences: Investment revisions after a tax hike

	Log Revision Ratio				
	(1)	(2)	(3)	(4)	(5)
Net-of-Tax Change	1.849 (1.119)	2.617 (1.070)	2.217 (1.177)	2.764 (1.138)	3.032 (1.312)
Constant	-0.033 (0.006)	-0.033 (0.006)	-0.033 (0.001)	-0.033 (0.001)	-0.032 (0.001)
Observations	34421	34421	34421	34421	34421
Firm FE	-	-	✓	✓	✓
Year FE	-	✓	-	✓	-
Year × State FE	-	-	-	-	✓
Year × Industry FE	-	-	-	-	✓

Notes: This table reports estimates from linear regressions of the log revision ratio on the percent change in the net-of-tax rate, defined as $\log(1 - \tau_t) - \log(1 - \tau_{t-1})$, as main explanatory variable. Industry fixed effects refer to the ifo industry classification, comparable to two-digit NACE industries. Standard errors in parentheses are clustered at the municipality level.

Table A.8 – Treatment effect heterogeneity: Volatility of revenue growth

	Downward Revision				Log Revision Ratio			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Tax Hike Indicator ×								
Low Revenue Growth Volatility	0.029 (0.015)	0.034 (0.015)			-0.012 (0.021)	-0.020 (0.022)		
High Revenue Growth Volatility	0.029 (0.015)	0.032 (0.016)			-0.045 (0.023)	-0.050 (0.024)		
Tax Hike Indicator ×								
Low Revenue Growth Volatility			0.022 (0.013)	0.024 (0.014)			-0.016 (0.018)	-0.019 (0.020)
High Revenue Growth Volatility			0.022 (0.013)	0.025 (0.014)			-0.052 (0.020)	-0.054 (0.023)
Constant	0.536 (0.001)	0.535 (0.001)	0.536 (0.001)	0.536 (0.001)	-0.033 (0.001)	-0.033 (0.001)	-0.033 (0.001)	-0.033 (0.001)
H0: Coefficients Equal: p-value	0.993	0.912	0.987	0.967	0.275	0.331	0.18	0.227
Observations	35155	35151	35155	35151	34281	34277	34281	34277
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	-	✓	-	✓	-	✓	-
Year × State FE	-	✓	-	✓	-	✓	-	✓
Year × Industry FE	-	✓	-	✓	-	✓	-	✓

Notes: This table reports estimates from linear regressions based on Equation (1.1), where the tax hike treatment is split into firms with low and high revenue growth volatility (split at median of firm-level standard deviation in revenue growth elicited in the ifo Investment Survey). The dependent variable is based on the ratio of realized investments over planned investments (elicited in fall the year before). “Downward Revision” is an indicator that is one if the ratio is below one. “Log Revision Ratio” is the log of the ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. The p-values at the bottom of each panel indicate whether the coefficients are statistically different from each other. Standard errors in parentheses are clustered at the municipality level.

Table A.9 – Treatment effect heterogeneity: Current revenue growth I

	Downward Revision				Log Inv. Revision			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Tax Hike Indicator ×								
No Recession ×								
No Strong Revenue Drop	0.018 (0.013)	0.024 (0.013)			-0.013 (0.019)	-0.022 (0.021)		
Strong Revenue Drop	0.009 (0.034)	-0.002 (0.034)			0.010 (0.053)	0.009 (0.054)		
Recession ×								
No Strong Revenue Drop	0.072 (0.026)	0.080 (0.029)			-0.085 (0.040)	-0.088 (0.042)		
Strong Revenue Drop	0.032 (0.037)	0.037 (0.038)			-0.067 (0.064)	-0.066 (0.068)		
Strong Revenue Drop	0.105 (0.008)	0.094 (0.009)	0.105 (0.008)	0.094 (0.008)	-0.190 (0.014)	-0.172 (0.014)	-0.191 (0.014)	-0.173 (0.014)
Tax Hike ×								
No Recession ×								
No Strong Revenue Drop			0.014 (0.011)	0.017 (0.013)			-0.021 (0.017)	-0.027 (0.020)
Strong Revenue Drop			0.007 (0.030)	-0.006 (0.031)			-0.004 (0.046)	0.002 (0.048)
Recession ×								
No Strong Revenue Drop			0.047 (0.019)	0.054 (0.022)			-0.077 (0.031)	-0.077 (0.034)
Strong Revenue Drop			0.004 (0.031)	0.008 (0.033)			-0.019 (0.052)	-0.017 (0.056)
Constant	0.518 (0.002)	0.520 (0.002)	0.519 (0.002)	0.520 (0.002)	-0.002 (0.003)	-0.004 (0.003)	-0.002 (0.002)	-0.004 (0.002)
H0: Coefficients Equal: p-value	0.379	0.354	0.243	0.256	0.817	0.787	0.347	0.376
Observations	35138	35138	35138	35138	34257	34257	34257	34257
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	-	✓	-	✓	-	✓	-
Year × State FE	-	✓	-	✓	-	✓	-	✓
Year × Industry FE	-	✓	-	✓	-	✓	-	✓

Notes: This table reports estimates from linear regressions based on Equation (1.1), where the tax hike treatment effect is estimated separately for each combination of recession and non-recession years and indicators of strong and weak revenue drops. A strong revenue drop is defined as a decline in revenues by more than 10% compared to the previous year. The dependent variable is based on the ratio of realized investments over planned investments (elicited in fall the year before). “Downward Revision” is an indicator that is one if the ratio is below one. “Log Revision Ratio” is the log of the ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. 1980-1982, 1992-1993, 2001-2003, and 2008-2009 are classified as recession years as defined by the German Council of Economic Experts. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. The p-values at the bottom of each panel indicate whether the coefficients are statistically different from each other. Standard errors in parentheses are clustered at the municipality level.

Table A.10 – Treatment effect heterogeneity: Current revenue growth II

	Downward Revision				
	(1)	(2)	(3)	(4)	(5)
Tax Hike Indicator ×					
No Recession ×					
No Strong Revenue Drop	0.024 (0.012)	0.018 (0.013)	0.024 (0.013)	0.027 (0.017)	0.034 (0.018)
Strong Revenue Drop	0.001 (0.033)	0.009 (0.034)	-0.002 (0.034)	-0.035 (0.057)	-0.009 (0.061)
Recession ×					
No Strong Revenue Drop	0.058 (0.026)	0.072 (0.026)	0.080 (0.029)	0.078 (0.035)	0.078 (0.039)
Strong Revenue Drop	0.011 (0.034)	0.032 (0.037)	0.037 (0.038)	-0.036 (0.073)	-0.020 (0.076)
Strong Revenue Drop	0.122 (0.008)	0.105 (0.008)	0.094 (0.009)	0.087 (0.013)	0.072 (0.014)
Constant	0.515 (0.005)	0.518 (0.002)	0.520 (0.002)	0.497 (0.002)	0.498 (0.002)
<i>N</i>	35139	35138	35138	21255	21193
Firm FE	-	✓	✓	✓	✓
Year FE	✓	✓	-	✓	-
Year × State FE	-	-	✓	-	✓
Year × Industry FE	-	-	✓	-	✓
Exclude Labor Drop	-	-	-	Yes, > 5%	Yes, > 5%

Notes: This table reports estimates from linear regressions based on Equation (1.1), where the tax hike treatment effect is estimated separately for each combination of recession and non-recession years, as well as indicators of strong and weak revenue drop observations. A strong revenue drop is defined as a decline in revenue by more than 10% compared to the previous year. In Columns (4) and (5), we drop firm observations that have a decrease in employees by more than 5% compared to the previous year. “Downward Revision” is an indicator that is one if the ratio of realized investments over planned investments (elicited in fall the year before) is below one. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. 1980-1982, 1992-1993, 2001-2003, and 2008-2009 are classified as recession years as defined by the German Council of Economic Experts. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. Standard errors in parentheses are clustered at the municipality level.

Table A.11 – Treatment effect heterogeneity: Financial constraints

	Downward Revision				Log Inv. Revision			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Tax Hike Indicator ×								
No Recession ×								
No Fin. Constr.	0.025 (0.015)	0.024 (0.016)			-0.009 (0.022)	-0.011 (0.024)		
Fin. Constr.	-0.004 (0.055)	-0.017 (0.057)			-0.051 (0.116)	-0.045 (0.121)		
Recession ×								
No Fin. Constr.	0.024 (0.029)	0.039 (0.030)			-0.040 (0.044)	-0.061 (0.047)		
Fin. Constr.	0.126 (0.065)	0.153 (0.070)			-0.066 (0.125)	-0.095 (0.127)		
Fin. Constr.	0.113 (0.017)	0.111 (0.017)	0.115 (0.017)	0.113 (0.017)	-0.225 (0.031)	-0.214 (0.032)	-0.226 (0.031)	-0.214 (0.031)
Tax Hike ×								
No Recession ×								
No Fin. Constr.			0.026 (0.014)	0.021 (0.016)			-0.023 (0.023)	-0.018 (0.025)
Fin. Constr.			-0.025 (0.050)	-0.039 (0.053)			-0.048 (0.101)	-0.045 (0.108)
Recession ×								
No Fin. Constr.			0.013 (0.023)	0.024 (0.025)			-0.019 (0.036)	-0.030 (0.040)
Fin. Constr.			0.066 (0.064)	0.089 (0.077)			-0.056 (0.107)	-0.066 (0.116)
Constant	0.550 (0.001)	0.550 (0.001)	0.551 (0.001)	0.551 (0.001)	-0.054 (0.002)	-0.054 (0.002)	-0.054 (0.002)	-0.054 (0.002)
H0: Coefficients Equal: p-value	0.168	0.138	0.449	0.426	0.849	0.802	0.746	0.767
Observations	23661	23640	23661	23640	23123	23101	23123	23101
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	-	✓	-	✓	-	✓	-
Year × State FE	-	✓	-	✓	-	✓	-	✓
Year × Industry FE	-	✓	-	✓	-	✓	-	✓

Notes: This table reports estimates from linear regressions based on Equation (1.1), where the tax hike treatment effect is estimated separately for each combination of recession and non-recession years, as well as indicators on whether the financing situation is reported to be a factor for a strong slowdown in investment volumes or not. To construct the financing indicator, we use a question from the fall survey (available since 1989), where firms rate on a scale from 1 (strong stimulus) to 5 (strong slowdown) different factors that influence investments in the current year, see Appendix A.2 for the exact wording. We construct an indicator that is one if a firm reports the highest category (5). The dependent variable is based on the ratio of realized investments over planned investments (elicited in fall the year before). “Downward Revision” is an indicator that is one if the ratio is below one. “Log Revision Ratio” is the log of the ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. 1980-1982, 1992-1993, 2001-2003, and 2008-2009 are classified as recession years as defined by the German Council of Economic Experts. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. The p-values at the bottom of each panel indicate whether the coefficients are statistically different from each other. Standard errors in parentheses are clustered at the municipality level.

**Table A.12 – Treatment effect heterogeneity:
Firm size and settlement structure**

	Downward Revision				Log Revision Ratio			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel (A): Heterogeneity by Firm Size</i>								
Tax Hike Indicator ×								
Small Firms	0.029 (0.017)	0.031 (0.017)			-0.031 (0.027)	-0.035 (0.028)		
Large Firms	0.028 (0.014)	0.034 (0.014)			-0.027 (0.017)	-0.036 (0.019)		
Tax Hike ×								
Small Firms			0.021 (0.015)	0.024 (0.016)			-0.040 (0.023)	-0.043 (0.025)
Large Firms			0.022 (0.012)	0.024 (0.013)			-0.030 (0.016)	-0.034 (0.018)
Constant	0.536 (0.001)	0.535 (0.001)	0.536 (0.001)	0.536 (0.001)	-0.033 (0.001)	-0.032 (0.001)	-0.033 (0.001)	-0.032 (0.001)
H0: Coefficients Equal	0.951	0.869	0.972	0.989	0.899	0.96	0.73	0.756
<i>Panel (B): Heterogeneity by Settlement Structure</i>								
Tax Hike Indicator ×								
Urban Area	0.027 (0.012)	0.030 (0.013)			-0.020 (0.018)	-0.028 (0.019)		
Rural Area	0.037 (0.023)	0.043 (0.025)			-0.069 (0.034)	-0.070 (0.037)		
Tax Hike ×								
Urban Area			0.019 (0.011)	0.022 (0.012)			-0.023 (0.016)	-0.027 (0.019)
Rural Area			0.029 (0.017)	0.032 (0.019)			-0.072 (0.026)	-0.072 (0.028)
Constant	0.536 (0.001)	0.535 (0.001)	0.536 (0.001)	0.536 (0.001)	-0.033 (0.001)	-0.032 (0.001)	-0.033 (0.001)	-0.032 (0.001)
H0: Coefficients Equal	0.688	0.641	0.64	0.649	0.199	0.314	0.106	0.184
Observations	35310	35310	35310	35310	34421	34421	34421	34421
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	-	✓	-	✓	-	✓	-
Year × State FE	-	✓	-	✓	-	✓	-	✓
Year × Industry FE	-	✓	-	✓	-	✓	-	✓

Notes: This table reports estimates from linear regressions based on Equation (1.1). In Panel (A), the tax hike treatment is split into small (< 250 employees) and large (≥ 250 employees) firms. In Panel (B) the treatment variables are interacted with indicators of urban and rural areas following the classification of the Federal Institute for Research on Building, Urban Affairs and Spatial Development (BBSR) that is mainly based on population density. The dependent variable is based on the ratio of realized investments over planned investments (elicited in fall the year before). “Downward Revision” is an indicator that is one if the ratio is below one. “Log Revision Ratio” is the log of the ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. The p-values at the bottom of each panel indicate whether the coefficients are statistically different from each other. Standard errors in parentheses are clustered at the municipality level.

Table A.13 – Treatment effect heterogeneity: Tax hike dynamics

	Downward Revision				Log Revision Ratio			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel (A): Heterogeneity by the Frequency of Tax Hikes</i>								
Tax Hike Indicator ×								
Few Tax Hikes	0.024 (0.015)	0.028 (0.016)			-0.014 (0.026)	-0.020 (0.029)		
Many Tax Hikes	0.031 (0.015)	0.036 (0.015)			-0.038 (0.020)	-0.046 (0.021)		
Tax Hike ×								
Few Tax Hikes			0.021 (0.013)	0.021 (0.014)			-0.030 (0.019)	-0.029 (0.022)
Many Tax Hikes			0.021 (0.013)	0.027 (0.014)			-0.038 (0.019)	-0.046 (0.022)
Constant	0.536 (0.001)	0.535 (0.001)	0.536 (0.001)	0.536 (0.001)	-0.033 (0.001)	-0.032 (0.001)	-0.033 (0.001)	-0.032 (0.001)
H0: Coefficients Equal:								
p-value	0.721	0.733	0.998	0.775	0.464	0.468	0.762	0.584
Observations	35310	35310	35310	35310	34421	34421	34421	34421
<i>Panel (B): Heterogeneity by Occurrence of a Tax Hike in the Last 5 Years</i>								
Tax Hike Indicator ×								
≥ 1 Hike in Last 5 Years	0.039 (0.019)	0.052 (0.021)			-0.050 (0.026)	-0.067 (0.029)		
No Hike in Last 5 Years	0.019 (0.013)	0.020 (0.013)			-0.012 (0.021)	-0.014 (0.022)		
Tax Hike ×								
≥ 1 Hike in Last 5 Years			0.030 (0.016)	0.044 (0.019)			-0.043 (0.022)	-0.062 (0.027)
No Hike in Last 5 Years			0.013 (0.011)	0.013 (0.012)			-0.023 (0.017)	-0.021 (0.019)
Constant	0.540 (0.001)	0.540 (0.001)	0.541 (0.001)	0.541 (0.001)	-0.040 (0.001)	-0.039 (0.001)	-0.039 (0.001)	-0.039 (0.001)
H0: Coefficients Equal:								
p-value	0.358	0.195	0.386	0.155	0.257	0.145	0.489	0.204
Observations	33220	33201	33220	33201	32375	32356	32375	32356
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	-	✓	-	✓	-	✓	-
Year × State FE	-	✓	-	✓	-	✓	-	✓
Year × Industry FE	-	✓	-	✓	-	✓	-	✓

Notes: This table reports estimates from linear regressions based on Equation 1.1. In Panel (A), the tax hike treatment variable is interacted with dummies splitting the sample into municipalities with few (≤ 3) and many (> 3) tax hikes over the entire sample period. In Panel (B), the tax hike treatment is split into cases where at least one tax hike has already occurred in the previous five years and where no tax hike occurred in the previous five years in the respective municipality. The dependent variable is based on the ratio of realized investments over planned investments (elicited in fall the year before). “Downward Revision” is an indicator that is one if the ratio is below one. “Log Revision Ratio” is the log of the ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. Standard errors in parentheses are clustered at the municipality level.

A.4 BACK-OF-THE-ENVELOPE CALCULATION

In the following, we present the assumptions underlying the back-of-the-envelope calculation used to approximate the investment loss for each additional Euro of tax revenue.

The median firm in our sample generates yearly revenues of 45 million Euro. Among the subsample of firms that can be linked to information on the cashflow/revenue ratio balance sheet data, the median profit margin is 4.4%. Assuming that this figure corresponds to all firms in the sample, this translates into 1.98 million Euro of aggregate profits. A one percentage point increase in the LBT increases the tax burden of the median firm—and thus overall tax revenues—by 19,800 Euro. Moreover, the median investment-revenue ratio amounts to 3% in the microdata of the ifo Investment Survey. Hence, the median firm invests approximately 1.4 million Euro each year. Given the estimated semi-elasticity of 3 (see Section 1.4), a one percentage point increase in the LBT is associated with decreased investment of the median firm by roughly 42,000 Euro. Finally, dividing 42,000 by 19,800 gives that 2.12 Euro of investment volume is lost for each additional Euro of tax revenue. In crisis years, we estimate a semi-elasticity of investments with respect to the LBT rate of 6. Assuming that the relation between the profit margin and investment-revenue ratio is the same in a recession, investments even decrease by 4.24 Euro for each additional Euro of tax revenue.²

It is furthermore necessary to also take the (long-term) behavioral response of firms into account: as tax increases decrease firm investment, future firm profits should also be reduced, resulting in lower tax revenues of the municipalities. Unfortunately, we cannot estimate the elasticity of firms' profits with respect to changes in investment based on our data. We circumvent this constraint by separately calculating the behavioral response for reasonable lower and upper bounds of this elasticity, i.e., assuming that foregone investment maps into foregone future profits with half of the median profit margin (2.2%) or with five times the median profit margin (22%). For the median firm, which lowered investment by 42,000 Euro, this translates into lower profits between 924 Euro and 9,240 Euro. As the average LBT rate is approximately 15%, this leads to an additional reduction in tax revenues between 139 Euro and 1,386 Euro. Taken together, we approximate that incorporating the behavioral response increases the investment loss for each additional Euro of tax revenue from 2.12 Euro to an estimate in the range between 2.14 Euro ($42000/(19800-139)$) and 2.28 Euro ($42000/(19800-1386)$).

²In fact, the profit margin decreases slightly more than the investment-revenue ratio in recessions. Incorporating this relation in the calculation would lead to an even higher loss of investments for each additional Euro of tax revenue in recessions.

From this approximation of the behavioral response, we can also derive the marginal value of public funds (MVPF) in the spirit of Hendren and Sprung-Keyser (2020), given as:

$$MVPF = \frac{\text{Beneficiaries' Willingness to Pay}}{\text{Net Cost to Government}}$$

In our setting, firms are the beneficiaries and their willingness to pay is equal to the change of the tax burden. The net cost of the government equals the change of tax revenues plus the additional revenue changes via the behavioral response. According to this, our estimates point at a MVPF in the range between $(19,800)/(19,800 - 139) = 1.01$ and $1.08 = (19,800)/(19,800 - 1,386)$, i.e., slightly above one.

A.5 CALCULATION OF EFFECTIVE TAX RATES

This appendix describes how we calculate effective tax rates used in the alternative specification presented in Section 1.4.2 based on the statutory LBT rates used in the main specification. When talking about effective tax rates, we always refer to marginal—rather than average—tax rates.³ Nicodème (2001) provides a helpful overview on different approaches of computing effective tax rates. Our procedure is guided by the classic framework of Hall and Jorgenson (1967), as, e.g., recently applied by Furno (2022). The key difference between the two concepts is that while the statutory tax rate is the one imposed by law on taxable profits, the effective tax rate is the percentage of profits actually paid by a company after taking into account deductions including depreciation of assets, exemptions, tax credits, and preferential rates. Note that in the case of the German case, the most important feature are depreciation rules (denoted with z below) while the other components play a negligible role. To compute effective tax rates in the setting of the LBT, we proceed as follows:

- We first obtain depreciation schedules separately for machinery m and buildings b , the two main types of investment for which different depreciation rules apply. To do so, we use information from the Oxford Corporate Tax Database (<https://oxfordtax.sbs.ox.ac.uk/cbt-tax-database>). Indeed, over our sample period of almost 40 years, the depreciation rules have changed repeatedly.
- To illustrate this change over time, we calculate the present discounted value (PDV) of a depreciation, denoted by z_m (for machinery) and z_b (for buildings), respectively. As the choice of the adequate discount rate is not innocuous in our setting, we employ the following two different specifications, whose resulting z_m and z_b are depicted in Figure A.18:

I. We follow Zwick and Mahon (2017) and set the discount rate to 7%.

II. We use time-varying interest rates for discounting to accommodate for the fact that over the sample period the interest rates on firm loans have been declining substantially, from close to 10% in 1980 to less than 2% in recent years, with considerable variation in between.⁴ In contrast to most other studies in the

³Note that in the German context marginal and average tax rates are approximately the same, since the tax rate is flat and there are only a few exceptions, e.g., no tax credits.

⁴The time series of average interest rates on firm loans displayed in Figure A.17 builds on three different charts provided by the Deutsche Bundesbank (German Central Bank), as the effective interest rate for non-financial corporations is only available since 2003. The breaks are indicated by the dashed vertical lines. Over our entire sample period, the average interest rate according to this graph has been 5.1%.

literature that rely on a single or few tax reforms within shorter time periods, time-variation in interest rates may have large implications for the PDV of a depreciation in our analysis that covers a period of almost 40 years.

- Next, we calculate the combined depreciation schedule, z , for each firm, i.e., a weighted average of z_m and z_b based on firms' respective share of investment in machinery and buildings. However, as we do not observe these investment shares in machinery and buildings for each firm in all years of the survey, we have to impute these values (in some years). We consider two distinct specifications for this imputation:

I. In each year, we assign the average share of investment in machinery and buildings based on aggregate data from the Federal Statistical Office of Germany (only available since 1990, imputed for the years before). This way, the investment shares vary over time, but are the same for all firms in our sample in a given year. Across years, the average share of investment in machinery amounts to 88%.

II. We use the firm-specific share of investment in machinery and buildings reported to the ifo Investment Survey. As this information is provided less frequently to the IVS compared to the overall level of investment, we use the firm-specific mean across all years if firms reported machinery and building investments at least three times. To retain the sample size, we replace missing values by the values obtained from method I.

- Having obtained the depreciation schedule z , the effective tax rate is then given by

$$\tau_{eff} = 1 - \frac{1 - \tau}{1 - z * \tau},$$

which only depends on z and the statutory LBT rate τ in the German case as there are no relevant tax credits in the LBT that would complicate the calculation.

- Finally, we calculate the change in the effective tax rate if a tax hike takes place in a given year. Here, we set z in both years (t_0 and t_{-1}) equal to the value of the tax hike year. Thereby, we isolate the effect of tax changes by making the arguably reasonable assumption that firms know the z value of the next year when forming their investment plans.⁵

⁵Note that while German municipalities have the power and discretion to change τ , i.e. the LBT rate, they cannot change the depreciation rules z that are determined at the Federal level.

The results show that, across all specifications, the variation captured by changes in effective tax rates is strongly associated with the underlying changes in the LBT rate as plotted in Figure A.19.

In alternative specifications, we also express changes in the costs for investment in terms of the user cost of capital instead of effective tax rates. This only requires a simple transformation:


$$UserCost = \frac{1 - z * \tau}{1 - \tau}$$

Hence,

$$\tau_{eff} = 1 - UserCost^{-1},$$


which means that switching from effective tax rates to a user cost approach will not impact our results apart from rescaling the magnitude of the coefficients. That the user cost of capital yields virtually the same results as using effective tax rates is also visible in Figure A.20, which plots the change in the user cost of capital against the change in effective tax rates for all tax hikes in our sample.

Figure A.17 – Time series of average interest rate on loans for firms




Notes: This figure shows a time series of the average lending rate for firms. From 2003 onward, the effective interest rate for non-financial corporations is used. For the year 1997 to 2002, the effective interest rate to firms for loans between 500,000 and 5 million Euro is used and adjusted upwards (roughly 1 p.p.) to ensure a smooth transition in 2003. For the years 1980 to 1996, the discount rate of the Deutsche Bundesbank is used and adjusted upwards (roughly 4 p.p.) to ensure a smooth transition in 1997. The two dashed vertical lines indicate the breaks in the time series. Source: Deutsche Bundesbank.

Figure A.18 – Present discounted value of depreciation: 7% vs time-varying interest rate




Notes: This figure shows values of the present discounted value of depreciation for machinery (z_m) and buildings (z_b) in the period 1980 to 2018. Depreciation schedules are obtained from the Oxford Corporate Tax Database. The solid line assumes a time-constant discount rate of 7% following Zwick and Mahon (2017), the dashed line calculates the PDV based on the time-varying interest rate on firm loans as displayed in Figure A.17.

Figure A.19 – Relation of changes in LBT rate and changes in effective tax rates



Notes: For each tax hike in our sample, this figure plots its size in terms of an effective tax hike (τ_{eff} ; y-axis) against its size as a statutory tax hike (x-axis). As we do not observe the investment shares in machinery and buildings for each firm in all years of the survey, we must impute these values. We consider two distinct specifications in which τ_{eff} is either calculated based on the average share of investment in machinery and buildings based on aggregate data from the Federal Statistical Office of Germany (Specification “I”) or on the firm-specific share of investment in machinery and buildings reported to the ifo Investment Survey whenever available (Specification “II”). In the left panel, we assume a time-constant discount rate of 7% following Zwick and Mahon (2017), in the right panel we calculate the PDV based on the time-varying interest rate on firm loans as displayed in Figure A.17.

Figure A.20 – Relation of changes in effective tax rates and changes in user cost of capital




Notes: This figure plots the change in the user cost of capital (multiplied with 100) against the change in effective tax rates for all tax hikes in our sample, assuming that the share of investment allocated to machinery and buildings is constant across firms, but varying over time (Specification I). In the left panel, we assume a time-constant discount rate of 7% following Zwick and Mahon (2017), in the right panel we calculate the PDV based on the time-varying interest rate on firm loans as displayed in Figure A.17.

B Appendix to Chapter 2

B.1 IFO BUSINESS SURVEY

B.1.1 Survey timing and inflation dynamics

Figure B.1 – Producer and consumer prices over time



Notes: German manufacturing sector's producer price index (PPI) and consumer price index (excl. energy, food). The vertical lines correspond to the dates of the supplementary pass-through questions in the ifo Institute's business survey: June 2021, April 2022, and October 2022. Source: Federal Reserve Economic Data (FRED).

B.1.2 Supplementary questions (translated)

Pass-through

- June 2021: "To what extent can you pass through higher purchase prices to your customers? We can pass through the higher purchase prices to our customers by ____%."
- April 2022: "To what extent can you pass through higher purchase prices for energy, raw materials, and intermediate inputs to your customers? We can pass through the higher purchase prices to our customers by ____%."

- October 2022: “To what extent have you so far passed through higher purchase prices for energy, raw materials, and intermediate inputs to your customers? ___%”
- October 2022: “To what extent do you plan to pass through higher purchase prices for energy, raw materials, and intermediate inputs to your customers in the next six months (including previous price changes)? ___%”

Limiting factors for pass-through

- What factors limit the flexibility of your firm in pricing? (Please check all that apply)
 - Competitive pressure
 - Weak demand / willingness to pay
 - Long-term contractual commitments
 - Administrative effort
 - Regulation / Pricing not under the firm’s control
 - Other, namely:

Duration of supply shortages

- October 2021: “If your production activity is currently limited by a shortage of raw materials or intermediate inputs: How long do you expect these supply shortages will persist?” ___ months
- May 2022: “If your production activity is currently limited by a shortage of raw materials or intermediate inputs: How long do you expect these supply shortages will persist?”___ months

Change in realized prices

- “The prices of our goods/services increased/stayed the same/decreased compared to [previous month].”

Change in revenues/demand

- “The demand situation has improved/not changed/deteriorated compared to [previous month].”

Range of orders

- “Our range of orders currently corresponds to an average production of ___ month(s).”

Capacity utilization

- “The utilization of our machines (normal full utilization = 100%) is currently” — % [tick box from 30% to 100 % in 5/10 p.p. steps, or enter value manually if larger than 100%]

Change in orders

- “Our range of orders increased/remained the same/decreased compared to [previous month].”

Production expectations for the next three months

- “Our production activity is expected to increase/remain roughly the same/decrease.”

Energy intensity

- “What percentage share of revenue do you estimate your company spent on energy costs in 2021 (energy intensity)” — %

Business state, expectation, and uncertainty

- “We assess our current business state as:” [continuous slider from 0 (bad) over 50 (satisfactory) to 100 (good)]
- “Expectations about the next six months: our expected business state in economic terms:” [continuous slider from 0 (rather worse) over 50 (the same) to 100 (rather better)]
- “We estimate the uncertainty regarding our business expectations in the next six months as:” [continuous slider from 0 (low) over 50 (average) to 100 (high)] — %

B.1.3 Supplementary questions (original)

Pass-through

- June 2021: “Zu welchem Grad können Sie die höheren Einkaufspreise an Ihre Kunden weitergeben? Wir können die höheren Einkaufspreise zu — % an unsere Kunden weitergeben.”
- April 2022: “Zu welchem Grad können Sie höhere Einkaufspreise für Energie, Rohstoffe und Vormaterialien an Ihre Kunden weitergeben? Wir können die höheren Einkaufspreise zu — % an unsere Kunden weitergeben.”
- October 2022: “Zu welchem Grad haben Sie die höheren Einkaufspreise für Energie, Rohstoffe und Vormaterialien bereits an Ihre Kunden weitergegeben?” —%
- October 2022: “Zu welchem Grad planen Sie, in den kommenden 6 Monaten die höheren Einkaufspreise für Energie, Rohstoffe und Vormaterialien (inklusive bisheriger Preisanpassungen) an Ihre Kunden weiterzugeben?” —%

Limiting factors for pass-through

- Welche Faktoren begrenzen den Spielraum Ihres Unternehmens bei der Preissetzung? (Bitte kreuzen Sie alles Zutreffende an)
 - Wettbewerbsdruck
 - schwache Nachfrage / Zahlungsbereitschaft
 - langfristige Vertragsbindung
 - administrativer Aufwand
 - Regulierung / Preissetzung liegt nicht beim Unternehmen
 - Sonstiges, und zwar:

Duration of supply shortages

- October 2021: “Falls Ihre Produktionstätigkeit zurzeit durch einen Mangel an Rohstoffen / Vormaterialien behindert wird: Was vermuten Sie, wie lange werden diese Probleme noch anhalten?” — Monate
- May 2022: “Falls Ihre Produktionstätigkeit zurzeit durch Mangel an Rohstoffen oder Vormaterialien/ Lieferengpässe/ Materialknappheit behindert wird: Was vermuten Sie, wie lange werden diese Probleme noch anhalten?” — Monate

Change in revenues/demand

- “Die Nachfragesituation hat sich im Vergleich zum [Vormonat] gebessert/nicht verändert/verschlechtert.”

Range of orders

- “Unsere Auftragsbestände entsprechen derzeit einer durchschnittlichen Produktion von — Monat(en)”

Capacity utilization

- “Die Ausnutzung unserer Anlagen (betriebsübliche Vollaussnutzung = 100%) beträgt gegenwärtig” — % [tick box from 30% to 100 % in 5/10 p.p. steps, or enter value manually if larger than 100%]

Change in orders

- “Unser Auftragsbestand ist im Vergleich zum [Vormonat] gestiegen/etwa gleich geblieben/gesunken”

Production expectations for the next three months

- “Unsere Produktionstätigkeit wird voraussichtlich steigen/etwa gleich bleiben/abnehmen.”

Energy intensity


- “Was schätzen Sie, welchen Anteil des Umsatzerlöses musste Ihr Unternehmen 2021 für Energiekosten aufwenden (Energieintensität)?” — %

Business state, expectation, and uncertainty

- “Wir beurteilen unsere Geschäftslage als:” [continuous slider from 0 (bad) over 50 (satisfactory) to 100 (good)]
- “Erwartungen für die nächsten 6 Monate: Unsere Geschäftslage wird in konjunktureller Hinsicht:” [continuous slider from 0 (rather worse) over 50 (the same) to 100 (rather better)]
- “Die Unsicherheit hinsichtlich unserer Geschäftsentwicklung in den nächsten 6 Monaten schätzen wir wie folgt ein:” [continuous slider from 0 (low) over 50 (average) to 100 (high)] — %


B.1.4 Additional results

Figure B.2 – Binned scatterplots of pass-through and expected effects of gas supply shortages




Notes: The figure shows binned scatterplots between the pass-through until 04/2022 and the expected decrease in production if gas supply were to be cut by 10% (navy) or 50% (maroon). The estimated slope coefficients and standard errors are reported in the figure. Source: ifo Institute's business survey.

Figure B.3 – Limiting factors for pass-through




Notes: The figure presents the share of respondents reporting that the following factors limit pass-through: competition, weak demand, long-term contracts, administrative effort/burden, regulatory reasons, and others. Multiple answers possible. Source: ifo Institute's business survey.

Figure B.4 – Heterogeneous pass-through across firms (2022/2023)




Notes: The figure shows the distribution of firm-level pass-through of cost changes to sales prices for October 2022 and planned pass-through until April 2023 (conditional on cost changes until October 2022). Source: ifo Institute's business survey.

Figure B.5 – Heterogeneous pass-through across firms (2021/2022)




Notes: The figure shows the distribution of firm-level pass-through of cost changes to sales prices for June 2021 (only manufacturing) and April 2022. Source: ifo Institute's business survey.

Figure B.6 – Gradual increase of pass-through over time (sector split)




Notes: The figure shows the change in pass-through of cost changes to sales prices from October 2022 to April 2023, conditional on cost changes until October 2022, separately for the manufacturing sector, services sector, and trade sector. Source: ifo Institute's business survey.

Figure B.7 – Cross-sectional variation in pass-through within and between sectors: Drop short time spent




Notes: The stacked bar figure shows the R^2 of sector fixed effects on 2-digit and 4-digit level of pass-through in October 2022 (blue). The remaining variation is within sectors (red). The decomposition is based on a subsample, where time spent on the survey is observable and time spent is above the 25th percentile. Source: ifo Institute's business survey.

Figure B.8 – Cross-sectional histogram time spent



Notes: The figure shows the histogram of time spent on the survey in minutes (October 2022). The information on time spent is only observable for firms who participate online. Source: ifo Institute's business survey.

Figure B.9 – Time-series of average firm-level uncertainty



Notes: The figure shows the time-series of average firm-level uncertainty, measured by a quantitative slider from 0-100 about the business uncertainty in the next six months. Source: ifo Institute's business survey.

Table B.1 – Pass-through until 10/2022 and return on sales in 2019 and 2022

	(1)	(2)	(3)	(4)
Return on Sales 2019	-0.80** (0.32)	-0.68** (0.32)	-0.84*** (0.31)	-0.69** (0.32)
Return on Sales 2022	1.25*** (0.30)	1.21*** (0.30)	1.20*** (0.29)	1.15*** (0.30)
Log Employees			2.76*** (0.74)	2.36*** (0.77)
Constant	47.1*** (1.33)	46.8*** (1.33)	34.9*** (3.59)	36.4*** (3.74)
Observations	1078	1078	1078	1078
R^2	0.018	0.089	0.030	0.097
Sector FE	No	Yes	No	Yes

Notes: The table reports estimates from linear regressions of the firm-level pass-through of cost changes to sales prices until 10/2022 on the return on sales (ROS) in 2019 and 2022. The ROS of the current year is elicited in September in the survey. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Source: ifo Institute's business survey.

Table B.2 – Pass-through until 10/2022 and uncertainty at industry level

	(1)	(2)
Uncertainty	1.28*** (0.27)	1.33*** (0.37)
Business State		1.22*** (0.39)
Business Expectations		-1.17** (0.45)
Constant	-61.2*** (18.8)	-79.6 (48.1)
Observations	49	49
R^2	0.267	0.476

Notes: The table reports estimates from linear regressions of the firm-level pass-through of cost changes to sales prices until 10/2022 on uncertainty aggregated at 2-digit industries (at least 20 obs. per industry). Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Source: ifo Institute's business survey.


Table B.3 – Pass-through until 10/2022 and uncertainty at firm level

	(1)	(2)	(3)	(4)
Uncertainty	0.18*** (0.023)	0.19*** (0.025)	0.093*** (0.023)	0.10*** (0.023)
Business State		0.26*** (0.025)	0.22*** (0.023)	0.19*** (0.023)
Business Expectations		-0.20*** (0.030)	-0.083*** (0.028)	-0.058** (0.028)
Constant	20.7*** (1.71)	15.3*** (2.65)	19.8*** (2.46)	19.4*** (2.46)
Observations	4689	4689	4689	4689
R^2	0.014	0.039	0.235	0.316
Industry FE	No	No	2-Digit	4-Digit

Notes: The table reports estimates from linear regressions of the firm-level pass-through of cost changes to sales prices until 10/2022 on uncertainty. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Source: ifo Institute's business survey.

B.2 PASS-THROUGH ESTIMATES IN THE LITERATURE

Figure B.10 – Box plots of pass-through estimates in the literature



Notes: Comparison of pass-through estimates of selected papers in the literature, listed in Appendix Table B.4.

Table B.4 – Pass-through estimates across studies

Paper	Specification	0	1	2	3	4	5	6	12	24
Dedola, Kristoffersen and Zullig (2021)	energy cost	0.04	0.03	0.13	0.18	0.27	0.26	0.34	0.64	0.68
	import cost shock	0.28	0.34	0.35	0.40	0.41	0.39	0.38	0.35	0.21
Yagi et al. (2022)	material cost, JPN	0.20			0.40			0.55	0.70	0.81
	material cost pass-through USA	0.32			0.52			0.80	1.06	1.15
	exchange-rate pass-through JPN	0.01			0.02			0.03	0.05	0.05
Nakamura and Zerom (2010)	commodity cost to retail prices	0.06			0.17			0.18	0.26	0.25
	commodity cost to wholesale prices	0.12			0.28			0.27	0.27	0.26
Gopinath and Itskhoki (2010)	all countries, high freq. adjustments		0.00	0.04	0.06	0.11	0.11	0.13	0.16	0.16
	all countries, low freq. adjustments		0.00	0.01	0.01	0.02	0.03	0.04	0.04	0.08
	high-income OECD high freq		0.04	0.10	0.11	0.16	0.19	0.21	0.35	0.44
	high-income OECD low freq		0.00	0.02	0.04	0.07	0.07	0.07	0.13	0.19
Gopinath, Itskhoki and Rigobon (2010)	goods priced in non-dollars		0.96	0.92	0.95	0.95	0.98	0.98	0.94	0.94
	aggregate		0.22	0.22	0.25	0.26	0.27	0.27	0.28	0.31
	goods priced in dollars		0.03	0.04	0.06	0.07	0.08	0.09	0.12	0.15
Neiman (2010)	intrafirm		0.00				0.04			
	arm's length		0.00				0.04			
Ahlander, Carlsson and Klein (2023)	OLS	0.25	0.29	0.33	0.33	0.29	0.20	0.35		
	IV	-0.03	0.05	0.30	0.30	0.30	0.20	0.35	0.40	0.45

Notes: Comparison of pass-through estimates across selected papers in the literature. Horizon from zero to 24 months. Full title and author details on the individual papers are deferred to references. For papers with more than one baseline pass-through estimate, the specification is specified.

B.3 IFO MANAGEMENT-SURVEY

B.3.1 Original questionnaire



Die letzten zwei Jahre waren weltweit geprägt von steigenden Preisen. Im Folgenden wird abgefragt, wie Ihr Unternehmen generell sowie in der aktuellen Situation die **Preissetzung** gestaltet.

Ihre Antworten helfen dabei, das Verständnis für die Inflationsentwicklung und entscheidende Einflussfaktoren zu verbessern.

Wie gut sind Sie persönlich über die Preisstrategien Ihres Unternehmens informiert?

gar nicht informiert ☐ ☐ ☐ ☐ ☐ sehr gut informiert

Wie groß müsste ein Anstieg der **Einkaufspreise** für Energie, Rohstoffe und Vormaterialien (in %) **grundsätzlich** sein, damit Ihr Unternehmen die **Preise anpasst**?

 %

Um wie viel Prozent sind Ihre **Einkaufspreise** für Energie, Rohstoffe und Vormaterialien in den **letzten 2 Jahren gestiegen**?

(grobe Schätzung genügt)

 %

Zu welchem **Grad** haben Sie die höheren **Einkaufspreise** für Energie, Rohstoffe und Vormaterialien bereits an Ihre **Kunden weitergegeben**?

(keine Weitergabe = 0%, vollständige Weitergabe = 100%, überproportionale Weitergabe > 100%)

 %

Welche wesentlichen **Faktoren begrenzen den Preissetzungsspielraum** Ihres Unternehmens?

Bitte priorisieren Sie.

wichtigster Faktor

zweitwichtigster Faktor

drittwichtigster Faktor

Wie hoch war der **Jahresüberschuss nach Steuern bzw. Verlust** Ihres Unternehmens in **% des Nettoumsatzes** in den Jahren 2018 und 2022?

(grobe Schätzung genügt)

2018:


 %

2022:

 %

Versetzen Sie sich bitte nun in die folgenden **hypothetischen Szenarien**.

Stellen Sie sich vor, dass die **Einkaufspreise** für Sie und Ihre Konkurrenten in der Branche **dauerhaft um 20%** steigen aufgrund eines globalen Angebotsschocks (siehe Abbildung). Alle weiteren Faktoren, wie beispielsweise die Zinsentwicklung und fiskalpolitische Maßnahmen, bleiben unverändert. Nehmen Sie als Ausgangspunkt eine **gute wirtschaftliche Lage mit normaler Auslastung** für Ihr Unternehmen und die Gesamtwirtschaft an.




Zu welchem Grad würden Sie die Kostenanstiege in diesem Szenario an Ihre Kunden weitergeben?

(keine Weitergabe = 0%, vollständige Weitergabe = 100%, überproportionale Weitergabe > 100%)

Geben Sie an, wie hoch die Weitergabe zu dem jeweiligen Zeitpunkt ist.

	4 Wochen	3 Monate	6 Monate	12 Monate	24 Monate
Weitergabe in % innerhalb der nächsten...					

Stellen Sie sich nun die gleichen Gegebenheiten vor wie in der vorhergehenden Frage, mit dem einzigen Unterschied, dass der Kostenanstieg nicht dauerhaft ist, sondern **temporär für 12 Monate** (siehe Abbildung).



Zu welchem Grad würden Sie die Kostenanstiege in diesem Szenario an Ihre Kunden weitergeben?


(keine Weitergabe = 0%, vollständige Weitergabe = 100%, überproportionale Weitergabe > 100%)

Geben Sie an, wie hoch die Weitergabe zu dem jeweiligen Zeitpunkt ist.

	4 Wochen	3 Monate	6 Monate	12 Monate
Weitergabe in % innerhalb der nächsten...				

Stellen Sie sich nun die gleichen Gegebenheiten vor wie in der vorhergehenden Frage, mit dem einzigen Unterschied, dass der Angebotsschock nicht dauerhaft sein muss, sondern die **Dauer** der Kostenerhöhung **ungewiss** ist (siehe Abbildung mit Beispielen für mögliche Preispfade).

Mit gleicher Wahrscheinlichkeit (~10 %) kann die Kostenerhöhung jeden Monat wieder zurückgehen.



Zu welchem Grad würden Sie die Kostenanstiege in diesem Szenario an Ihre Kunden weitergeben?

Nehmen Sie an, dass zu dem jeweiligen Zeitpunkt die Kostenerhöhung noch vorhanden ist.

(keine Weitergabe = 0%, vollständige Weitergabe = 100%, überproportionale Weitergabe > 100%)

	4 Wochen	3 Monate	6 Monate	12 Monate	24 Monate
Weitergabe in % innerhalb der nächsten...					

Zu welchem Grad würden Sie die Kostenanstiege weitergeben, wenn **nur Ihr Unternehmen** und **nicht die gesamte Branche** betroffen ist? Gehen Sie von den gleichen Rahmenbedingungen wie in den vorhergehenden Fragen aus. Bitte geben Sie die Preisweitergabe **nach 6 Monaten** in diesen Szenarien an.

Permanenter

Kostenanstieg (bei dem branchenweiten Kostenanstieg haben Sie % angegeben):

Temporärer

Kostenanstieg (bei dem branchenweiten Kostenanstieg haben Sie % angegeben):

Unsichere Dauer

(bei dem branchenweiten Kostenanstieg haben Sie % angegeben):

Anmerkungen:

Wurde aufgrund von starken Kostenerhöhungen das Produktangebot umstrukturiert?

☐ ja

☐ nein

☐ nicht zutreffend

B.3.2 English translation

The last two years have been globally marked by rising prices. The following survey asks how your company approaches price setting generally as well as in the current situation. Your answers will help improve understanding of inflation dynamics and crucial decision-making factors.

1. How well-informed are you personally about your company's pricing strategies?

uninformed ○ ○ ○ ○ ○ well informed

2. What would be the minimum cost increase for energy, raw materials, and intermediate input goods (in %) for your company to adjust prices?

_____ %

3. By how much (in %) have input prices for energy, raw materials, and intermediate input goods increased in the last two years?

_____ %

(Rough estimate is sufficient)

4. To what extent have you already passed through the higher input prices for energy, raw materials, and intermediate input goods to your customers?

_____ %

(No pass-through = 0%, full pass-through = 100%, disproportionate pass-through > 100%)

5. What key factors limit your company's pricing flexibility? Please prioritize them.

Most important factor: _____

Second most important factor: _____

Third most important factor: _____


6. What was your company's annual net profit or loss as a percentage of net sales in 2018 and 2022?

_____ %

(Rough estimate is sufficient)

7. Please now imagine the following hypothetical scenarios.

Suppose that purchase prices for you and your competitors in the industry permanently increase by 20% due to a global supply shock (see figure). All other factors, such as interest rates and fiscal policy measures, remain unchanged. Assume [good economic conditions]¹ with normal capacity utilization for your company and the overall economy.




To what extent would you pass through the cost increase to your customers in this scenario? Indicate the level of pass-through at the respective point in time.

	4 weeks	3 months	6 months	12 months	24 months
Pass-through in % after ...					

(No pass-through = 0%, full pass-through= 100%, disproportionate pass-through > 100%)

8. Now, imagine the same circumstances as in the previous question, with the only difference that the cost increase is not permanent but temporary for 12 months (see figure).

¹Economic conditions vary across participants: good economic conditions vs. bad economic conditions.




To what extent would you pass through the cost increase to your customers in this scenario? Indicate the level of pass-through at the respective point in time.

Pass-through in % after ... 4 weeks 3 months 6 months 12 months

(No pass-through = 0%, full pass-through= 100%, disproportionate pass-through > 100%)

9. Now, imagine the same circumstances as in the previous question, with the only difference that the cost increase is not necessarily permanent, but the duration of the cost increase is uncertain (see figure with examples of possible price paths). There is an equal probability (10%) that the cost increase could peak off every month.



To what extent would you pass through the cost increase to your customers in this scenario? Indicate the level of pass-through at the respective point in time. Assuming that at each respective point in time, the cost increase is still present.

	4 weeks	3 months	6 months	12 months	24 months
Pass-through in % after ...	☐	☐	☐	☐	☐

(No pass-through = 0%, full pass-through= 100%, disproportionate pass-through > 100%)

10. To what extent would you pass through cost increases if only your company, and not the entire industry, is affected? Assuming the same conditions as in the previous questions, please indicate the degree of price pass-on after 6 months in these scenarios.

- Permanent cost increase (for the industry-wide increase you stated ☐%):
_____%
- Temporary cost increase (for the industry-wide increase you stated ☐%):
_____%
- Uncertain duration (for the industry-wide increase you stated ☐%): _____%

11. Has your product portfolio been restructured due to significant cost increases?

Yes ☐

No ☐


Not applicable ☐

12. If yes, in what manner?

Thank you for your participation!


B.3.3 Additional results

Figure B.11 – Threshold of cost increases for pass-through




Notes: The figure plots the cumulative density function of the minimum required cost increase (in %) for a firm to change prices. 5%, 10%, and 20% cost threshold levels are highlighted by the dashed lines.
Source: ifo Institute's Management-Survey.

Figure B.12 – Heterogeneity of pass-through across sectors and firm size



Notes: The figure plots the pass-through of cost-push shocks across horizons, split by sector (left column) and size (right column). The rows refer to the nature of the shock, differing between permanent, transitory, and uncertain duration shocks, in blue, green, and red colors, respectively. The whiskers reflect the 68% confidence interval (one standard error). Source: ifo Institute's Management-Survey.

Figure B.13 – Aggregate vs idiosyncratic pass-through across sectors and firm size



Notes: The figure shows the pass-through of permanent vs. idiosyncratic cost-push shocks after six months, depending on the nature of the shock, differing between permanent, transitory, and uncertain duration shocks. The left column splits firms by sector, and the right column by firm size. The whiskers reflect the 68% confidence interval (one standard error). Source: ifo Institute's Management-Survey.

Table B.5 – Factors limiting firm’s price setting (open-ended question)

Factor	%
Competition	41
Customers	16
Market	10
Contracts	9
Demand	3
Regulation	2
Fixed prices	2
Energy	2
Input prices	2
Employees	2
Others	11

Notes: The table presents the share of respondents (in %), stating that the above factor limits the firm’s price setting. The results are based on an open-ended survey question, see Appendix B.3.2 for the exact wording (question 5). We focus on the “most important factor”, the first out of three text boxes. The ten categories group similar answers. *Others* summarizes all factors stated once and cannot be assigned to the other categories. Source: ifo Institute’s Management-Survey.

Table B.6 – Expected duration of the shock and pass-through in April 2022

	(1)	(2)	(3)	(4)	(5)	(6)
Duration Supply (Oct. 21)	0.68*** (0.22)	0.72** (0.28)	0.83*** (0.31)			
Duration Supply (June 22)				0.38*** (0.12)	0.51*** (0.17)	0.56** (0.22)
Range of Orders (April 22)		-0.97** (0.39)	-0.94** (0.47)		-0.97*** (0.28)	-1.09*** (0.39)
Capacity Utilization (April 22)		0.49*** (0.11)	0.65*** (0.14)		0.13 (0.087)	0.39*** (0.13)
Change in Revenues (April 22)		2.90 (2.28)	3.81 (2.70)		4.72** (1.95)	5.55** (2.65)
Change in Orders (April 22)		2.05 (2.11)	-1.21 (2.74)		2.49 (2.01)	0.25 (2.88)
Production Expectations (April 22)		-1.02 (2.50)	-1.56 (2.84)		3.55 (2.27)	-0.71 (3.08)
Log Employees		0.61 (0.97)	-0.26 (1.25)		1.54* (0.88)	2.36* (1.24)
Cost Changes (March-May 21)			0.11** (0.054)			0.16* (0.079)
Energy Intensity 21			-0.93*** (0.25)			-0.59** (0.27)
Constant	45.1*** (1.75)	5.05 (9.62)	-1.16 (11.7)	44.7*** (1.26)	29.7*** (7.51)	6.96 (12.3)
Observations	951	732	561	1095	878	511
R^2	0.202	0.262	0.287	0.179	0.221	0.265

Notes: The table reports estimates from linear regressions of the firm-level pass-through of cost changes to sales prices in April 2022 on the expected duration of supply shortages in months (elicited in October 2021 and June 2022) in the manufacturing sector. Columns 2 and 5 control for the range of orders in months, capacity utilization in percent, the qualitative change in revenues, qualitative change in orders, qualitative production expectations, and log employees. Columns 3 and 6 additionally control for the energy intensity (measured as the share of energy costs in terms of revenues) and cost changes between March and May 2021. All specifications control for 4-digit industry fixed effects. Appendix B.1 provides the wording of the specific questions. Clustered standard errors at the 4-digit industry level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Source: ifo Institute's business survey.

Discussion of other variables in Table B.6 The coefficients of the control variables in Table B.6 are also of interest to itself. For instance, the negative coefficient on the range of orders in months indicates the existence of nominal rigidities. Firms with an extended range of orders cannot pass through increased costs due to existing contracts with fixed prices. In addition, a more significant cost increase is associated with a higher pass-through. Intuitively, large cost increases imply large adjustment gains relative to the fixed costs of adjusting prices. Lastly, there is a negative effect of energy exposure on pass-through. This may be attributed to two potential factors: first, energy costs are part

Table B.7 – Change in expected duration of the shock and pass-through until April 2023

	(1)	(2)	(3)	(4)	(5)	(6)
Δ Duration Supply	0.48** (0.21)	0.56** (0.24)	0.59** (0.26)			
Δ Duration Supply (winsorized)				0.70** (0.30)	0.69** (0.31)	0.74** (0.33)
Pass-Through (April 22)	0.40*** (0.049)	0.42*** (0.055)	0.45*** (0.060)	0.40*** (0.048)	0.42*** (0.055)	0.45*** (0.059)
Range of Orders (April 22)		1.21*** (0.45)	1.18** (0.46)		1.22*** (0.45)	1.19** (0.46)
Capacity Utilization (April 22)		0.11 (0.13)	0.061 (0.16)		0.11 (0.13)	0.063 (0.16)
Change in Revenues (April 22)		-1.18 (2.37)	-1.54 (2.42)		-1.13 (2.36)	-1.49 (2.41)
Change in Orders (April 22)		1.50 (2.94)	0.52 (3.02)		1.55 (2.93)	0.54 (2.99)
Production Expectations (April 22)		-6.17** (2.93)	-4.93 (3.43)		-6.04** (2.93)	-4.77 (3.44)
Log Employees		0.95 (1.20)	0.68 (1.48)		0.91 (1.21)	0.63 (1.49)
Energy Intensity 21			-0.59** (0.26)			-0.58** (0.26)
Cost Changes (March-May 21)			0.020 (0.100)			0.013 (0.10)
Constant	44.9*** (2.54)	25.2** (11.5)	31.9** (14.9)	44.3*** (2.52)	25.0** (11.5)	31.7** (14.9)
Observations	592	467	379	592	467	379
R^2	0.341	0.387	0.413	0.342	0.388	0.414

Notes: The table reports estimates from linear regressions of the firm-level pass-through of cost changes to sales prices until April 2023 on the change of expected duration of supply shortages in months from October 2021 to June 2022 in the manufacturing sector. In Columns 4-6, the change in the duration of supply shortages is winsorized at the 5th and 95th percentiles. Columns 2 and 5 control for the range of orders in months, capacity utilization in percent, the qualitative change in revenues, qualitative change in orders, qualitative production expectations, and log employees. Columns 3 and 6 additionally control for the energy intensity (measured as the share of energy costs in terms of revenues) and cost changes between March and May 2021. All specifications control for 4-digit industry fixed effects. Appendix B.1 provides the wording of the specific questions. Clustered standard errors at the 4-digit industry level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Source: ifo Institute's business survey.

of firms' fixed costs and are therefore to a first order irrelevant for pricing off marginal costs; second, firms' energy costs are very volatile and hence, as just discussed, firms have a lower incentive to respond to transitory shocks.

B.3.4 Quantifying the pass-through dynamics and implications for the Phillips curve

The pass-through matrix, Ψ , which is identified in partial equilibrium, delivers a sufficient statistic for the generalized PC slope coefficient matrix $\mathbf{K}$.² The PC reflects the inflation, π , response to output, $\hat{y}$, depending on the coefficient matrix $\mathbf{K}$ and the elasticity of intertemporal substitution and intertemporal labor supply elasticity, σ and φ , respectively:

$$\pi = (\varphi + \sigma) \underbrace{(\mathbf{I} - \mathbf{L})\Psi(\mathbf{I} - \Psi)^{-1}}_{\mathbf{K}} \hat{y}, \quad (\text{B.1})$$

where Ψ reflects the pass-through matrix, a function of the adjustment probability θ and the discount factor β , $\mathbf{I}$ the identity matrix and $\mathbf{L}$ the lag-matrix.³ The implied PC coefficients from our survey experiment are presented on a monthly level in Appendix Table B.8.

Table B.8 – Implied coefficients of Phillips curve slope $\mathbf{K}$

	nominal rigidities	nominal + real rigidities	
		calvo counterfactual	based on idios. shock PT
	$\theta=0.37$	$\theta=0.22$	$\theta=0.37; \chi=0.625$
$\mathbf{K}$	0.2186	0.0628	0.1366

Notes: The table presents the implied coefficients of the Phillips curve slope based on $\mathbf{K} = (\mathbf{I} - \mathbf{L})\Psi(\mathbf{I} - \Psi)^{-1}$ for different values of θ . In the third column, the $\mathbf{K}$ based on nominal rigidities is multiplied by χ , capturing the extent of real rigidities.

²The formula for $\mathbf{K}$ in Equation B.1 is the solution of a fixed point problem that translates nominal into real marginal costs and thereby accounts for general equilibrium effects (Auclert et al., 2024).

³The pass-through matrix, Ψ , for a time-dependent (e.g., Calvo) model from Auclert et al. (2024) equals:

$$\Psi \equiv \frac{1}{\sum_{s \geq 0} (1 - \theta)^s \sum_{s \geq 0} \beta^s (1 - \theta)^s} \begin{pmatrix} 1 & 0 & 0 & \dots \\ (1 - \theta)^1 & 1 & 0 & \dots \\ (1 - \theta)^2 & (1 - \theta)^1 & 1 & \dots \end{pmatrix} \begin{pmatrix} 1 & \beta(1 - \theta)^1 & \beta^2(1 - \theta)^2 & \dots \\ 0 & 1 & \beta(1 - \theta)^1 & \dots \\ 0 & 0 & 1 & \dots \end{pmatrix}$$

We set the discount factor β to .9966 in accordance with the monthly frequency.

In the first column, we show the implied PC coefficient based on the nominal rigidities observed in the survey experiment. The implied PC slope coefficient is with 0.22 relatively steep. In the second and third columns, we display the PC coefficients that also consider real rigidities based on two approaches: Either relying on the “Calvo counterfactual” - as described above - or using the PC coefficient based on nominal rigidities and adjusting for micro real rigidities in a second step based on the pass-through of the idiosyncratic shock scenario.

Following Auclert et al. (2024), in an environment with micro real rigidities, the PC coefficient extends to:

$$\pi = (\varphi + \sigma)\chi\mathbf{K}\hat{y},$$

where χ reflects the micro real rigidities for values less than one and downscales the slope of the PC. We can derive χ by using the idiosyncratic shock results presented in Figure 2.8. Focusing on the permanent shock, we obtain $\chi = 0.625$ in the high-demand environment.⁴ The degree of micro real rigidities is likely to be even higher in the low-demand environment or for transitory shocks (as the gap between idiosyncratic and aggregate shocks is wider in these scenarios).

The PC coefficient gets substantially flatter for both approaches that account for real rigidities. Based on the “Calvo counterfactual,” it shrinks to 0.06 and based on the idiosyncratic shock PT approach to 0.14. Overall, our estimate is in the range of previous work, as summarized in Appendix Table B.9.⁵ However, a recent study by Hazell et al. (2022) estimates an even flatter slope, suggesting additional frictions at play, such as macro real rigidities.

Table B.9 – Estimates of K in the literature

	ARRS (2023)			RW (1997)	Gali (2008)	NS (2014)	HHNS (2023)
	GL	NS	Calvo				
<i>K</i>	1.71	0.47	0.08	0.0095	0.0425	0.00385	0.0031

Notes: The table summarizes estimates in the literature. ARRS: Auclert et al. (2024) with calibration to GL (Goloso and Lucas, 2007), NS (Nakamura and Steinsson, 2014), and a Calvo setting; RW: Woodford and Rotemberg (1997); Gali (2008); NS: Nakamura and Steinsson (2014); HHNS: Hazell et al. (2022).

⁴ χ is calculated as pass-through of the idiosyncratic shock divided by pass-through of the aggregate shock after 24 months (Figure 2.8) serving as a full pass-through counterfactual.

⁵To relate unemployment to marginal costs, we divided the unemployment coefficient by two following Auclert et al. (2024).

B.4 ADDITIONAL MODEL DETAILS AND RESULTS

B.4.1 Calibration

Table B.10 – Calibration for the menu-cost model with micro real rigidities

Parameter	Description	Value
β	Discount factor	$0.96^{1/12}$
γ	Relative risk aversion	1.5
ϕ	Inverse of Frisch elasticity of labor supply	1
L^*	Steady-state labor share	$1/3$
sm	Material input share	0.7
$\bar{\theta}$	Elasticity of demand	4
μ	Trend inflation (monthly)	0.0015906
σ_S	St. dev. of nominal demand shock	0.0015883
σ_A	St. dev. of productivity shock	0.046
ρ_A	Persistence of productivity shock	0.7
ρ_S	Persistence of nominal demand shock	1


The trend inflation rate μ and standard deviation of the nominal demand σ_S are calculated as the average and standard deviation of the German Harmonised Index of Consumer Prices (HICP) for 1997 to 2023.

B.4.2 Simulated method of moments

We simulate data for a broad range of menu cost values and real rigidities, estimate the pass-through dynamics and frequency of price changes for each parameter combination, and calculate the weighted value function for each pair correspondingly. We weight both moments—the frequency of price changes $freq^{data}$ and the level of the impulse response function to the aggregate shock $IRF_h^{data} \quad \forall h \in (1, 3, 6, 12, 24)$ —according to their relative squared deviations from the target and equally to minimize the target function: $\frac{1}{(1+\sum^H IRF_h^{data})^2} \sum_h^H (IRF_h^{model} - IRF_h^{data})^2 + \frac{1}{(1+freq^{data})^2} (freq^{model} - freq^{data})^2$, where h reflects the respective horizon.

B.4.3 Shape of the demand and profit functions


Figure B.14 – Relative demand and profit functions varying extent of micro real rigidities



Notes: The figure shows the relative demand and profit functions for different values of the super-elasticity parameter $\bar{\epsilon}$.


B.4.4 Results of estimation

Figure B.15 – Values of the weighted function



Notes: The figure shows the values of the weighted function on the vertical axis depending on the degree of real rigidities (0-25) and menu costs (0.001-0.01). The right panel focuses on the area around the minimum. The minimum is marked by the yellow cross. The values of the weighted function above 0.05 and 0.003 are winsorized in the left and right panels, respectively.

Figure B.16 – Effect of nominal demand shock on output



Notes: The solid dark-blue line reflects the model-implied impulse response of real output to a permanent aggregate shock with the estimated degree of real rigidities. The solid orange line is the model-implied impulse response of real output to a transitory aggregate shock with the estimated degree of real rigidities. The dashed light-blue line is the model-implied counterfactual impulse response of real output to a permanent aggregate shock with no real rigidities. The dashed light-orange line is the model-implied counterfactual impulse response of real output to a transitory aggregate shock with no real rigidities.

C Appendix to Chapter 3

C.1 STYLIZED MODEL – DETAILS AND PROOFS

Section C.1.1 presents the household sector, Section C.1.2 derives the solution of the model, and Section C.1.3 provides the proof of Proposition 1. Section C.1.4 discusses the extension for price expectations.

C.1.1 Households

The representative household maximizes a standard utility function of the form

$$U = \ln(C) + \ln(M) - \frac{1 + N^{1+\varphi}}{1 + \varphi},$$

where consumption C is a Dixit-Stiglitz bundle consisting of goods from all islands with an elasticity of substitution γ . M is the (exogenously fixed) money supply, $N = \int N^i di$ total hours worked, and φ is the inverse of the Frisch elasticity. Given the budget constraint $CP + M = WN$ and market clearing, we obtain the demand function for individual goods as $Y^i = (P^i/P)^{-\gamma} M/P$. Optimal labor supply is given by $N^\varphi C = W/P$ and optimal money holdings by $C = M/P$.

C.1.2 Solution

Throughout, we consider a linear approximation to the equilibrium conditions of the model. Lower-case variables denote percentage deviations from steady state, and variables without island subscripts refer to economy-wide values. The wage is set at the beginning of the period based on all available information. As information sets of agents are perfectly aligned at this point, we use the expectation operator $F_0 = F_0^i$ to denote (common) stage-one expectations. Linearized expected labor supply can be combined with expected labor demand $F_0(w-p) = F_0(\varepsilon)$, the linearized production function $F_0(y^i) = F_0(n^i + \varepsilon)$, optimal money holdings $F_0(c) = m - F_0(p)$, market clearing, and symmetry assumptions to obtain $w = m$. We furthermore normalize the commonly known initial value of technology, i.e.,

before the realization of ε , to zero. The mentioned equilibrium conditions, together with $F_0(\varepsilon) = 0$, also imply $F_0(y^i) = 0$. After w is set, firms observe their own productivity and the public signal.

Linearized demand for firm j 's goods is given by

$$y^i = -\gamma p^i + (\gamma - 1)p + m . \quad (\text{C.1})$$

Given the production function $y^i = a^i + n^i$, firms set prices according to

$$p^i = w - a_t^i . \quad (\text{C.2})$$

Firm-specific technology equals $a^i = \varepsilon + \eta^i$ with $\int \eta^i di = 0$. Aggregating over all producers gives the aggregate price index

$$p = w - \varepsilon . \quad (\text{C.3})$$

The public signal is $s = p + e$. Without loss of generality we can define ω_1 and ω_2 such that $\sigma_\eta^2 = \omega_1 \text{Var}(a^i)$ and $\sigma_\varepsilon^2 = (1 - \omega_1) \text{Var}(a^i)$ as well as $\sigma_\varepsilon^2 = (1 - \omega_2) \text{Var}(s)$ and $\sigma_e^2 = \omega_2 \text{Var}(s)$, where $\text{Var}(s)$ refers to the variance of s conditional on w . Firms correctly observe the variances of their own technology $\text{Var}(a^i)$ and that of the public signal $\text{Var}(s)$, but not that of its components. Firms therefore form their expectation regarding ε after observing a^i and s according to

$$F^i(\varepsilon) = \delta a^i - \rho(s - w) = \delta(\varepsilon + \eta^i) + \rho(\varepsilon - e) . \quad (\text{C.4})$$

Here, F^i refers to subjective expectations of firm i after observing its own technology and the public signal,

$$\delta = \frac{\hat{\sigma}_e^2}{\hat{\sigma}_\eta^2 + \hat{\sigma}_e^2 + \frac{\hat{\sigma}_\eta^2 \hat{\sigma}_e^2}{\hat{\sigma}_\varepsilon^2}} \quad \text{and} \quad \rho = \frac{\hat{\sigma}_\eta^2}{\hat{\sigma}_\eta^2 + \hat{\sigma}_e^2 + \frac{\hat{\sigma}_\eta^2 \hat{\sigma}_e^2}{\hat{\sigma}_\varepsilon^2}} ,$$

where $\hat{\sigma}^2$ refers to perceived variances that result from a perceived weight $\hat{\omega}_1 = \omega_1 \Upsilon$ of idiosyncratic technology in own technology. $\Upsilon = 1$ represents rational expectations. For $\sigma_a^2 = \text{Var}(a^i) = \text{Var}(s) = \sigma_s^2$, the above equations collapse to those in the main text, using equation (C.3). Expectation formation is consistent with the structure of technology and the signal, such that $\hat{\sigma}_\varepsilon^2 = (1 - \hat{\omega}_1) \text{Var}(a^i) = (1 - \hat{\omega}_2) \text{Var}(s)$, i.e., $1 - \hat{\omega}_2 = (1 - \hat{\omega}_1) \text{Var}(a^i) / \text{Var}(s)$.

C.1.3 Proof of Proposition 1

Observing market clearing and equation (C.3), aggregate output results as

$$y = m - p = m - w + \varepsilon = \varepsilon . \quad (\text{C.5})$$

The gap between expected own and aggregate output can then be calculated using equations (C.1), (C.2), and (C.3), as

$$F^i(y^i) - F^i(y) = -\gamma (p^i - F^i(p)) = \gamma F^i(\eta^i) . \quad (\text{C.6})$$

Forecast revisions are given by the change in expectations between before and after receiving the private and public signals. Before receiving the signals, firms expect their output to equal aggregate output since $F_0(\eta^i) = 0 \forall i$. Equation (C.5) implies $F_0(y) = F_0(\varepsilon) = 0$ and $F^i(y) = F^i(\varepsilon)$. Using this together with equation (C.6) gives

$$FR^i = F^i(y^i) - F_0(y) = F^i(y^i) - F^i(y) + F^i(y) = \gamma F^i(\eta^i) + F^i(\varepsilon) .$$

Because of equation (C.4) and

$$F^i(\eta^i) = (1 - \delta)(\varepsilon + \eta^i) + \rho(s - w) ,$$

we can write the above as

$$FR^i = \gamma(1 - \delta)(\varepsilon + \eta^i) + \delta(\varepsilon + \eta^i) + \rho(s - w)(\gamma - 1) \equiv X\varepsilon + X^\eta \eta^i - X^e e ,$$

with $X^\eta = (\gamma - 1)(1 - \delta) + 1 > 0$, $X^e = -\rho(\gamma - 1) < 0$, $X = X^\eta + X^e > 0$.

Calculating the expectation error of firms for idiosyncratic output, using the island-specific demand equation (C.1) and the price-level equation (C.3), yields

$$FE^i = y^i - F^i(y^i) = \Lambda (\varepsilon - F^i(\varepsilon)) = \Lambda [(1 - \delta - \rho)\varepsilon - \delta\eta^i + \rho e] . \quad (\text{C.7})$$

The effect $\Lambda = -(\gamma - 1)$ of the expectation error regarding aggregate technology innovations on the expectation error regarding own output is negative.

The sign of β_1 in regression (3.4) can then be determined in two steps. Since both independent variables, forecast revisions and the signal, are correlated, we first regress the forecast revisions on the signal, yielding the regression coefficient

$$Coe f_1 = \frac{\text{Cov}(FR^i, s)}{\text{Var}(s)} = -\frac{X^\eta \sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_e^2} - X^e . \quad (\text{C.8})$$

The residual of this regression is $FR^i - Coef_1 s$. Thus, the coefficient β_1 in equation (3.4) is given by

$$\beta_1 = \frac{\text{Cov}(FE^i, FR^i - Coef_1 s)}{\text{Var}(FR^i - Coef_1 s)}.$$

The sign of β_1 equals the sign of the numerator, which can be written as

$$\text{Cov}(FE^i, FR^i) - Coef_1 \text{Cov}(FE^i, s) = \Lambda[X^\eta R_\eta + (X^e + Coef_1)R_e],$$

with $R_e \equiv (1 - \delta - \rho)\sigma_\varepsilon^2 - \rho\sigma_e^2$ and $R_\eta \equiv (1 - \delta - \rho)\sigma_\varepsilon^2 - \delta\sigma_\eta^2$. These two terms are proportional to the covariances of the forecast error with the public signal and idiosyncratic technology, respectively. Both covariances are zero under rational expectations ($\Upsilon = 1$), such that $R_e = R_\eta = \beta_1 = 0$. If firms underestimate the signal-to-noise ratio (proportional to the covariance of a^i with ε) of micro news, we get $R_\eta > 0$. In contrast, an underestimation of the signal-to-noise ratio of macro news (proportional to the covariance of s with ε) yields $R_e > 0$. Island illusion ($1 < \Upsilon < 1/\omega_{1/2}$) therefore implies $R_\eta, R_e > 0$. Hence, because of equation (C.8), β_1 is negative if

$$R_\eta - \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_e^2} R_e > 0 \quad \Leftrightarrow \quad \frac{R_e}{R_\eta} = \frac{\hat{\sigma}_\eta^2/\sigma_a^2}{\hat{\sigma}_e^2/\sigma_s^2} < \frac{\sigma_\varepsilon^2 + \sigma_e^2}{\sigma_\varepsilon^2}. \quad (\text{C.9})$$

The sign of β_2 in equation (3.4) can equivalently be derived by first regressing the forecast revision on the signal, which gives the coefficient

$$Coef_2 = \frac{\text{Cov}(FR^i, s)}{\text{Var}(FR^i)} = -\frac{X\sigma_\varepsilon^2 + X^e\sigma_e^2}{X^2\sigma_\varepsilon^2 + (X^\eta)^2\sigma_\eta^2 + (X^e)^2\sigma_e^2},$$

such that β_2 equals

$$\beta_2 = \frac{\text{Cov}(FE^i, s - Coef_2 FR^i)}{\text{Var}(s - Coef_2 FR^i)}.$$

Its sign depends on the sign of the numerator, which can be written as

$$\text{Cov}(FE^i, s) - Coef_2 \text{Cov}(FE^i, FR^i) = -\Lambda[(1 + Coef_2 X^e)R_e + Coef_2 X^\eta R_\eta].$$

With $\Upsilon = 1$ (rational expectations), we have $R_e = R_\eta = \beta_2 = 0$. Note that $(1 + Coef_2 X^e)R_e + Coef_2 X^\eta R_\eta$ has the same sign as $[X\sigma_\varepsilon^2 + X^\eta\sigma_\eta^2]R_e - [X\sigma_\varepsilon^2 + X^e\sigma_e^2]R_\eta$, which is positive if

$$\frac{R_e}{R_\eta} > \frac{X\sigma_\varepsilon^2 + X^e\sigma_e^2}{X\sigma_\varepsilon^2 + X^\eta\sigma_\eta^2} \quad \Leftrightarrow \quad \frac{R_e}{R_\eta} = \frac{\hat{\sigma}_\eta^2/\sigma_a^2}{\hat{\sigma}_e^2/\sigma_s^2} > \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \gamma\sigma_\eta^2}. \quad (\text{C.10})$$

Combining equations (C.9) and (C.10), we obtain $\beta_1 < 0$ and $\beta_2 > 0$ for

$$\underbrace{\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \gamma\sigma_\eta^2}}_{<1} < \frac{\hat{\sigma}_\eta^2/\sigma_a^2}{\hat{\sigma}_e^2/\sigma_s^2} < \underbrace{\frac{\sigma_\varepsilon^2 + \sigma_e^2}{\sigma_\varepsilon^2}}_{>1} \quad (\text{C.11})$$

or, expressed in terms of weights

$$\frac{1 - \omega_1}{1 + \omega_1(\gamma - 1)} < \frac{\omega_1}{\omega_2} < \frac{1}{1 - \omega_2} .$$

Condition (C.11) states that the ratio of the subjective signal-to-noise ratios of the two signals needs to stay within certain bounds. It is automatically fulfilled for the assumption of $\text{Var}(a^i) = \text{Var}(s)$, made in the main text. This case corresponds to firms observing their own technology and, additionally, one price forecast based on the technology of one island, made public by, e.g., a forecasting firm on that island. In this case, $(\hat{\sigma}_\eta^2/\sigma_a^2)/(\hat{\sigma}_e^2/\sigma_s^2) = 1$. Furthermore, condition (C.11) is always fulfilled for high degrees of island illusion ($\Upsilon \rightarrow \infty$) and/or a high weight of the idiosyncratic component in technology ($\sigma_\eta^2/\sigma_a^2 \rightarrow 1$) and/or a high weight of noise in the public signal ($\sigma_e^2/\sigma_s^2 \rightarrow 1$), where the two last conditions are simultaneously given by a low volatility of aggregate technology ($\sigma_\varepsilon^2 \rightarrow 0$). The first part is also fulfilled for a high demand elasticity ($\gamma \rightarrow \infty$).

Regarding the derivatives of β_1 and β_2 , note that we can write

$$\beta_1 = \frac{\Lambda}{X^\eta} \frac{R_\eta - \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_e^2} R_e}{\sigma_\varepsilon^2 \sigma_e^2 / (\sigma_\varepsilon^2 + \sigma_e^2) + \sigma_\eta^2} = \frac{\Lambda}{\gamma + \delta/(1 - \delta)} \left[\frac{\sigma_\varepsilon^2 \sigma_e^2}{\sigma_\varepsilon^2 + \sigma_e^2} - \frac{\delta}{1 - \delta} \left(\sigma_\eta^2 + \frac{\sigma_\varepsilon^2 \sigma_e^2}{\sigma_\varepsilon^2 + \sigma_e^2} \right) \right] .$$

Its derivative w.r.t. Υ is negative whenever the derivative of $\delta/(1 - \delta)$ w.r.t. Υ is negative. This is the case if

$$\frac{\hat{\sigma}_\eta^2/\sigma_a^2}{\hat{\sigma}_e^2/\sigma_s^2} < \frac{\sigma_e^2 + \sigma_\varepsilon^2}{\hat{\sigma}_\varepsilon^2} ,$$

which is given by condition (C.11). For the derivative of β_2 w.r.t. Υ , we can write

$$\beta_2 = \frac{\text{Cov}(FE^i, s - \text{Coe}f_2 FR^i)}{\text{Var}(s - \text{Coe}f_2 FR^i)} = -\Lambda \frac{\sigma_\eta^2 \sigma_\varepsilon^2 - [\sigma_e^2 \sigma_\varepsilon^2 + \gamma \sigma_\eta^2 (\sigma_\varepsilon^2 + \sigma_e^2)] \rho / X^\eta}{\sigma_\varepsilon^2 \sigma_e^2 / \bar{\delta}} ,$$

where $\bar{\delta}$ is δ under rational expectations, and hence independent of Υ . Hence, β_2 depends on Υ only via ρ/X^η , which can be written as

$$\frac{1 - \omega_2 \Upsilon}{\gamma + \omega_2(1/\omega_1 - \Upsilon)} .$$

The derivative of this term is negative whenever

$$\frac{1 - \omega_1 \Upsilon}{1 + \omega_1 \Upsilon (\gamma - 1)} < \frac{\omega_1}{\omega_2},$$

which is given if condition (C.11) is fulfilled, such that β_2 depends positively on Υ . ■

C.1.4 Price expectations

To derive the expectations of firms regarding their future prices, we need to specify how technology evolves after the period of pre-set prices. We follow Lorenzoni (2009) and assume that aggregate technology ε is a random walk while the idiosyncratic component η^i is white noise. From above, we have $F_+^i(w) = m$, where F_+^i refers to expectations regarding the time after the pre-set price period. With expected labor demand $F_+^i(w - p) = F_+^i(\varepsilon)$, it follows that $F_+^i(p^i) = F_+^i(p) = m - F_+^i(\varepsilon)$, where $F_+^i(\varepsilon) = F^i(\varepsilon)$ because of the random-walk assumption. Thus, the forecast revision (before and after receiving the signals) regarding the change in own future prices is $F_+^i(p - p^i) = F^i(\eta^i)$ and the forecast error is $(p - p^i) - F_+(p - p^i) = -(\varepsilon - F(\varepsilon))$, see equation (C.7). The forecast revision and the forecast error are, hence, the same as for production expectations, but divided by $\gamma - 1$ and for $\gamma \rightarrow \infty$. It follows that β_1 remains unchanged and β_2 turns into $\beta_2/(\gamma - 1)$ (note that β_2 is proportional to $\gamma - 1$). Hence, we obtain the same results for price expectations as in Proposition 1 for production expectations.

C.2 SURVEY ON INFLATION AND GROWTH EXPECTATIONS

Table C.1 – Relevant questions from SIGE

Question	Name	Introduced	Wording
Q1	expected change in own price	1999q4	For the next 12 months, what do you expect will be the average change in your firm's prices?
Q2	expected inflation (12 months ahead)	1999q4	In July consumer price inflation, measured by the 12-month change in the harmonized index of consumer prices was 8.4 percent in Italy and 8.9 percent in the euro area. What do you think it will be in Italy in September 2023?
Q3	realized change in own prices	2002q4	In the last 12 months, what has been the average change in your firm's prices?
Q4	expected inflation (24 months ahead)	2009q2	In July consumer price inflation, measured by the 12-month change in the harmonized index of consumer prices was 8.4 percent in Italy and 8.9 percent in the euro area. What do you think it will be in Italy in September 2024?
Q5	expected inflation (6 months ahead)	2010q4	In July consumer price inflation, measured by the 12-month change in the harmonized index of consumer prices was 8.4 percent in Italy and 8.9 percent in the euro area. What do you think it will be in Italy in March 2023?
Q6	expected inflation (12 months ahead, no infl. update)	2012q3	What do you think consumer price inflation in Italy, measured by the 12-month change in the harmonized index of consumer prices, will be in September 2023?


Notes: Wording taken from the September 2022 questionnaire. Questions elicit growth rates in percentage points. We construct firms' forecast error for their own prices by subtracting their expectation reported in quarter t (Q1) from their realization reported in quarter $t + 4$ (Q3). For micro news, we generally consider the first difference of firms' one-year-ahead expectation for the change in their own prices (Q1). For macro news, we generally consider the first difference of firms' one-year-ahead inflation expectation (Q2). Since 2009 and 2010, the SIGE features additional questions about firms' inflation expectations six months and two years ahead (Q5 and Q4). We use these questions to construct macro news that uniquely refer to the same forecasting target. See main text for more details.

Table C.2 – Macro news and forecast revisions

	Forecast revision for			
	$\pi_{t+4,t}^i$		$\pi_{t+4,t}^i$ net of $\gamma_i \Delta \pi_t$	
	(1)	(2)	(3)	(4)
Macro News				
Forecast revision for $\pi_{t+4,t}$	0.292*** (0.033)		0.230*** (0.031)	
Inflation surprise		0.861*** (0.191)		0.636*** (0.206)
Observations	22,970	8,389	22,970	8,389
R ²	0.007	0.003	0.004	0.001
Within R ²	0.007	0.003	0.005	0.001

Notes: The macro component of forecast revisions for firms' own prices. We regress firms' forecast revisions for their own prices on macro news. Columns (1) and (2) use firms' forecast revisions for their own prices as reported in the SIGE. In Columns (3) and (4), these revisions are purged of effects of recent changes in inflation. Firm-fixed effects are always included, and standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure C.1 – Representativeness of survey and sample



Notes: Distribution of firms' geographical location, employees, export orientation and sector in the full SIGE (survey) and the sample used in our baseline specification.

Table C.3 – Response of price expectations to news: Robustness

	Forecast error about firm's own prices				
	(1)	(2)	(3)	(4)	(5)
Micro News					
Forecast revision for $\pi_{t+4,t}^i$	-0.438*** (0.019)	-0.457*** (0.026)		-0.431*** (0.027)	
Forecast revision for $\pi_{t+4,t}^i$ net of $\gamma_i \Delta \pi_t$			-0.451*** (0.026)		-0.428*** (0.028)
Macro News					
Forecast revision for $\pi_{t+4,t}$	0.071* (0.043)				
Forecast revision for $\pi_{t+4,t}$ (2q v 4q)		0.305*** (0.059)	0.290*** (0.059)		
Forecast revision for $\pi_{t+4,t}$ (4q v 8q)				0.198*** (0.046)	0.190*** (0.046)
Observations	16,675	8,858	8,858	9,067	9,067
R ²	0.084	0.082	0.080	0.069	0.068
Within R ²	0.094	0.098	0.097	0.085	0.085

Notes: Alternative measures of micro and macro news. Column (1) estimates the baseline specification where we cut off the sample after 2019. Columns (2) and (3) use as macro news firms' forecast revisions for a fixed forecasting date. Specifically, we subtract from firms' current two-quarters-ahead inflation forecast their four-quarter-ahead forecasts reported two quarters ago. For Columns (4) and (5), we proceed similarly but compare current four-quarter-ahead forecasts to eight-quarter-ahead forecasts reported four quarters ago. Firm-fixed effects are always included. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.4 – Additional regression results from the SIGE

(A) without winsorizing survey variables				
	Forecast error about firm's own prices			
	(1)	(2)	(3)	(4)
Micro News				
Forecast revision for $\pi_{t+4,t}^i$ (all obs.)	-0.431*** (0.029)		-0.401*** (0.043)	
Forecast revision for $\pi_{t+4,t}^i$ net of $\gamma_i \Delta \pi_t$ (all obs.)		-0.438*** (0.022)		-0.417*** (0.032)
Macro News				
Forecast revision for $\pi_{t+4,t}$ (all obs.)	0.316*** (0.088)	0.305*** (0.089)		
Inflation surprise			3.994*** (0.792)	3.925*** (0.806)
Observations	18,284	18,284	6,640	6,640
R ²	0.065	0.067	0.053	0.056
Within R ²	0.070	0.071	0.061	0.065

Notes: Baseline specifications as in Table 3.2, except that we do not winsorize the survey variables. Standard errors are clustered at the firm level. *** p<0.01, ** p<0.05, * p<0.1.

(B) dropping top and bottom 1% of survey variables				
	Forecast error about firm's own prices			
	(1)	(2)	(3)	(4)
Micro News				
Forecast revision for $\pi_{t+4,t}^i$ (no outliers)	-0.348*** (0.016)		-0.347*** (0.033)	
Forecast revision for $\pi_{t+4,t}^i$ net of $\gamma_i \Delta \pi_t$ (no outliers)		-0.339*** (0.016)		-0.320*** (0.032)
Macro News				
Forecast revision for $\pi_{t+4,t}$ (no outliers)	0.245*** (0.038)	0.222*** (0.037)		
Inflation surprise			2.506*** (0.313)	2.424*** (0.312)
Observations	17,634	17,587	6,445	6,426
R ²	0.053	0.048	0.057	0.050
Within R ²	0.061	0.055	0.070	0.060

Notes: Baseline specifications as in Table 3.2, except that we drop the bottom and top 1% of survey variables (rather than winsorizing at these values). Standard errors are clustered at the firm level. *** p<0.01, ** p<0.05, * p<0.1.

Additional regression results from the SIGE, continued

(C) alternative definitions of micro news

	Forecast error about firm's own prices			
	(1)	(2)	(3)	(4)
Micro News				
Forecast revision for $\pi_{t+4,t}^i$ net of time-fixed effects	-0.441*** (0.018)		-0.443*** (0.031)	
Forecast revision for $\pi_{t+4,t}^i$ net of $\gamma\Delta\pi_t$		-0.424*** (0.018)		-0.423*** (0.031)
Macro News				
Forecast revision for $\pi_{t+4,t}$	0.212*** (0.050)	0.259*** (0.050)		
Inflation surprise			2.564*** (0.409)	2.967*** (0.404)
Observations	18,284	18,284	6,640	6,640
R ²	0.078	0.072	0.082	0.074
Within R ²	0.087	0.081	0.097	0.090


Notes: Specifications as in Table 3.2, using alternative definitions of news. Column (1) purges firms' forecast revisions of time-fixed effects. Column (2) purges forecast revisions of recent changes in inflation with homogeneous loading γ (instead of γ_i). In both columns, macro news is a firm's forecast revision for inflation. Columns (3) and (4) use inflation surprises as macro news, micro news as in Columns (1) and (2). Standard errors are clustered at the firm level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

(D) Univariate results

	Forecast error about firm's own prices					
	(1)	(2)	(3)	(4)	(5)	(6)
Micro News						
Forecast revision for $\pi_{t+4,t}^i$	-0.417*** (0.019)					
Forecast revision for $\pi_{t+4,t}^i$ net of $\gamma_i\Delta\pi_t$		-0.412*** (0.018)				
Macro News						
Forecast revision for $\pi_{t+4,t}$			0.150*** (0.055)			
Inflation surprise				2.756*** (0.401)		
Forecast revision for $\pi_{t+4,t}$ (6m v 12m)					0.264*** (0.056)	
Forecast revision for $\pi_{t+4,t}$ (12m v 24m)						0.157*** (0.046)
Observations	18,284	18,284	18,284	7,196	9,486	9,754
R ²	0.068	0.067	0.001	0.015	0.009	0.005
Within R ²	0.078	0.077	0.001	0.013	0.006	0.003

Notes: Univariate versions of the baseline specification. Firm-fixed effects are always included. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure C.2 – Firm-level regressions



Notes: We estimate our baseline specification (micro and macro news are firms' forecast revisions (first difference) for changes in their own prices and aggregate inflation, respectively) for each firm separately and show the univariate distribution of firm-level news coefficients. Grey area represents insignificant estimates, light green area represents estimates significant at the 10% level, dark green area indicates significance at the 5% level. Bars with less than three observations per group are not shown due to confidentiality.


C.3 IFO BUSINESS CLIMATE SURVEY

Table C.5 – Relevant questions from ifo survey

Label	Name	Question	Possible answers
Q1	Expected state of business (qualitative)	Plans and Expectations for the next 6 months: Our business situation will be	rather more favorable [1] not changing [0] rather less favorable [-1]
Q2	Expected state of business (quantitative)	Expectations for the next 6 months: In cyclical regards our state of business will be	slider with range 0 [be rather less favorable] to 100 [rather more favorable]
Q3	Realized state of business (qualitative)	Current situation: We evaluate our state of business to be	good [1] satisfiable [0] bad [-1]
Q4	Realized state of business (quantitative)	Current situation: We consider our state of business to be	slider with range good [100] to bad [0]
Q5	Realized production	Review - tendencies in [t-1]: Compared to [t-2] our production	increased [1] stayed about the same [0] decreased [-1]
Q6	Expected production	Plans and Expectations for the next 3 months: Our production is expected to be	increasing [1] not changing [0] decreasing [-1]
Q7	Macro importance	How important is the general economic development in Germany for your business situation?	very important [1] important [2] not as important [3] less important [4] unimportant [5]


Notes: Most recent wording of relevant questions from the ifo survey taken from the EBDC Questionnaire manual. t denotes the month of the survey, so in July Q5 asks about the change in June compared to May.

Figure C.3 – Average forecast revisions and production growth



Notes: Figure displays the average, seasonally adjusted forecast revision (rolling mean over 6 months) in green and year-on-year growth of manufacturing production in black (administrative data).

Figure C.4 – Distribution of firm-level responses to news




Notes: The figure shows results for expectations about firms' business situation (6-month horizon, quantitative data). Grey area represents insignificant estimates, light green area represents estimates significant at the 10% level, dark green area indicates significance at the 5% level.

Table C.6 – Macro news and forecast revisions

	$\hat{\beta}$	$SE(\hat{\beta})$
Macro News	0.008	0.001
Macro News		
× 1. Quartile by employees	0.007	0.002
× 2. Quartile by employees	0.008	0.002
× 3. Quartile by employees	0.008	0.002
× 4. Quartile by employees	0.008	0.001
Macro News		
× Firm age < 20 years	0.007	0.003
× Firm age $\geq$ 20 years	0.006	0.001
Macro News		
× Time in survey < half a year	0.015	0.007
× Time in survey $\geq$ half a year	0.008	0.001
Macro News		
× Low business-cycle exposure	0.005	0.002
× Medium business-cycle exposure	0.008	0.002
× High business-cycle exposure	0.006	0.003

Notes: Reaction of forecast revisions to macro news. Firms' forecast revisions are regressed on macro news, interaction terms, and firm-fixed effects for each interaction variable separately. For (quartiles of) the number of employees, we rely on annual questions in the ifo survey. For firm age, we rely on a one-time question about the year the firm was founded. To compute the firm age, we subtract from the year of response the year of foundation. For business-cycle exposure, we rely on a one-time question, where firms rank the importance of general economic developments in Germany for their business on a five-point scale from very important [1] to unimportant [5]. Business-cycle exposure is high when the response was very important [1], medium when the response was important [2], and low otherwise [3-5]. Standard errors are clustered at the firm level.

Figure C.5 – Response to concurrent and lagged news



Notes: Estimates based on Equation (3.2). Black lines represent point estimates, grey areas correspond to 95% confidence intervals.

Table C.7 – Production expectations—robustness

Variation	Details	Micro coeff.	Macro coeff.
1) Micro news (forecast revisions)			
Use only revisions from zero	table C.10a	−0.110***	0.030***
As above and set small errors ($\pm \frac{1}{3}$) to zero	table C.10b	−0.086***	0.023***
2) Forecast error (Bachmann, Elstner and Sims, 2013)			
Set small errors ($\pm \frac{1}{3}$) to zero	table C.10c	−0.128***	0.018***
Above only for no-change expectations	Table C.10d	−0.192***	0.018***
3) Estimation (OLS)			
Ordered logit	Table C.10e	−1.24***	0.11***
4) Macro component of forecast revision (real-time indicators)			
Fixed effect by time	Table C.10f	−0.194***	0.021***
Fixed effect by time and sector	Table C.10g	−0.196***	0.021***
5) Macro news (surprise component in ifo index)			
Surprise component in manuf. orders	Table C.10h	−0.208***	0.005***
First difference of ifo index	Table C.10i	−0.208***	0.002***
Average forecast revision	Table C.10j	−0.209***	0.345***
Average forecast revision by sector ^a	Table C.10k	−0.211***	0.216***
First difference of stock market index	Table C.10l	−0.208***	0.328***

Notes: Each row corresponds to a variation of the specification for which we report results in Table 3.4, see Appendix C.3 for details. Micro coefficient and macro coefficient are the estimates on micro and macro news. ^a In this specification, macro news is the time and sector average. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C.8 – Business situation as outcome variable (ifo survey)

	Firms' forecast errors about their business situation			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for $y_{t+6,t}^i$	-0.441*** (0.004)			
Forecast Revision for $y_{t+6,t}^i$ net of $\gamma_i \Gamma_t$		-0.453*** (0.004)	-0.450*** (0.004)	
Macro News				
Surprise component of the ifo index	0.857*** (0.044)	0.795*** (0.044)		0.697*** (0.044)
Observations	153,398	153,398	153,398	153,398
R ²	0.31864	0.30652	0.30357	0.25466
Within R ²	0.08861	0.07240	0.06845	0.00303

Notes: Results based on Equation (3.1); observations are pooled across firms. The table shows results for the quantitative business expectations (6-month horizon) for German firms as surveyed monthly in the ifo survey. Macro news is the surprise component of the ifo index. Column (1): micro news measured by forecast revisions (while controlling for macro news). Columns (2) and (3): micro news represents forecast revisions net of real-time observable aggregate developments, measured by macroeconomic indicators Γ_t with idiosyncratic reaction coefficient γ_i (see Footnote 29 for more details). All specifications include firm-fixed effects. Standard errors are clustered at the firm level. *** p<0.01, ** p<0.05, * p<0.1.

Table C.9 – Prices as outcome variable (ifo survey)

	Firms' forecast errors about their price			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for $p_{t+3,t}^i$	-0.175*** (0.001)			
Forecast Revision for $p_{t+3,t}^i$ net of $\gamma_i \Gamma_t$		-0.186*** (0.001)	-0.185*** (0.001)	
Macro News				
Surprise component of the ifo index	0.010*** (0.0005)	0.009*** (0.0005)		0.009*** (0.0005)
Observations	298,400	298,400	298,400	298,400
R ²	0.15826	0.14937	0.14746	0.07552
Within R ²	0.09112	0.08151	0.07945	0.00178

Notes: Results based on Equation (3.1); observations are pooled across firms. The table shows results for the qualitative price expectations (3-month horizon) for German firms as surveyed monthly in the ifo survey. Macro news is the surprise component of the ifo index. Column (1): micro news measured by forecast revisions (while controlling for macro news). Columns (2) and (3): micro news represents forecast revisions net of real-time observable aggregate developments, measured by macroeconomic indicators Γ_t with idiosyncratic reaction coefficient γ_i (see Footnote 29 for more details). All specifications include firm-fixed effects. Standard errors are clustered at the firm level. *** p<0.01, ** p<0.05, * p<0.1.

Table C.10 – Alternative specifications

(A) Expectations: only no-change forecasts

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for $y_{t+3,t}^i$	-0.091*** (0.003)			
Forecast Revision for $y_{t+3,t}^i$ net of $\gamma_i \Gamma_t$		-0.110*** (0.003)	-0.112*** (0.003)	
Macro News				
Surprise component of the ifo index	0.030*** (0.0008)	0.030*** (0.0008)		0.030*** (0.0009)
Observations	205,962	205,962	205,962	205,962
R ²	0.17355	0.17605	0.16728	0.16331
Within R ²	0.02310	0.02605	0.01569	0.01100

Notes: Set-up as in Table 3.4, except that we only use observations where firms expect no production change. Firm-fixed effects are always included. Standard errors are clustered at firm level. *** p<0.01, ** p<0.05, * p<0.1.

(B) Expectations: only no-change forecasts and small errors set to zero

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for $y_{t+3,t}^i$	-0.072*** (0.002)			
Forecast Revision for $y_{t+3,t}^i$ net of $\gamma_i \Gamma_t$		-0.086*** (0.002)	-0.088*** (0.002)	
Macro News				
Surprise component of the ifo index	0.024*** (0.0008)	0.023*** (0.0008)		0.024*** (0.0008)
Observations	205,962	205,962	205,962	205,962
R ²	0.14081	0.14270	0.13592	0.13288
Within R ²	0.01729	0.01945	0.01170	0.00823

Notes: Set-up as in Table 3.4, except that we only use observations where firms expect no production change and set small forecast errors ($\pm \frac{1}{3}$) to zero. Firm-fixed effects are always included. Standard errors are clustered at firm level. *** p<0.01, ** p<0.05, * p<0.1.

Alternative specifications, continued.

(C) Forecast error: small errors set to zero

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for $y_{t+3,t}^i$	-0.115*** (0.001)			
Forecast Revision for $y_{t+3,t}^i$ net of $\gamma_i \Gamma_t$		-0.128*** (0.002)	-0.128*** (0.002)	
Macro News				
Surprise component of the ifo index	0.018*** (0.0006)	0.018*** (0.0006)		0.018*** (0.0006)
Observations	302,737	302,737	302,737	302,737
R ²	0.11352	0.11278	0.10838	0.07974
Within R ²	0.04103	0.04022	0.03547	0.00449

Notes: Set-up as in Table 3.4, except small forecast errors ($\pm \frac{1}{3}$) are set to zero. Firm-fixed effects are always included. Standard errors are clustered at firm level. *** p<0.01, ** p<0.05, * p<0.1.

(D) Forecast error: small errors set to zero if no change expected

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for $y_{t+3,t}^i$	-0.176*** (0.001)			
Forecast Revision for $y_{t+3,t}^i$ net of $\gamma_j \Gamma_t$		-0.192*** (0.001)	-0.191*** (0.001)	
Macro News				
Surprise component of the ifo index	0.018*** (0.0006)	0.018*** (0.0006)		0.017*** (0.0006)
Observations	302,737	302,737	302,737	302,737
R ²	0.14684	0.14143	0.13768	0.07495
Within R ²	0.08113	0.07529	0.07125	0.00369

Notes: Set-up as in Table 3.4, except that we set small forecast errors ($\pm \frac{1}{3}$) to zero when expectations are zero. Firm-fixed effects are always included. Standard errors are clustered at firm level. *** p<0.01, ** p<0.05, * p<0.1.

Alternative specifications, continued.

(E) Estimation: Ordered Logit rather than OLS

Term	Estimate	Standard Error	t-value	Coefficient type	exp(estimate)
Micro News	-1.24	0.01	-158.19	coefficient	0.29
Macro News	0.11	0.00	37.16	coefficient	1.12
-4/3 -1	-6.04	0.03	-173.89	scale	0.00
-1 -2/3	-3.56	0.01	-337.00	scale	0.03
-2/3 -1/3	-2.45	0.01	-370.14	scale	0.09
-1/3 0	-1.27	0.00	-280.89	scale	0.28
0 1/3	1.52	0.00	314.78	scale	4.57
1/3 2/3	2.71	0.01	373.96	scale	15.10
2/3 1	3.91	0.01	321.66	scale	49.88
1 4/3	6.66	0.05	144.17	scale	782.37

Notes: Results using ordered logit to estimate the effect of micro news and macro news on the production forecast error. The last column shows the odds ratios. Rows 3 to 10 depict the cut points of the latent variable. The full, pooled sample is used.

(F) Micro news: absorb macro component of forecast revision with time-fixed effect

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for $y_{t+3,t}^i$	-0.191*** (0.001)			
Forecast Revision for $y_{t+3,t}^i$ net of $\gamma_i \Gamma_t$		-0.194*** (0.001)	-0.194*** (0.001)	
Macro News				
Surprise component of the ifo index	0.022*** (0.0007)	0.021*** (0.0007)		0.021*** (0.0007)
Observations	302,737	302,737	302,737	302,737
R ²	0.16260	0.16471	0.16015	0.08967
Within R ²	0.08471	0.08701	0.08202	0.00498

Notes: Set-up as in Table 3.4, except that we absorb the macro component from forecast revisions by means of time-fixed effects. Firm-fixed effects are always included. Standard errors are clustered at firm level. *** p<0.01, ** p<0.05, * p<0.1.

Alternative specifications, continued.

(G) Micro news: absorb macro comp. of forecast revision with time-sector-fixed effect

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for $y_{t+3,t}^i$	-0.191*** (0.001)			
Forecast Revision for $y_{t+3,t}^i$ net of $\gamma_i \Gamma_t$		-0.196*** (0.001)	-0.196*** (0.001)	
Macro News				
Surprise component of the ifo index	0.022*** (0.0007)	0.021*** (0.0007)		0.021*** (0.0007)
Observations	302,737	302,737	302,737	302,737
R ²	0.16260	0.16555	0.16100	0.08967
Within R ²	0.08471	0.08793	0.08295	0.00498

Notes: Set-up as in Table 3.4, except that we absorb the macro component from forecast revisions by means of time-sector-fixed effects. Firm-fixed effects are always included. Standard errors are clustered at firm level. *** p<0.01, ** p<0.05, * p<0.1.

(H) Macro news: manufacturing orders rather than ifo index

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for $y_{t+3,t}^i$	-0.190*** (0.001)			
Forecast Revision for $y_{t+3,t}^i$ net of $\gamma_i \Gamma_t$		-0.208*** (0.001)	-0.208*** (0.001)	
Macro News				
Surprise component of the ifo index	0.005*** (0.0003)	0.005*** (0.0003)		0.005*** (0.0003)
Observations	298,586	298,586	298,586	298,586
R ²	0.15828	0.15383	0.15286	0.08580
Within R ²	0.08023	0.07536	0.07431	0.00103

Notes: Set-up as in Table 3.4, except macro news is constructed from the median professional forecast of manufacturing orders. Firm-fixed effects are always included. Standard errors are clustered at firm level. *** p<0.01, ** p<0.05, * p<0.1.

Alternative specifications, continued.

(1) Macro news: first difference of ifo index rather than ifo index surprise

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for $y_{t+3,t}^i$	-0.190*** (0.001)			
Forecast Revision for $y_{t+3,t}^i$ net of $\gamma_i \Gamma_t$		-0.208*** (0.001)	-0.208*** (0.001)	
Macro News				
Surprise component of the ifo index	0.002*** (0.0002)	0.002*** (0.0003)		0.001*** (0.0003)
Observations	301,185	301,185	302,737	301,185
R ²	0.15737	0.15318	0.15313	0.08505
Within R ²	0.07908	0.07450	0.07435	0.00004

Notes: Set-up as in Table 3.4, except macro news is constructed with the first difference of the ifo index. Firm-fixed effects are always included. Standard errors are clustered at firm level. *** p<0.01, ** p<0.05, * p<0.1.

(J) Macro news: average forecast revisions rather than ifo index

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for $y_{t+3,t}^i$	-0.194*** (0.001)			
Forecast Revision for $y_{t+3,t}^i$ net of $\gamma_i \Gamma_t$		-0.209*** (0.001)	-0.208*** (0.001)	
Macro News				
Surprise component of the ifo index	0.502*** (0.019)	0.345*** (0.018)		0.308*** (0.018)
Observations	302,737	302,737	302,737	302,737
R ²	0.16186	0.15526	0.15313	0.08681
Within R ²	0.08389	0.07668	0.07435	0.00187

Notes: Set-up as in Table 3.4, except macro news is constructed with average production forecast revisions. Firm-fixed effects are always included. Standard errors are clustered at firm level. *** p<0.01, ** p<0.05, * p<0.1.

Alternative specifications, continued.

(K) Macro news: average forecast revisions for each sector rather than ifo index

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for $y_{t+3,t}^i$	-0.196*** (0.001)			
Forecast Revision for $y_{t+3,t}^i$ net of $\gamma_i \Gamma_t$		-0.211*** (0.001)	-0.208*** (0.001)	
Macro News				
Surprise component of the ifo index	0.326*** (0.013)	0.216*** (0.011)		0.129*** (0.012)
Observations	302,737	302,737	302,737	302,737
R ²	0.16169	0.15506	0.15313	0.08580
Within R ²	0.08371	0.07646	0.07435	0.00076

Notes: Set-up as in Table 3.4, except macro news is constructed with average production forecast revisions for each sector. Firm-fixed effects are always included. Standard errors are clustered at firm level. *** p<0.01, ** p<0.05, * p<0.1.

(L) Macro news: first difference of stock market index rather than ifo index surprise

	Firms' forecast errors about their production			
	(1)	(2)	(3)	(4)
Micro News				
Forecast Revision for $y_{t+3,t}^i$	-0.190*** (0.001)			
Forecast Revision for $y_{t+3,t}^i$ net of $\gamma_i \Gamma_t$		-0.208*** (0.001)	-0.208*** (0.001)	
Macro News				
Surprise component of the ifo index	0.371*** (0.014)	0.328*** (0.014)		0.328*** (0.014)
Observations	302,737	302,737	302,737	302,737
R ²	0.15999	0.15518	0.15313	0.08716
Within R ²	0.08185	0.07659	0.07435	0.00224

Notes: Set-up as in Table 3.4, except macro news is constructed with the first difference of the German stock market index DAX. Firm-fixed effects are always included. Standard errors are clustered at firm level. *** p<0.01, ** p<0.05, * p<0.1.

D Appendix to Chapter 4

D.1 IFO BUSINESS SURVEY: SURVEY QUESTIONS (TRANSLATED)

Business expectation (monthly)

- Expectations about the next six months: our expected business state in economic terms: [continuous slider from 0 (rather worse) over 50 (the same) to 100 (rather better)]

Business uncertainty (monthly)

- We estimate the uncertainty regarding our business expectations in the next six months as: [continuous slider from 0 (low) over 50 (average) to 100 (high)]

Business state (monthly)

- We assess our current business state as: [continuous slider from 0 (bad) over 50 (satisfactory) to 100 (good)]

Investment plans (November wave)

- In the next year ($t+1$ compared to t) our investments will likely ...
increase o not change o decrease o

Sales expectations (October wave): short run

- By how much in percentage terms will your sales change in the fourth quarter of 2023 compared to the third quarter of 2023?
 - In the best possible case:

- In the worst possible case:
- Taking into account all contingencies and risks, I expect for the second quarter of 2014 all in all a change of:

Sales expectations (October wave): medium run


- By how much in percentage terms will your sales change in 2024 compared to 2023?
 - In the best possible case:
 - In the worst possible case:
 - Taking into account all contingencies and risks, I expect for the second quarter of 2014 all in all a change of:

Sales expectations (October wave): long run

- By how much in percentage terms will your sales change on average in the three years 2025 to 2027? (average annual rate of change)
 - In the best possible case:
 - In the worst possible case:
 - Taking into account all contingencies and risks, I expect for the second quarter of 2014 all in all a change of:


D.2 ADDITIONAL FIGURES

Figure D.1 – Firms’ macroeconomic tail risk expectations and mean GDP growth expectations




Notes: The figure relates subjective firm-level expectations of the occurrence of a large decline in annual GDP ($\geq 5\%$) in the next five years, elicited in October 2022, to first moment GDP growth expectations for 2020 (in orange) and 2021 (in blue), elicited in August 2020. The lines represent the prediction from a kernel-weighted local polynomial regression. Confidence intervals refer to the 95% level. The long dashed lines show the GDP forecast for Germany of the Bundesbank from June 2020 (Deutsche Bundesbank, 2020) and the short dashed lines show the GDP forecast for Germany of the median professional forecaster from June 2020 (Focus Economics, 2020).

Figure D.2 – Relation of subjective probability of macro tail event occurrence and firm exposure beliefs




Notes: This figure displays binned scatter plots between the subjective probability of the occurrence of a large decline in annual GDP ($\geq 5\%$) in the next five years and firm exposure expectations to a drop in GDP of 5% in terms of real sales changes. The figures in the first column relate subjective firm-level expectations of the occurrence of a large decline in annual GDP ($\geq 5\%$) in the next five years elicited in October 2022, April 2023, and October 2023 to firm exposure expectations to a drop in GDP of 5% in terms of real sales changes elicited in October 2023. The figures in the second column relate the change in subjective probability of the occurrence of a large decline in annual GDP ($\geq 5\%$) in the next five years from October 2022 to October 2023 to firm-specific exposure expectations to a drop in GDP of 5% in terms of real sales changes elicited in October 2023. The first row uses the expected worst-case exposure scenario, the second row the expected average exposure scenario, and the third row the expected best-case exposure scenario, see Section 4.2.3 for more details on firm exposure beliefs.

Figure D.3 – Relation of macro tail risk to unemployment rate expectations




Notes: This figure displays binned scatter plots between subjective firm-level expectations of the occurrence of a large decline in annual GDP ($\geq 5\%$) in the next five years (elicited in October 2022/2023) and subjective firm-level unemployment rate expectations one year ahead (first row: elicited in September 2022/2023, second row: elicited in December 2022/2023). Panel (B) and (D) absorb firm fixed effects.

Figure D.4 – Relation of macro tail risk to unemployment rate uncertainty




Notes: This figure displays binned scatter plots between subjective firm-level expectations of the occurrence of a large decline in annual GDP ($\geq 5\%$) in the next five years (elicited in October 2022/2023) and subjective firm-level uncertainty about the unemployment rate one year ahead (first row: elicited in September 2022/2023, second row: elicited in December 2022/2023), measured on a scale from 1 (low) to 5 (high). Panel (B) and (D) absorb firm fixed effects.

**Figure D.5 – Firms’ macro tail risk expectations and investment plans:
All four macro tail risk scenarios**




Notes: This figure shows the point estimates and corresponding 90% confidence interval based on Equation (4.4) with and without fixed effects for the four different macro tail risk scenarios. The dependent variable is an indicator that is one for a planned decline in investment in the next year compared to the previous year, elicited in the November wave 2022 and 2023. The macro tail risk variables are firm-level expectations of occurrence in the next five years, elicited in the October wave 2022 and 2023. Section 4.2 discusses the macro tail risk expectations questions in more detail. Standard errors are clustered at the firm level.

Figure D.6 – Firms’ macro tail risk expectations and investment plans; only within firm variation - split by average exposure




Notes: This figure shows binned a scatterplot based on Equation (4.4). The dependent variable is an indicator that is one for a planned decline in investment in the next year compared to the previous year, elicited in the November wave 2022 and 2023. The explanatory variable shown on the horizontal axis is subjective probability of the occurrence of a large decline in annual GDP ($\geq 5\%$) in the next five years, elicited in the October wave 2022 and 2023. All explanatory variables are interacted with an indicator for above median average-case exposure to a macro tail event. Standard errors are clustered at the firm level.

Figure D.7 – Business expectations - only within firm variation




Notes: This figure shows binned a scatterplot based on Equation (4.4). The dependent variable is a slider (0-100) on business expectations for the next six months, elicited in the October wave 2022 and 2023. The explanatory variable shown on the horizontal axis is subjective probability of the occurrence of a large decline in annual GDP ($\geq 5\%$) in the next five years, elicited in the October wave 2022 and 2023. All explanatory variables are interacted with an indicator for above median average-case exposure to a macro tail event. Standard errors are clustered at the firm level.

Figure D.8 – Robustness check: Business expectations and uncertainty for different horizons




Notes: This figure displays robustness checks of the baseline regression specification, see Equation (4.4). “baseline” refers to the baseline specification, see Table D.3. In the other specifications, controls include quantitative sales growth expectations (proxy for subjective first moments), sales growth uncertainty (measured as the span between best- and worst-case scenarios, a proxy for subjective second moments), and the realized growth rate in the previous quarter to control for current conditions. ‘short-term’ refers to forecasts for the next quarter, ‘mid-term’ refers to forecast for the next year, and ‘long-term’ refers to average yearly sales growth expectations two to four years ahead. The transparent red estimates are based on the smaller sample of firms that also report long-term sales growth expectations. 90% confidence intervals are shown. Standard errors are clustered at the firm level.

Figure D.9 – Robustness check: First and second moment macro controls (unemployment rate)




Notes: This figure displays robustness checks of the baseline regression specification, see Equation (4.4). Unemployment rate expectations controls consist of firms’ expectations of the unemployment rate in September and December over a 12-month horizon and the certainty of that forecast on a scale of 1 (low) to 5 (high). 90% confidence intervals are shown. Standard errors are clustered at the firm level.

Figure D.10 – Vignette: Baseline outlooks




Notes: This figure shows the baseline line vignette. Both outlooks have the same mean (2) and variance (9), but different skewness (A: neutral skewness 0 ; B: negatively skewed -2.6). In Section 4.3.3, the corresponding wording of the question is displayed.

Figure D.11 – Vignette: Shifted outlooks




Notes: This figure shows the alternative, shifted vignette. Both outlooks have the same mean (4) and variance (9), but different skewness (A: neutral skewness 0 ; B: negatively skewed -2.6). In Section 4.3.3, the corresponding wording of the question is displayed.

Figure D.12 – Local projections: Impact on production of investment goods (other macro moments)



Notes: This figure shows local projections of investment good production in response to an interquartile-range increase in the average macro expectations index (*MacroAvg* index) in the left panel and to an interquartile-range increase in the upside risk macro expectations index (*MacroUp* index) in the right panel based on Equation 4.5. Confidence bands are shown at the 68% and 90% level based on lag augmented standard errors (Montiel Olea and Plagborg-Møller, 2021).

Figure D.13 – Local projections: Impact on production of investment goods (robustness)



Notes: This figure shows local projections of investment good production in response to an interquartile-range increase in the macro tail risk expectations index (*MacroTail* index) based on Equation 4.5. In Panel (A), it is controlled for the average macro expectations index (*MacroAvg* index), and in Panel (B) additionally for the macro upside risk expectations index (*MacroUp* index). The solid blue line shows the baseline estimates from Figure 4.10. Confidence bands are shown at the 90% level based on lag augmented standard errors (Montiel Olea and Plagborg-Møller, 2021). The short (long) dashed red (green) line shows the results based on a specification of three (one) lags of all variables. The long dashed orange line includes a dummy for the Great Recession period. The dotted dashed yellow line controls for the European Economic Policy Uncertainty Index (available since 1987). The dotted brown line controls for the log of the DAX (German stock market index) and the German VIX (standard deviation of the daily realized returns of the DAX for each month).

D.3 ADDITIONAL TABLES

Table D.1 – Relation of macro tail risk expectations and sales expectations

	Sales Expectations		Sales Uncertainty		
	(1)	(2)	(3)	(4)	(5)
Subj. Prob. GDP Drop $\geq 5\%$	-0.91*** (0.17)	-0.35* (0.20)	1.33*** (0.19)	0.36* (0.20)	0.39* (0.21)
Average Sales Forecast					0.081** (0.033)
Subj. Prob. Financial Crisis	-0.72*** (0.16)	-0.023 (0.18)	1.32*** (0.19)	0.47** (0.20)	0.47** (0.20)
Average Sales Forecast					0.080** (0.032)
Subj. Prob. Pandemic	-0.18 (0.17)	0.12 (0.19)	0.18 (0.18)	0.26 (0.20)	0.25 (0.20)
Average Sales Forecast					0.079** (0.033)
Subj. Prob. Geop. Escalation	-0.45*** (0.14)	-0.14 (0.18)	0.76*** (0.16)	0.15 (0.18)	0.16 (0.18)
Average Sales Forecast					0.080** (0.033)
Observations	5268	5268	5268	5268	5268
Q-on-Q Sales Growth ($t - 1$)	Yes	Yes	Yes	Yes	Yes
Firm FE	No	Yes	No	Yes	Yes
Year X Industry FE	Yes	Yes	Yes	Yes	Yes

Notes: The table presents the estimation results based on linear regressions of firms' subjective probability of a GDP drop of at least 5% / financial crisis / pandemic / geopolitical escalation in the next five years on firms' sales expectations for the next year (Columns 1-2) and firms' sales uncertainty for the next year (Columns 3-5). Sales uncertainty is measured as the span between best- and worst-case scenarios. Appendix D.1 provides the wording of the questions. The macro tail risk expectations are normalized by their cross-sectional standard deviation. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table D.2 – Distribution of firms’ expected exposure to a 5% decline in GDP

	Change in real sales, when GDP drops by 5%		
	Worst-case scenario	Average-case scenario	Best-case scenario
p10	-5%	0%	5%
p25	-15%	-5%	0%
p50	-25%	-10%	-5%
p75	-40%	-20%	-10%
p90	-60%	-30%	-20%

Notes: This table presents the expected worst-case, average-case, and best-case exposure scenarios in case of a 5% decline in German GDP in terms of changes in real sales for firms across different percentiles of the exposure distribution. The survey question is shown in Section 4.2.

Table D.3 – Macroeconomic tail risk expectations and investment decisions

	Indicator for a Decline in Investment					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel (a): Baseline Controls (Visual Analog Sliders)</i>						
Prob. GDP Drop $\geq 5\%$	0.0025*** (0.00031)	0.0014*** (0.00030)	0.0012*** (0.00045)	0.0012** (0.00048)		
× low worst-case β					0.00015 (0.00060)	0.000026 (0.00073)
× high worst-case β					0.0020*** (0.00064)	0.0018*** (0.00069)
Observations	2982	2982	2982	2982	2982	2982
<i>Panel (b): Alternative Controls (Quantitative Sales Expectations)</i>						
Prob. GDP Drop $\geq 5\%$	0.0023*** (0.00044)	0.0019*** (0.00042)	0.0015*** (0.00055)	0.0018*** (0.00063)		
× low worst-case β					0.00090 (0.00070)	0.00031 (0.00086)
× high worst-case β					0.0020** (0.00082)	0.0026*** (0.00094)
Observations	1372	1372	1372	1372	1372	1372
Controls	No	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	No	Yes	No
Firm FE	No	No	Yes	Yes	Yes	Yes
Year X Industry FE	No	No	No	Yes	No	Yes

Notes: This table presents the regression results based on Equation (4.4). The dependent variable is an indicator that is one for a planned decline in the next year compared to the previous year, elicited in the November wave 2022 and 2023. The macro tail risk variables are firm-level expectations of an annual GDP drop $\geq 5\%$ in the next five year, elicited in the October wave 2022 and 2023. In Panel (A), controls include business expectations (proxy for first moment expectations), business uncertainty (proxy for second moment expectations), and business situation, all measured on a visual analog scale ranging from 0 to 100. In Panel (B), controls include quantitative sales expectations for the next year (proxy for first moment expectations), sales uncertainty (measured as the span between best- and worst-case scenarios, a proxy for second moment expectations), and the realized growth rate in the previous quarter. Appendix D.1 provides the wording of the questions. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table D.4 – Revisions in macroeconomic tail risk expectations and investment decisions

	Downward Investment Revision Indicator			
	(1)	(2)	(3)	(4)
Δ GDP Drop $\geq 5\%$	0.00088** (0.00038)			
Δ Fin. Crisis		0.0016*** (0.00046)		
Δ Pandemic			0.0016** (0.00063)	
Δ Geop. Escalation				0.00074* (0.00045)
Δ Business expectations	-0.00053 (0.00068)	-0.00049 (0.00068)	-0.00064 (0.00068)	-0.00066 (0.00068)
Δ Business uncertainty	0.00011 (0.00062)	0.00011 (0.00061)	0.00014 (0.00062)	0.00014 (0.00062)
Δ Business situation	-0.0020*** (0.00058)	-0.0020*** (0.00058)	-0.0020*** (0.00058)	-0.0020*** (0.00058)
Constant	0.24*** (0.015)	0.25*** (0.015)	0.23*** (0.013)	0.23*** (0.012)
Observations	1433	1433	1433	1433

Notes: This table shows regression results based on revisions of macroeconomic tail risk expectations and revisions of investment plans. The dependent variable is an indicator that is one if a firm invested less in 2023 and initially planned in November 2022. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table D.5 – Expected probability of a large decline in GDP and quantitative investment

	$\Delta \log \text{Investment (2022 - 2021)}$				$\Delta \log \text{Investment (2023 - 2021)}$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Prob. GDP $\geq$ 5% (Oct. 2022)	-0.0019*** (0.00056)	-0.0019*** (0.00056)	-0.0017*** (0.00063)	-0.0021*** (0.00066)			0.00074 (0.0010)	0.00076 (0.0011)
Prob. GDP $\geq$ 5% (Apr. 2023)			-0.00064 (0.00081)	-0.00030 (0.00093)	-0.0032** (0.0012)	-0.0032*** (0.0012)	-0.0028** (0.0013)	-0.0021 (0.0015)
Prob. GDP $\geq$ 5% (Oct. 2023)				-0.00051 (0.00089)				-0.0012 (0.0014)
$\Delta \log \text{Sales (2022 - 2021)}$		0.54*** (0.100)	0.54*** (0.10)	0.59*** (0.10)				
$\Delta \log \text{Sales (2023 - 2021)}$						0.48*** (0.081)	0.48*** (0.081)	0.51*** (0.085)
Observations	1961	1961	1961	1728	1557	1557	1557	1444
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Additional controls	No	No	Yes	Yes	No	No	Yes	Yes
R^2	0.13	0.14	0.15	0.16	0.16	0.18	0.18	0.20

Notes: The table shows the regression results based on the expected probability of a large decline in GDP and change of log quantitative investment volumes from 2021 to 2022 (Columns 1-4) and from 2021 to 2023 (Columns 5-8), elicited in March 2023 and 2024 respectively. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table D.6 – Correlations of macro beliefs based on newspaper data and other indices

	<i>MacroTail</i>	<i>MacroAvg</i>	<i>MacroUp</i>
<i>MacroTail</i>	1.00	-0.72	-0.85
<i>MacroAvg</i>	-0.72	1.00	0.74
<i>MacroUp</i>	-0.85	0.74	1.00
VIX (DAX)	0.29	-0.50	-0.25
EPU (Europe)	0.56	-0.15	-0.54
VIX (S&P 500)	0.36	-0.51	-0.33
CBOE SKEW (S&P 500)	0.11	0.33	-0.03

Notes: This table presents correlation coefficients between macro expectations, classified by GPT 3.5, and other uncertainty indices, at the monthly level. All correlations are significant at the 5% level, except the correlation between CBOE SKEW (S&P 500) and *MacroUp*.

Table D.7 – Macro expectations and news topics

	(1)	(2)	(3)	(4)	(5)	(6)
	<i>MacroTail</i> $\neq 0$		<i>MacroAvg</i> $\neq 0$		<i>MacroUp</i> $\neq 0$	
<i>Relative to news about the aggregate economy:</i>						
Financial	-0.0023 (0.0023)	-0.0032 (0.0023)	-0.021*** (0.0027)	-0.025*** (0.0027)	-0.017*** (0.0040)	-0.020*** (0.0040)
Idiosyncratic	-0.0096*** (0.0018)	-0.0097*** (0.0018)	0.020*** (0.0021)	0.019*** (0.0021)	-0.065*** (0.0031)	-0.063*** (0.0031)
International Politics	-0.039*** (0.0020)	-0.039*** (0.0020)	-0.086*** (0.0023)	-0.085*** (0.0023)	-0.084*** (0.0035)	-0.081*** (0.0034)
National Politics	-0.024*** (0.0026)	-0.024*** (0.0026)	-0.098*** (0.0031)	-0.097*** (0.0031)	-0.018*** (0.0046)	-0.018*** (0.0046)
Observations	187342	187342	187342	187342	187342	187342
Time FE	No	Yes	No	Yes	No	Yes

Notes: This table presents regressions that relate a change in macro expectations, classified by GPT 3.5, to news topic, classified by a dictionary approach, at the article level. Table D.8 lists the dictionaries. Robust standard errors in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table D.8 – Dictionary for classifying articles into news topics

Topic	Words
Aggregate economy	business cycle, GDP, inflation, unemployment, employment, growth, interest rate, industry, monetary policy, budget, government spending, tax cut, consumer spending, trade, export, import, trade balance, current account, economic forecast, stimulus package, forecasts, forecast, gross capital investment, national debt, money supply, labor market, wage development, productivity, economic policy, economic reforms, market economy, social market economy, globalization, international relations, world economy, economic development, industrial production, investment, trade agreement, export promotion, economic research, business survey, corporate profits, fiscal policy, central bank, economic report, economic analysis, economic advisor, economic development, economic reporting, labor costs, trade volume, investment activity, deficit, price, commodity prices, commodity, environmental economics, job security, technology sector, consumer prices, business climate, money market policy, public debt, economic indices, economic research institutes, trade wars, economic data, infrastructure investments, inflation rate, trade policy, production costs, business friendliness, labor productivity, trading partner, industrial development, social expenditures, price stability, credit conditions, employment prospects, job growth, economic relations, capital flow, labor shortage, fiscal discipline, industrial policy, employment level, industrial location, economic outlook, consumer sentiment, budget deficit, economic power, economic indicators, domestic product, technology, finance, energy, healthcare, automotive, real estate, bubble, finance minister, economic minister, central bank governor, economy, G20, consumption, dollar, DAX, Dow Jones, Nasdaq, trade minister, World Bank, production, ECB, USA, China, money, recession, boom, crisis, worldwide.
Financial economy	stock exchange, stock market, stocks, finance, financial market, financial crisis, securities, investor, investment, investors, trading, prices, index, indices, performance, dividends, investment strategy, portfolio, asset allocation, risk management, return, ECB, Fed, interest rate, interest rate policy, volatility, volatile, financial instruments, capital, capital market, investment, financial products, investment funds, central bank, monetary policy, ETF, options, futures, certificates, derivatives, hedge funds, commodities, bonds, precious metals, foreign exchange market, forex trading, currencies, currency, exchange rates, forex, asset classes, asset management, investment strategies, trading, securities trading, stock exchange trading, financial analysis, market analysis, chart analysis, fundamental analysis, technical analysis, trend, market movements, stock prices, market news, financial news, financial reporting, analyst recommendations, stock exchange reports, trading volume, trading activities, investor behavior, capital flow, investor sentiment, trading strategies, investment opportunities, investment potential, return expectations, risk appetite, investment opportunities, market efficiency, return on capital, trading platforms, financial institutions, market liquidity, trading fees, financial market regulation, financial market participants, market trends, stock performance, investment goals, market capitalization, investor portfolios, portfolio, investment decisions, trading strategies, market volatility, investment strategies, investment opportunities, stock indices, stock valuations, stockbroker, financial advisor, trading algorithms, investment vehicles, investment, banks, bank, banking supervision, investment instruments, financial market analysts, investment portfolios, investment options, stock traders, market developments, trading platforms, capital management, financial market experts, securities markets, cryptocurrencies, Bitcoin, DAX, Deutsche Bank, Commerzbank, Bundesbank, BaFin, savings banks, cooperative banks, HypoVereinsbank, KfW Bankengruppe, Allianz, Siemens, BMW, SAP, Volkswagen, Deutsche Börse, Deutsche Telekom, Bayer, BASF, Deutsche Post, Lufthansa

Idiosyncratic (firms)	company, firm, corporation, CEO, board of directors, shareholder, share, revenue, profit, loss, merger, merger negotiations, takeover, insolvency, restructuring, expansion, layoffs, product, market, new hires, production outage, market launch, market withdrawal, business model, industry leader, corporate strategy, product development, market position, competitive advantage, market leadership, market share, growth prospects, corporate goals, business unit, corporate governance, annual report, analyst report, company valuation, stock price, dividend, invest, capital increase, financial result, balance sheet, quarterly report, annual financial statement, fiscal year, liquidity, capital, auditor, compliance, accounting, corporate policy, revenue development, corporation, stake, legal, lawyer, cost structure, economic situation, business partner, customer relations, employer, employer side, supply chain, supplier, sales channel, marketing strategy, advertising campaign, customer loyalty, marketing, customer satisfaction, brand image, product quality, customer service, customer acquisition, complaint management, sales strategy, trade partner, contract conclusion, partnership, joint venture, alliance, cooperation, supply contract, exclusivity, negotiating partner, contract extension, contract design, compensation, wage negotiation, bonus payment, collective bargaining, salary structure, salary increase, performance incentive, compensation package, salary policy, works council, employee survey, working time model, personnel management, employee development, further training, occupational safety, working conditions, employee satisfaction, Volkswagen, Siemens, BMW, Daimler, Adidas, Bosch, Lufthansa, Deutsche Bank, SAP, BASF, Apple, Amazon, Microsoft, Google, Facebook, Tesla, Alibaba, Tencent, industry, sector, service, Samsung, Toyota, retail, food, health, pharmaceuticals, healthcare services, telecommunications, entertainment, media, consumer electronics, automotive, technology, steel, metal, software, fashion, mail order, e-commerce, investors, customers, industry, board of directors, profitability, GmbH, factory
International politics	foreign minister, European Union, international relations, refugee, conflict, Erdoğan, Moscow, sanction, worldwide, Russia, Israel, France, Iran, Turkey, UN, soldiers, war, politics, Trump, China, Chinese, president, Russian, world, Ukraine, American, Putin, international
National politics	president, prime minister, federal chancellor, chancellor, head of state, finance minister, trade minister, head of government, law, bill, parliament, opposition leader, party chairperson, G20, federal government, government, politics, banks, coalition, parliament, bill, elections, election, party, politician, ministerial council, political parties, government program, opposition leader, members of parliament, election campaign, constitution, democracy, federal president, domestic politics, social policy, environmental policy, education policy, tax policy, social minister, environment minister, education minister, health policy, security policy, immigration policy, defense policy, state affairs, coalition negotiations, cabinet meeting, opposition party, legislation, vote, federal council, Bundestag, governing coalition, head of state, interior minister, justice minister, family policy, labor market policy, Merkel, SPD, CDU, union, FDP, Greens, environment, regulation, expenditures, federal bank, budget, tax, union, decision, municipality, community

Notes: This table presents the dictionaries that classify articles into news topics. The words were generated from ChatGPT with the prompt: “Imagine you are reading a Handelsblatt article and want to distinguish whether the article is about macroeconomic developments, idiosyncratic news, purely political developments or articles about the financial market. Name 50 words that you would interpret as an indicator for a XXX article.”. Additionally, I added words from prompts asking for relevant companies etc. Finally, I separated politics into national and international by prompting “Look at this word list and extract all keywords that indicate news exclusively on German politics and words that indicate international politics:”.


D.4 DRIVERS OF LEVEL DIFFERENCES BETWEEN FIRMS' MACROECONOMIC TAIL RISK EXPECTATIONS

What are the drivers of the level differences in macro tail risk expectations across firms? To answer this question, I regress the firm fixed effects on a set of firm characteristics and additionally control for industry fixed effects. Figure D.14 shows the respective point estimates and 95% confidence interval of the joint regression.

I find that standard firm characteristics such as firm size, whether the firm is family owned or not, and whether it has a large export share are not significantly correlated with average macro tail risk expectations. Firms in former East Germany have, on average, slightly higher macro tail risk expectations, especially for the financial crisis scenario and the geopolitical escalation scenario. Also, being more affected by the COVID-19 pandemic is associated with higher macro tail risk expectations across the board. This is suggestive evidence of scarring effects: Experiencing a tail event leads to a particularly strong update of tail risk expectations. Low trust in experts advising policymakers is also positively associated with higher macro tail risk expectations. This is consistent with the idea that firms differ in their confidence that policymakers can avert a major crisis. In this context, it also makes sense that the pandemic scenario shows a lower correlation, as pandemics themselves are difficult to avoid.

Finally, the regression results show that firms that were founded during the turbulent period 1920-1945 (Hyperinflation/Great Depression/WW2) are significantly more likely to expect a large GDP decline or a financial crisis than firms founded between 1946-1989, a period without a realized macro tail event. This is suggestive evidence of institutional memory: While managers from 1920-1945 are no longer in the firms, the potentially turbulent founding histories are embedded in the firm's culture (Malmendier and Nagel, 2011; Bouwman and Malmendier, 2015). Note that the R^2 of this regression is still relatively low with 4.5-7.9% . This means that the level differences between firms are mostly due to other unobserved characteristics.

Figure D.14 – Drivers of level differences between firms’ macroeconomic tail risk expectations



Notes: The figure shows the coefficients of a joint regression, relating the fixed effect of firms’ subjective macro tail risk expectations—elicited in October 2022, April 2023, and October 2023—to firm characteristics. The question on the respective macro tail risk is shown in the text. It is controlled for 2-digit NACE industry fixed effects. Confidence intervals refer to the 95% level based on robust standard errors.

D.5 FIRMS’ EXPOSURE BELIEFS: EVIDENCE FROM AN OPEN-ENDED SURVEY QUESTION

This section documents the results of an open-ended survey text question on firms’ exposure beliefs to a 5% change in GDP in Germany. In Section 4.2.3, I introduce the survey question and present the main findings.

I received 2220 answers and read through all of them. Based on that, I created a codebook that classifies the responses into different categories. A research assistant then hand-coded all responses into the categories according to the codebook. I exclude firms that are categorized in the groups “politics”, “general pessimism”, or “other” because these managers mostly use the text box to complain about their current situation or to comment on something else unrelated to macro tail risk exposure.

Assessability There are 545 managers who directly report on the assessability of their exposure to a 5% drop in GDP. The vast majority of these responses (93%) report significant difficulties in predicting their exposure, with close to 40% of them even stating that

it is (almost) impossible to predict. This finding aligns with the concept of Knightian uncertainty, where firms are unable to quantify the firm-specific risk associated with a macro tail event.

Exposure in numbers A significant portion of firms also provided numerical estimates regarding expected sales changes in this macro tail event. More than half of these responses (54%) presented non-average exposure scenarios, such as ranges or max/min scenarios. Meanwhile, 29% provided a single number, and 17% offered a single number while explicitly noting the associated uncertainty. This further reflects the uncertainty surrounding the calculation of an average exposure.

Exposure strength (in words) A large share of firms express the estimated strength of their potential exposure. Among those firms, 14% expect close to zero exposure, 79% anticipate a negative impact, and 7% even foresee a positive impact. Within the negative impact group, the majority of firms expect a potentially disproportionately strong impact.

Reference to past events When firms are asked about sharing their thoughts on exposure, they often refer to their realized exposure in past macro tail events. The COVID-19 pandemic is referenced twice as much as the Great Recession, which is consistent with learning from experience and a higher weight of more recent experiences (Malmendier and Nagel, 2011, 2016). Additionally, 57 responses refer to other historical events, such as the 2001 downturn or more recent years.

Shock type What kind of shock do firms consider when asked about a 5% drop in GDP? The majority of responses identify different types of demand deficiencies, including general low demand, sectoral spillovers, and precautionary saving motives among customers. Nonetheless, there are also 67 responses that mention geopolitical escalations, typically a war between Taiwan and China, as an event that leads to a significant drop in GDP. Other supply-side factors are mentioned more indirectly through references to higher prices. Note that many of the demand channels mentioned could be the result of negative supply shocks, as discussed in Guerrieri et al. (2022).

Firm reactions to a macro tail event realization Many managers share in the open-ended text box their reactions to a macro tail event or generally how other firms might respond (the “(customers)” version). The most common response is reducing labor input, including the use of short-time work, which, however, seems insufficient in such a

crisis. Additionally, managers suggest reducing investment, general business shrinkage, and overall cost reductions as reactions to a macro tail event. Firms also mention potential closures, and some discuss the possibility of switching production activities abroad.

Further responses There are also numerous responses (244) where managers indicate that they will be affected but do not provide further details. Many respondents also argue from an industry-wide perspective, consistent with the absence of firm-level exposure information. Some firms note that they are partially shielded from a significant GDP decline due to a diversified export structure. Others indicate that the impact on their business critically depends on government spending ('fiscal dependence'). Lastly, some firms explicitly refer to two different exposure scenarios, often describing a close-to-zero exposure scenario and a very strong exposure scenario.

Table D.10 – Codebook

Category	Explanation	Examples
Assessability		
Not assessable	Mentioning that it is impossible or almost impossible to assess exposure	“not possible to estimate” “An estimate is hardly possible. The past 2-3 years have shown that regional events can have a major impact on global contexts. [...]” “Cannot be estimated, as too many factors in this scenario are unpredictable.”
Difficult to assess	Mentioning that it is difficult to assess	“Very difficult to assess ” “It is difficult to estimate, but as we manufacture capital goods in the broadest sense, the effect is likely to be disproportionately high.”
Assessable	Mentioning that it is possible to assess exposure	“The impact of a possible macroeconomic crisis in Germany on our real sales can be estimated to a certain extent [...]”
Easy to assess	Mentioning that it is easy to assess exposure	“The effects are easy to assess because we have precisely defined our customer base and can therefore estimate their behavior from past experience.”
Expected sales changes (percent changes in numbers)		
One number	Change in sales in one number	“Sales down 25%”

One number and expressing uncertainty	Change in sales in one number, but expressing uncertainty	“Sales will slump by around 20%.” “Turnover is likely to halve.”
Range	Reporting expected sales change between two numbers	“Sales decline of between 25 and 40% possible.”
Min exposure scenario number	Minimum sales change	“Impact 5% GDP on own sales > 20%.”
Max exposure scenario number	Maximum sales change	“Difficult to estimate, sales losses of up to 30% conceivable”

Reference to past exposure

Great Recession reference	Reference to the Great Recession / Financial Crisis (2008/2009)	“We already had this situation in the 2009 financial crisis. That would mean a 30% slump in sales.”
COVID-19 reference	Reference to the COVID-19 pandemic	“Not assessable as it depends very much on which sectors are affected. Example: Corona crisis and booming construction industry.” “Very difficult to estimate: Both the corona crisis (sharp drop in GDP) and the mild recession in 2023 (merely the absence of growth) have each led to double-digit percentage declines in sales. [...]”
Other historic reference	Reference to other experiences	“Past crises of this magnitude have shown that the automotive industry is also affected, but recovers relatively quickly. We anticipate a 15-25% loss of sales in the relevant annual period”

Exposure strength (in words)

Positive impact	Positive impact on business	“In our sector, crises tend to boost our sales because certain consulting needs become necessary in order to reduce costs, optimize processes and increase sales. In principle, we always benefit.”
(Almost) no impact	No or almost no impact on own business	“We have a broad base, so there are always sectors that work. We have not really been affected by slumps so far.”

Disproportionately low impact	Minor negative impact on own business	“It is very difficult for us to assess the impact of such a crisis. As we operate in relatively resilient markets, as past crises have shown, we expect only minor, short-term declines in sales in such crises.”
Proportional impact	1:1 impact of GDP on own business	“We will lose sales in proportion to the decline in GDP.”
Disproportionately high impact	Particularly strong negative impact on own business	“Would have a catastrophic impact” “Significant drop in sales”

Shock type

Demand shock (sector spillover)	Mentioning low demand from other sectors	“If GDP falls, this means that industry and commerce will suffer a sharp drop in sales and reduce their investments. This would have a direct impact on the success of our company.”
Demand shock (precautionary customers)	Mentioning low demand because of low purchasing power / uncertainty	“A decline of 5% would be a dramatically large decline. The psychological impact on the economy would be more serious than the actual economic consequences. For the demand for my services, this could mean a drop in demand of up to 90%.”
Demand shock other	Other demand factors, mentioning lower prices	“Market demand for cosmetics, which are otherwise quite crisis-proof, will then also decline significantly, with a drop in orders for us.”
Supply shock (war / geopolitics)	Mentioning wars or international politics	“This is very difficult for us to assess. We assume that Ukraine will lose the war, climate catastrophes will become more frequent and China will probably dare to invade Taiwan. This will lead to a global economic crisis that will also hit our economy hard.” “An escalation of the China Taiwan conflict would have a significant impact [...]”
Supply shock other	Other supply factors, mentioning higher prices	“Competition would increase dramatically, and falling sales and rising unit costs would make things difficult.”

Reaction to macro tail event

Reduce labor	Decline in number of employees, fire employees	“Slump in sales and layoffs of employees. [...]”
Reduce labor (customers / general)	Decline in employees of customers, general high unemployment	“The probable job losses will lead to a reduction in purchasing power and stand in the way of domestic demand. This will block purchases by citizens. [...]”
Short-time work	Use of short-time work	“Our order situation will deteriorate and we will consider short-time working”
Reduce investment	Reduction in investment	“[...] stop of investments.”
Reduce investment (customers / general)	Reduction in investment of customers, general lower investment	“The impact is very high. Less is being invested and that affects us directly.”
Shrink	Overall decline in activities	“Adjustment of capacities and personnel requirements”
Reduce costs	cutting costs	“Our sales would fall by double digits, but much less than 20%. In such a scenario, we would not be planning any layoffs at present, but would be aiming for significant cost reductions.”
Potential firm closure	Potential closure of firm / bankruptcy	“If the conditions get too bad, the company closes down and I become a privateer and sell spare parts on my own.” “endangers the company”
Potential firm closure (customers / general)	Potential firm closures of customers, general more exits	“[...] as some customers and suppliers would certainly not survive such a crisis.”
Go abroad	Move the business to other countries	“Crisis mode with job cuts and relocation of production abroad”

Other responses

Exposed	Only stating that the firm would be affected	“Certainly an impact on our sales! ”

Industry	Relating to the exposure of the entire industry	“This would have a full impact because we are active in the leisure sector, where savings are always made first.”
Diversified	Diversified business	“Such a scenario cannot be predicted, but as a lesson from corona, we know how your own business model can be completely turned upside down within 7 days without even having made a mistake. Our lesson: diversification and innovation.”
High export	High export share	“The impact of a possible macroeconomic crisis in Germany on our real sales can be estimated to a certain extent, as we can control our foreign sales as a globally operating company”
Fiscal dependence	Dependent on fiscal response in the crisis	“We are heavily dependent on public sector investment. Macroeconomic crises that could lead to a decline in public investment would also have an impact on our sales.”
Two different scenarios	Discussion of two different explicit scenarios, range in words	“As a manufacturer of consumer goods, we are naturally affected by a macroeconomic crisis. On the other hand, there are strong market adjustments in our industry in such phases, which can even lead to an improvement in real sales despite declining macroeconomic figures.”
Politics	Frustration about politics	“If we had a government instead of the current missionaries, such a question would not arise. Honestly, Scholz and Co. are really weak and very clumsy.”
General pessimism	Overall pessimism about low growth	“This macroeconomic crisis is already here. Turnover is falling. Entire industries are coming to a standstill or moving away. Sooner or later, we will probably have to relocate the rest abroad or close down in order to limit the damage.”
Other	Other comments	“Demographic change will have a significant impact on the future.”

Notes: This table provides the codebook which is used to classify the answers of firms to the open-ended question “*To what extent can you generally forecast the effects of such a potential macroeconomic crisis (GDP collapses by around 5% within a year) on your real sales? Please let us know your thoughts on this macroeconomic scenario here:*”. The codebook provides an explanation of the categories, and example extracts from open-text responses.

Table D.9 – Classifications

Categories	Number of responses
Accessibility	
Not assessable	189
Difficult to assess	292
Assessable	17
Easy to assess	22
Expected sales changes (percent changes in numbers)	
One number	92
One number and expressing uncertainty	53
Range	111
Min exposure scenario number	26
Max exposure scenario number	34
Reference to past exposure	
Great Recession reference	48
COVID-19 reference	90
Other historic reference	56
Exposure strength (in words)	
Positive impact	44
(Almost) no impact	84
Disproportionately low	160
Proportional impact	54
Disproportionately high impact	272
Shock type	
Demand shock (sector spillover)	103
Demand shock (precautionary customers)	145
Demand shock other	158
Supply shock (war / geopolitics)	68
Supply shock other	42
Reaction to macro tail event	
Reduce labor	81
Reduce labor (customers)	15
Short-time work	28
Reduce investment	27
Reduce investment (customers)	67
Shrink	32
Reduce costs	20
Potential firm closure	95
Potential firm closure (customers)	19
Go abroad	17
Other responses	
Exposed	244
Industry	116
Diversified	11
High export	73
Fiscal dependence	44
Two different scenarios	99

Notes: This table presents the classification results based on the answers of firms to the open-ended question “*To what extent can you generally forecast the effects of such a potential macroeconomic crisis (GDP collapses by around 5% within a year) on your real sales? Please let us know your thoughts on this macroeconomic scenario here:*”. The classification is done by a research assistant following the Codebook, see Table D.10.

D.6 MODEL APPENDIX

Shareholder value maximization (risk aversion) This alternative model differs from the model in Section 4.3.1 in two ways. First, I introduce risk aversion as firms maximize shareholder value. They discount future cash flows using state prices, and those take into account investors' risk aversion in the form of risk neutral probabilities. Hence, there is risk aversion about aggregate developments. Second, the world is entirely Bayesian. This means that firms can assign probabilities to the respective firm-specific exposure scenarios in the low productivity state. I introduce risk aversion by using risk-neutral transition probabilities q . These can be derived by scaling the actual probability of the low productivity state by a stochastic discount factor m : $q_{h \rightarrow l} = m * p_{h \rightarrow l}$

In the following, I test whether this model also predicts a central role of the worst-case exposure for firms investment decisions. For simplicity, a firm has either productivity $A^l - s$ or $A^l + s$ in the case of a macro tail event, where $2 \times s$ defines the range of firm-specific exposure. The expected productivity depends on the risk-neutral probability of transitioning to the macro tail risk event q_l and the conditional probability $p_{w|l}$ of realizing the lowest productivity $A^l - s$ in this case:

$$\mathbb{E}[A_{t+1}|A^h] = (1 - q_l)A^h + q_l p_{w|l}(A^l - s) + q_l(1 - p_{w|l})(A^l + s) \quad (D.1)$$

$$\frac{\partial}{\partial q_l} = -A^h + p_{w|l}\beta_w(A^l - s) + (1 - p_{w|l})(A^l + s) \quad (D.2)$$

$$\frac{\partial}{\partial s \partial q_l} = 1 - 2p_{w|l} \quad (D.3)$$

These equations show that as long as the distribution of exposure is symmetric ($p_{w|l} = 1/2$), changes in the severity of the worst-case scenario do not affect the optimal level of capital. This symmetry in exposure implies that for every firm facing its worst-case scenario, there is another experiencing its best-case scenario. Consequently, these idiosyncratic risks cancel each other out across firms, effectively diversifying away the impact on expected aggregate productivity. Note that this is not the case for ambiguity, where the law of large numbers does not hold. Intuitively, the investor faces independent sources of uncertainty, represented by a set of exposure scenarios, and therefore acts as if the mean productivity of each firm is lower. For further examples and a more detailed discussion, see Ilut and Saijo (2021) and Epstein and Schneider (2008). A formal analysis of the law of large numbers for i.i.d. ambiguous random variables can be found in Epstein and Schneider (2003a).

I now examine the quantitative relevance of this alternative model with risk aversion relative to the baseline model with ambiguity aversion for firms' investment decisions beyond the first moment. To do so, I compute q_l for concrete examples. Investors' risk aversion can be thought of as a representative household that owns all firms and consumes firms' profits. Households maximize lifetime utility. For simplicity, I assume CRRA preferences $u(c) = \frac{c^{1-\eta}}{1-\eta}$. I back out the risk-neutral probability q_l , such that the utility of expected aggregate output based on q_l equals the utility of the certainty equivalent. For $\eta = 1$ (log utility), we have

$$\ln((1 - q_l)h + q_l l) = (1 - p_l) \ln(h) + p_l \ln(l)$$

$$q_l = \frac{h^{(1-p_l)}(h^{p_l} - l^{p_l})}{h - l}$$

and for $\eta \neq 1$, we have


$$\frac{((1 - q_l)h + q_l l)^{1-\eta}}{1 - \eta} = (1 - p_l) \frac{h^{1-\eta}}{1 - \eta} + p_l \ln(l)$$

$$q = \frac{h - (-p_l h^{1-\eta} + h^{1-t} + p_l l^{1-\eta})^{1/(1-\eta)}}{h - l}.$$

h (l) refers to output units in the high (low) productivity state. Following the parameterization in Section 4.3.1, I set $A^h = 1$, $A^l = 0.9$, $\beta_{avg} = 1$, $\beta_{worst} = 2.5$, $\beta_{best} = 0.5$ and $p = 0.1$. Since aggregate productivity scales output linearly and preferences are CRRA, we simply set h and l so that the ratio $l/h = 0.9$. I consider different degrees of risk aversion ($\eta = 1, \eta = 2, \eta = 5, \eta = 10$). For instance, the implied risk-neutral probability with log utility is $q \approx 0.1048$.

How risk averse does the investor need to be to achieve a similar decline in the optimal level of capital as in the model with ambiguity aversion? In Figure D.15, I show the deviation from the risk and ambiguity neutral capital stock. The optimal level of capital hardly changes (less than -0.1%) when I introduce shareholder maximization with a risk aversion parameter of $\eta = 1$ for investors. In contrast, ambiguity aversion about the firm-specific exposure to the macro tail risk event significantly reduces the optimal capital stock by 2.3%. A risk aversion parameter of $\eta = 2$ produces only a tenth of the effect compared to ambiguity aversion. Even with $\eta = 5$ or $\eta = 10$, the effect sizes are an order of magnitude smaller. To generate an effect size as large as with ambiguity aversion, we would have to assume extremely high risk aversion, with η close to 20.

Figure D.15 – Effect of different degrees of risk aversion on optimal capital stock



Notes: The figure shows the deviation from the risk and ambiguity neutral capital stock in % for different degrees of risk aversion (in blue) and the baseline model with ambiguity aversion about firm-specific exposure to macro tail risk events (in orange). The filled blue bars show results for the model with shareholder value maximization taking into account investors' risk aversion. The outlined bars show the results for an entrepreneur-type model. Both models are explained in the text. The parameterization follows Section 4.3.1: $A^h = 1$, $A^l = 0.9$, $\beta_{avg} = 1$, $\beta_{worst} = 2.5$, $\beta_{best} = 0.5$ and $p = 0.1$.

Another extension of the model is an entrepreneur-type model in which an entrepreneurs' only source of income for consumption is her own firm's output. In this case, there is additional risk aversion about firm-specific output in the macro tail risk event. To analyze this possibility, I additionally introduce risk neutral probabilities for the conditional probabilities of realizing a high/low exposure to the macro tail risk event. I set the actual probabilities so that the average exposure scenario equals the expectation of worst-case and best-case scenarios.¹

The outlined bars in Figure D.15 show the additional effect on the optimal capital stock in the entrepreneur-type model. It is still the case that reasonable values of risk aversion do not generate strong negative effects on investment beyond the first moment. Only with η close to 10, an effect size similar to ambiguity aversion can be achieved. While this is

¹This entrepreneurial model with idiosyncratic risk aversion of the owner would predict a negative effect on investment beyond the first moment when the subjective probability of macro tail risk increases, driven by the firm's worst-case exposure. This is because the conditional risk-neutral probability of realizing the worst case $q_{w|l}$ is greater than $p_{w|l}$. However, I can rule out this possibility because I show empirically that the effects are not stronger for risk-averse managers, see Section 4.3.


a very stylized model, it demonstrates that introducing risk aversion in a Bayesian world instead of ambiguous exposure to macro tail risk events would have only small negative effects on investment beyond the first moment, given reasonable levels of risk aversion.

Bankruptcy risk This alternative model introduces bankruptcy risk to the firms' choice problem which increases exponentially with a decline in productivity. Upon bankruptcy, a firm incurs additional costs equating to its capital stock k . Hence, the following term is added in the optimization problem

$$\text{Potential bankruptcy costs:} \quad \frac{k_{t+1}}{a \times \mathbb{E}[A_{t+1}]^b}, \quad (\text{D.4})$$

where $a > 0$ and $b > 1$ translate the expected productivity in the next period into the probability of bankruptcy. In Panel (A) of Figure D.16, the probability of bankruptcy is plotted against the subjective probability of macro tail risk, where $a = 100$ and $b = 30$. The corresponding optimal capital stock is shown in Panel (B). Since the probability of bankruptcy increases exponentially with lower productivity, the optimal capital stock also falls in a concave manner with an increase in the subjective probability of macro tail risk. To obtain an economically significant effect of bankruptcy risk on the optimal capital stock that increases with macro tail risk, a relatively high probability of bankruptcy risk in a macro tail risk event is necessary, see Panel (A).

Figure D.16 – Effect of bankruptcy risk in the stylized model



Notes: The figure shows the probability of bankruptcy (Panel (A)) and the effect of bankruptcy costs on the choice of optimal capital (Panel (B)) in relation to the subjective probability of a macro tail event.

D.7 ADDITIONAL QUANTITATIVE MODEL INFORMATION

Numerical implementation I solve the problem of the firm using value function iteration over a discrete grid of points. I specify a 350-point capital grid with an upper bound that is large enough to be non-binding. The capital grid is constructed recursively to ensure that firms can always adjust their capital at no cost if they set investment equal to zero: $k_i = k_{i-1}/(1 - \delta)$. Specifically, I generate two 175-point capital grids, where one is shifted by $\delta/2$ to create a more narrow grid. The aggregate productivity process is already discrete in the model and following a three-state Markov process. I discretize the autoregressive idiosyncratic productivity process following the standard method of Tauchen (1986) into a 13-state Markov process. I consider five firm-types with different worst-case exposures to the macro tail risk event. The model is solved separately for each type of firm.

Calibration details I parameterize the model based on standard parameters from the literature and informed by my empirical analysis at a quarterly level. In the baseline calibration, I set the annual real interest rate r to 2%. The decreasing return to capital is parameterized with $\alpha = 0.33$, with capital depreciating at 12% annually. Following other papers in the literature, I assume a high persistence of log idiosyncratic productivity shocks by setting $\rho = 0.95$. The standard deviation of the idiosyncratic innovations is 0.05 and thereby between the volatility of the high and low idiosyncratic uncertainty state in Bloom et al. (2018).

The corporate income tax rate is set to $\tau = 0.3$, which lies between the average statutory corporate income tax rate (slightly above) and the average effective corporate income tax rate in Germany (slightly below); see Fuest and Neumeier (2023). I calibrate the fixed production costs and linear production costs so that the majority of firms are profitable in normal times, but most will incur losses in a macro tail event. Panel (B) in Figure D.17 shows that the overall investment response is very similar, regardless of how these parameters are calibrated. These additional parameters are only important for the fiscal policy counterfactuals that hinge on whether a firm is profitable.

The model features firm-specific exposure to macro tail risk events. In a world with ambiguous exposure, firms act based on the expected worst-case scenario. I calibrate the parameterization of the five firm types using the elicited worst-case exposures from survey

Table D.11 – Parameter calibration in the model (quarterly level)


Technology		
r	0.005	Annual interest rate of 2%
α	0.33	Decreasing return to scale
δ	0.03	Depreciation of capital stock
ρ_z	0.95	Persistence of idiosyncratic productivity
σ_z	5%	Standard deviation of innovations to idiosyncratic productivity
Additional production costs		
γ_0	0.7	Fixed production costs
γ_1	0.12	Linear production costs
τ	0.3	Corporate income tax rate
Macro tail risk exposure		
$\beta_{worstp10}$	0.5	0.5× avg. exposure of median firm (survey)
$\beta_{worstp25}$	1.5	1.5× avg. exposure of median firm (survey)
$\beta_{worstp50}$	2.5	2.5× avg. exposure of median firm (survey)
$\beta_{worstp75}$	4	4× avg. exposure of median firm (survey)
$\beta_{worstp90}$	6	6× avg. exposure of median firm (survey)
Adjustment costs		
ϕ_1	0.1%	Fixed capital adjustment cost
ϕ_2	6	Quadratic capital adjustment cost
ϕ_3	97.5%	Resale capital adjustment cost
Transition probabilities of macro productivity states		
$\pi_{lowA highA,low\_risk}$	0.01	Probability of macro tail event in low macro risk state
$\pi_{lowA highA,high\_risk}$	0.04	Probability of macro tail event in high macro risk state
$\pi_{lowA lowA}$	0.8	Probability of staying in macro tail event

Notes: The table provides the parameters relating to technology, additional production costs, macro tail risk exposure, adjustment costs, and aggregate transition probabilities. The transition probability from *lowA* to *highA,high_risk* is 0.1 and to *highA,low_risk* 0.1. The transition probability to stay in the current state is 0.95 for *highA,high_risk* and *highA,low_risk*. The probability to transition from *highA,low_risk* to *highA,high_risk* is 0.04 and 0.01 vice versa.

data at the 10th, 25th, 50th, 75th, and 90th percentiles, see Table D.2. For instance, the median firm in the survey expects a 2.5 times greater drop in sales in the worst-case scenario compared to the average-case scenario.

The model includes a series of capital adjustment costs. For quadratic adjustment costs, I follow Salgado, Guvenen and Bloom (2023) and set $\phi_1 = 6$. To achieve more dynamics in the change of firms that plan to decrease investment, which is my main outcome variable in Section 4.3.2, I set fixed and resale capital adjustment costs lower than Salgado, Guvenen and Bloom (2023) and set $\phi_1 = 0.001$ and $\phi_3 = 0.975$. Note that to generate sizable

Figure D.17 – Macro tail risk news shocks: Dynamics of the aggregate capital stock – sensitivity analysis



Notes: The figures show the aggregate change in capital after a macro tail risk news shock. Note that only macro tail risk increases without a realization of the macro tail risk event. The left panel shows sensitivity analysis regarding the parameterization of the capital adjustment costs. The green dashed line is based on an alternative parameterization with higher fixed adjustment costs ($\phi_1 = 0.5\%$ instead of $\phi_1 = 0.1\%$). The blue dashed line is based on an alternative parameterization with higher quadratic adjustment costs ($\phi_2 = 8$ instead of $\phi_1 = 6$). The orange dashed line is based on an alternative parameterization with higher capital resale costs ($\phi_3 = 90\%$ instead of $\phi_3 = 97.5\%$). The right panel shows sensitivity analysis regarding the parameterization of the additional production costs and corporate tax rate. The green dashed line is based on an alternative parameterization with lower running costs ($\gamma_1 = 0.1$ instead of $\gamma_1 = 0.12$). The blue dashed line is based on an alternative parameterization with lower fixed costs ($\gamma_0 = 0.5$ instead of $\gamma_0 = 0.7$). The orange dashed line is based on an alternative parameterization without these additional productions and corporate taxes ($\gamma_0 = \gamma_1 = \tau = 0$).

negative investment effects of volatility shocks through the “wait-and-see” mechanism these non-convex costs need to be high. In the present model, the mechanism does not hinge on these costs. Nonetheless, we observe in the sensitivity analysis in Figure D.17 that higher fixed capital adjustment costs would further increase the negative effect of higher macro tail risk expectations on investment.

The macro productivity state in normal times is $A_h = 1$ and in the macro tail risk event $A_l = 0.925$, implying a 7.5% drop in productivity in the macro tail event for the median firm. I set the probability of remaining in the macro tail risk event to $\pi_{low A|low A} = 0.8$, implying an average length of a macro tail risk event of five quarters. This is equal to the length of the Great Recession 2008/2009 defined by the German Council of Economic Experts. The recession length of the COVID-19 recession was shorter, while the Great Depression and World War II were significantly longer.

D.8 BACK-OF-THE-ENVELOPE CALCULATION TO MAP THE EMPIRICAL ESTIMATES TO THE MODEL

- Baseline regression estimate β is 0.25 (Table D.3 Column 1)
- Baseline regression estimate β is 0.12 after controlling for first and second moment expectations (Table D.3 Columns 3-4).
- In the mode simulation, macro tail risk increases from 18% to 56% (five year horizon); a 38 percentage point (p.p.) increase.
- This increase in macro tail risk predicts an increase in the share of investment decreases by $0.25 \times 38 = 9.5$ p.p. and $0.12 \times 38 = 4.6$ p.p.
- Now, we map the indicator of investment decrease to a quantitative change in investment:
 - A planned decrease in investment in the survey predicts a decline in capital stock in balance sheet data by around 6% (see Section 2). Assuming an investment-to-capital (I/K) ratio of 1/12 implies a 72% decrease in investment.
 - A planned decrease in investment strongly predicts a realized decrease in investment in the survey. Using the planned decrease as in instrument for the realized decrease to estimate the effect on changes in quantitative investment (elicited in the survey) leads to a decline by 95%, see Table D.12.
- The implied overall effects are:
 - $9.5\% \times 72\% = 6.8\%$ decline in investment.
 - $9.5\% \times 95\% = 9\%$ decline in investment.
- The implied effects beyond first- and second-moments are:
 - $4.6\% \times 72\% = 3.3\%$ decline in investment.
 - $4.6\% \times 95\% = 4.4\%$ decline in investment.

Table D.12 – Prediction power of planned decrease in investment on quantitative changes in investment

	(1) First stage $\mathbb{1}\{\Delta Investment_{i,t} < 0\}$	(2) IV $\Delta \log(Investment_{i,t})$
$\mathbb{1}\{\Delta Planned Investment_{i,t+1} < 0\}$	0.311*** (0.0168)	
$\mathbb{1}\{\Delta Investment_{i,t} < 0\}$		-0.954*** (0.105)
Constant	0.203*** (0.00988)	0.337*** (0.0326)
F-Stat	343.4	
Time FE	Yes	Yes

Notes: The table presents the estimation results on the prediction power of a qualitative decrease in investment on quantitative changes in investment. The first column shows the first stage: a planned decrease in investment is highly predictive for a realized decrease in investment. The second column shows the effect of the instrumented decrease in investment on the change in log investment: . * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

D.9 ADDITIONAL LLM INFORMATION

GPT Prompt: You are a manager in a company who has to make macroeconomic forecasts for Germany based on information from newspaper articles.

Here is a newspaper article: [article]

Three forecasts are required, given the information from the article.

1. Assessment of strong recession:

- **Assessment:** Does the likelihood of a strong recession (Gross Domestic Product declines of at least 5%) in Germany over the next 12 months increase or decrease?
- **Confidence:** How confident are you in this assessment? (0 = not confident, 1 = confident)
- **Explanation:** General explanation for this assessment, maximum 15 words.

The format for the response is as follows:

- {increases, decreases, neutral}
- {0,1}
- {Words: maximum 15}

2. Assessment of strong upturn:

- Assessment: Does the likelihood of a strong upturn (Gross Domestic Product increases by at least 3%) in Germany over the next 12 months increase or decrease?
- Confidence: How confident are you in this assessment? (0 = not confident, 1 = confident)
- Explanation: General explanation for this assessment, maximum 15 words.

The format for the response is as follows:

- {increases, decreases, neutral}
- {0,1}
- {Words: maximum 15}

3. Assessment on average:

- Assessment: In the most likely scenario, how will the economic development in Germany evolve in the next 12 months?
- Confidence: How confident are you in this assessment? (0 = not confident, 1 = confident)
- Explanation: General explanation for this assessment, maximum 15 words.

The format for the response is as follows:

- {negative, positive, neutral}
- {0,1}
- {Words: maximum 15}

Please write the responses for the strong recession, the strong upturn, and the average economic development directly underneath each other, without additional text.

GPT Prompt (German):

Stell Dir vor Du bist ein Manager in einem Unternehmen, der herausfinden muss, wie die Informationen aus Nachrichten sich auf die deutsche Wirtschaft auswirken.

Hier is ein Zeitungsartikel: [article]

Es sind drei Prognosen gefragt, gegeben der Informationen aus dem Artikel.

1. Einschätzung starker Rezession:

- Einschätzung: Erhöht oder verringert sich die Wahrscheinlichkeit einer starken Rezession (Bruttoinlandsprodukt fällt um mindestens 5%) in Deutschland in den nächsten 12 Monaten?
- Konfidenz: Wie sicher bist Du Dir bei dieser Einschätzung? (0 = nicht sicher, 1 = sicher)
- Erklärung: Generelle Erklärung zu dieser Einschätzung, maximal 15 Wörter.

Das Format für die Antwort ist folgendes:

- {erhöht, verringert, neutral}
- {0,1}
- {Wörter: maximal 15}

2. Einschätzung starker Aufschwung:

- Einschätzung: Erhöht oder verringert sich die Wahrscheinlichkeit eines starken Aufschwungs (Bruttoinlandsprodukt steigt um mindestens 3%) in Deutschland in den nächsten 12 Monaten?
- Konfidenz: Wie sicher bist Du Dir bei dieser Einschätzung? (0 = nicht sicher, 1 = sicher)
- Erklärung: Generelle Erklärung zu dieser Einschätzung, maximal 15 Wörter.

Das Format für die Antwort ist folgendes:

- {erhöht, verringert, neutral}
- {0,1}
- {Wörter: maximal 15}

3. Einschätzung im Durchschnitt:

- Einschätzung: Wie wird sich im wahrscheinlichsten Szenario die wirtschaftliche Entwicklung in Deutschland in den nächsten 12 Monaten entwickeln?
- Konfidenz: Wie sicher bist Du Dir bei dieser Einschätzung? (0 = nicht sicher, 1 = sicher)
- Erklärung: Generelle Erklärung zu dieser Einschätzung, maximal 15 Wörter.

Das Format für die Antwort ist folgendes:

- {negativ, positiv, neutral}
- {0,1}
- {Wörter: maximal 15}

Bitte schreibe die Antworten für die starke Rezession, den starken Aufschwung und die durchschnittliche wirtschaftliche Entwicklung direkt untereinander, ohne zusätzlichen Text.

Article example: *Headline: Putin rejects Bush's plans (2001/18/06)*

Despite differences on key security issues, US President George W. Bush and Russian President Vladimir Putin demonstrated unity at their first meeting in Slovenia. Russia and America are "not enemies", Putin emphasized at Brdo Castle near the Slovenian capital Ljubljana on Saturday. The differences weighed "far less than what unites us", he added after the two-and-a-half hour meeting.[...] However, the two heads of state remained at odds over the concept for a new global security system.[...] "We have found a good basis on which to build our relations," Putin emphasized. The existing differences of opinion were not of a "fundamental nature". Another controversial topic at the summit was the eastward expansion of NATO, which Moscow opposes. Handelsblatt (2001)

GPT classification:

- macro tail risk: decreases
- average macro expectations: neutral
- upside risk expectations: neutral

GPT explanation for decrease in macro tail risk:

"Unity between the USA and Russia could help to stabilize the global security situation."

Bibliography

- Abdellaoui, Mohammed, Peter Klibanoff, and Lætitia Placido.** 2015. “Experiments on Compound Risk in Relation to Simple Risk and to Ambiguity.” *Management Science*, 61(6): 1306–1322.
- Afrouzi, Hassan.** 2023. “Strategic Inattention, Inflation Dynamics and the Non-Neutrality of Money.” NBER Working Paper 31796.
- Afrouzi, Hassan, Spencer Yongwook Kwon, Augustin Landier, Yueran Ma, and David Thesmar.** 2023. “Overreaction in Expectations: Evidence and Theory.” *Quarterly Journal of Economics*, 138(3): 1713–1764.
- Aher, Gati V., Rosa I. Arriaga, and Adam Tauman Kalai.** 2023. “Using Large Language Models to Simulate Multiple Humans and Replicate Human Subject Studies.” PMLR International Conference on Machine Learning.
- Ahlander, Edvin, Mikael Carlsson, and Mathias Klein.** 2023. “Price Pass-Through Along the Supply Chain: Evidence from PPI and CPI Microdata.” Working Paper.
- Alfaro, Iván, Nick Bloom, and Xiaoji Lin.** 2024. “The Finance Uncertainty Multiplier.” *Journal of Political Economy*, 132(2): 577–615.
- Almeida, Heitor, Murillo Campello, and Michael S. Weisbach.** 2004. “The Cash Flow Sensitivity of Cash.” *Journal of Finance*, 59(4): 1777–1804.
- Altig, David, Jose Maria Barrero, Nicholas Bloom, Steven J Davis, Brent Meyer, and Nicholas Parker.** 2022. “Surveying Business Uncertainty.” *Journal of Econometrics*, 231(1): 282–303.
- Ameriks, John, Joseph Briggs, Andrew Caplin, Matthew D. Shapiro, and Christopher Tonetti.** 2020. “Long-Term-Care Utility and Late-in-Life Saving.” *Journal of Political Economy*, 128(6): 2375–2451.

- Amiti, Mary, Oleg Itskhoki, and Jozef Konings.** 2019. “International Shocks, Variable Markups, and Domestic Prices.” *Review of Economic Studies*, 86(6): 2356–2402.
- Andrade, Philippe, Olivier Coibion, Erwan Gautier, and Yuriy Gorodnichenko.** 2022. “No Firm Is an Island? How Industry Conditions Shape Firms’ Expectations.” *Journal of Monetary Economics*, 125: 40–56.
- Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart.** 2022. “Subjective Models of the Macroeconomy: Evidence From Experts and Representative Samples.” *Review of Economic Studies*, 89(6): 2958–2991.
- Andre, Peter, Ingar Haaland, Chris Roth, Mirko Wiederholt, and Johannes Wohlfart.** 2024. “Narratives About the Macroeconomy.” SAFE Working Paper 426.
- Ang, Andrew, Joseph Chen, and Yuhang Xing.** 2006. “Downside Risk.” *Review of Financial Studies*, 19(4): 1191–1239.
- Angeletos, George-Marios, and Zhen Huo.** 2021. “Myopia and Anchoring.” *American Economic Review*, 111(4): 1166–1200.
- Angeletos, George-Marios, Zhen Huo, and Karthik A. Sastry.** 2021. “Imperfect Macroeconomic Expectations: Evidence and Theory.” *NBER Macroeconomics Annual*, 35: 1–86.
- Arndt, Sarah, and Zeno Enders.** 2023. “Shock Transmissions in Different Inflation Regimes.” Working Paper.
- Aruoba, S Boragan, Eugene Oue, Felipe Saffie, and Jonathan L Willis.** 2022. “Real Rigidities, Firm Dynamics and Monetary Non-Neutrality: Role of Demand Shocks.” Working Paper.
- Auclert, Adrien, Rodolfo D Rigato, Matthew Rognlie, and Ludwig Straub.** 2024. “New Pricing Models, Same Old Phillips Curves?” *Quarterly Journal of Economics*, 139(1): 121–186.
- Auerbach, Alan J.** 2006. “Who Bears the Corporate Tax? A Review of What We Know.” *Tax Policy and the Economy*, 20: 1–40.
- Auerbach, Alan J., and Yuriy Gorodnichenko.** 2013. “Output Spillovers from Fiscal Policy.” *American Economic Review*, 103(3): 141–46.

- Auer, Raphael A, and Raphael S Schoenle.** 2016. “Market Structure and Exchange Rate Pass-Through.” *Journal of International Economics*, 98: 60–77.
- Azeredo da Silveira, Rava, and Michael Woodford.** 2019. “Noisy Memory and Over-Reaction to News.” *AEA Papers and Proceedings*, 109: 557–61.
- Bachmann, Rüdiger, and Christian Bayer.** 2014. “Investment Dispersion and the Business Cycle.” *American Economic Review*, 104(4): 1392–1416.
- Bachmann, Rüdiger, and Peter Zorn.** 2020. “What Drives Aggregate Investment? Evidence from German Survey Data.” *Journal of Economic Dynamics and Control*, 115: 103873.
- Bachmann, Rüdiger, and Steffen Elstner.** 2015. “Firm Optimism and Pessimism.” *European Economic Review*, 79: 297–325.
- 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, David Baqaee, Christian Bayer, Moritz Kuhn, Andreas Löschel, Benjamin Moll, Andreas Peichl, Karen Pittel, and Moritz Schularick.** 2022a. “What if? The Economic Effects for Germany of a Stop of Energy Imports from Russia.” ECONtribute Policy Brief No. 028.
- Bachmann, Rüdiger, Kai Carstensen, Manuel Menkhoff, and Martin Schneider.** 2022b. “Umsatzunsicherheit und Umsatzerwartungen deutscher Firmen zur Zeit des Ukraine-Kriegs: Gas und Gasausfall spielen eine eher geringe Rolle.” *ifo Schnelldienst*, 75(06): 33–36.
- Bachmann, Rüdiger, Kai Carstensen, Manuel Menkhoff, and Martin Schneider.** 2024a. “Firm Expectations and Uncertainty in Normal Times and Times of Crisis.” Mimeo.
- Bachmann, Rüdiger, Kai Carstensen, Stefan Lautenbacher, and Martin Schneider.** 2020. “Uncertainty Is More Than Risk - Survey Evidence on Knightian and Bayesian Firms.” Stanford University Mimeo.

- Bachmann, Rüdiger, Kai Carstensen, Stefan Lautenbacher, Manuel Menkhoff, and Martin Schneider.** 2024b. “Uncertainty and Change: Survey Evidence of Firms’ Subjective Beliefs.” Working Paper.
- Bachmann, Rüdiger, Steffen Elstner, and Atanas Hristov.** 2017. “Surprise, Surprise – Measuring Firm-Level Investment Innovations.” *Journal of Economic Dynamics and Control*, 83: 107–148.
- Bachmann, Rüdiger, Steffen Elstner, and Eric R Sims.** 2013. “Uncertainty and Economic Activity: Evidence from Business Survey Data.” *American Economic Journal: Macroeconomics*, 5(2): 217–49.
- Backus, David, Mikhail Chernov, and Ian Martin.** 2011. “Disasters Implied by Equity Index Options.” *Journal of Finance*, 66(6): 1969–2012.
- Ba, Cuimin, J. Aislinn Bohren, and Alex Imas.** 2023. “Over- and Underreaction to Information.” Working Paper.
- Baker, Scott R., Nicholas Bloom, and Steven J. Davis.** 2016. “Measuring Economic Policy Uncertainty.” *Quarterly Journal of Economics*, 131(4): 1593–1636.
- Balleer, Almut, Sebastian Link, Manuel Menkhoff, and Peter Zorn.** 2024. “Demand or Supply? Price Adjustment Heterogeneity during the COVID-19 Pandemic.” *International Journal of Central Banking*, 20(1): 93–158.
- Ball, Laurence, and David Romer.** 1990. “Real Rigidities and the Non-Neutrality of Money.” *Review of Economic Studies*, 57(2): 183–203.
- Ball, Laurence, and Sandeep Mazumder.** 2011. “Inflation Dynamics and the Great Recession.” *Brookings Papers on Economic Activity*, 42(1): 337–405.
- Banca d’Italia.** 2022. “Survey on Inflation and Growth Expectations.” Methods and Sources: Methodological Notes.
- Barattieri, Alessandro, Susanto Basu, and Peter Gottschalk.** 2014. “Some Evidence on the Importance of Sticky Wages.” *American Economic Journal: Macroeconomics*, 6(1): 70–101.
- Barrero, Jose Maria.** 2022. “The Micro and Macro of Managerial Beliefs.” *Journal of Financial Economics*, 143(2): 640–667.

- Barro, Robert J.** 1972. “A Theory of Monopolistic Price Adjustment.” *Review of Economic Studies*, 39(1): 17–26.
- Barro, Robert J.** 2006. “Rare Disasters and Asset Markets in the Twentieth Century.” *Quarterly Journal of Economics*, 121(3): 823–866.
- Barro, Robert J., and José F. Ursúa.** 2012. “Rare Macroeconomic Disasters.” *Annual Review of Economics*, 4(1): 83–109.
- Basu, Susanto.** 1995. “Intermediate Goods and Business Cycles: Implications for Productivity and Welfare.” *American Economic Review*, 85(3): 512–531.
- Becker, Sascha 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.
- Beck, Günter W., and Sarah M. Lein.** 2020. “Price Elasticities and Demand-Side Real Rigidities in Micro Data and in Macro Models.” *Journal of Monetary Economics*, 115: 200–212.
- Bertrand, Marianne, and Antoinette Schoar.** 2006. “The Role of Family in Family Firms.” *Journal of Economic Perspectives*, 20(2): 73–96.
- Beutel, Johannes, and Valentin Stockerl.** 2023. “Tail Risk Neglect: A Reality Check.” Working Paper.
- Bianchi, Francesco, Cosmin Ilut, and Hikaru Saijo.** 2024a. “Diagnostic Business Cycles.” *Review of Economic Studies*, 91(1): 129–162.
- Bianchi, Francesco, Cosmin Ilut, and Hikaru Saijo.** 2024b. “Smooth Diagnostic Expectations.” NBER Working Paper 32152.
- Bianchi, Francesco, Cosmin Ilut, and Martin Schneider.** 2018. “Uncertainty Shocks, Asset Supply and Pricing over the Business Cycle.” *Review of Economic Studies*, 85(2): 810–854.
- Bianchi, Francesco, Sydney C. Ludvigson, and Sai Ma.** 2022. “Belief Distortions and Macroeconomic Fluctuations.” *American Economic Review*, 112(7): 2269–2315.

- Bils, Mark, Peter J. Klenow, and Benjamin A. Malin.** 2012. “Reset Price Inflation and the Impact of Monetary Policy Shocks.” *American Economic Review*, 102(6): 2798–2825.
- Blanco, Andres, Corina Boar, Callum Jones, and Virgiliu Midrigan.** 2022. “Non-linear Inflation Dynamics in Menu Cost Economies.” Working Paper.
- Blesse, Sebastian, Philipp Doerrenberg, and Anna Rauch.** 2019. “Higher Taxes on Less Elastic Goods? Evidence from German Municipalities.” *Regional Science and Urban Economics*, 75: 165–186.
- Blinder, Alan, Elie RD Canetti, David E Lebow, and Jeremy B Rudd.** 1998. *Asking About Prices: A New Approach to Understanding Price Stickiness*. Russell Sage Foundation.
- Bloom, Nicholas.** 2009. “The Impact of Uncertainty Shocks.” *Econometrica*, 77(3): 623–685.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry.** 2018. “Really Uncertain Business Cycles.” *Econometrica*, 86(3): 1031–1065.
- Bordalo, Pedro, John J Conlon, Nicola Gennaioli, Spencer Yongwook Kwon, and Andrei Shleifer.** 2023. “How People Use Statistics.” NBER Working Paper 31631.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer.** 2022. “Salience.” *Annual Review of Economics*, 14: 521–544.
- Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer.** 2019. “Diagnostic Expectations and Stock Returns.” *Journal of Finance*, 74(6): 2839–2874.
- Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer.** 2020. “Overreaction in Macroeconomic Expectations.” *American Economic Review*, 110(9): 2748–82.
- Born, Benjamin, Zeno Enders, and Gernot J. Müller.** 2024. “On Fire, News, and Expectations.” In *Handbook of Economic Expectations in Historical Perspective*, ed. Ingo Köhler, Laetitia Lenel, Alexander Nützenadel and Jochen Streb, forthcoming. Routledge.

- Born, Benjamin, Zeno Enders, Gernot J. Müller, and Knut Niemann.** 2023. “Firm Expectations About Production and Prices: Facts, Determinants, and Effects.” In *Handbook of Economic Expectations.*, ed. Rüdiger Bachmann, Giorgio Topa and Wilbert van der Klaauw, 355–383. Elsevier.
- Born, Benjamin, Zeno Enders, Manuel Menkhoff, Gernot J. Müller, and Knut Niemann.** 2024. “Firm Expectations and News: Micro v Macro.” CEPR Discussion Paper 17768.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess.** 2024. “Revisiting Event Study Designs: Robust and Efficient Estimation.” *Review of Economic Studies*, 91(6): 3253–3285.
- Bouchaud, Jean-Philippe, Philipp Krüger, Augustin Landier, and David Thesmar.** 2019. “Sticky Expectations and the Profitability Anomaly.” *Journal of Finance*, 74(2): 639–674.
- Bouwman, Christa H. S., and Ulrike Malmendier.** 2015. “Does a Bank’s History Affect Its Risk-Taking?” *American Economic Review: Papers & Proceedings*, 105(5): 321–325.
- Broer, Tobias, and Alexandre Kohlhas.** 2024. “Forecaster (mis-)behavior.” *Review of Economics and Statistics*, 106(5): 1334–1351.
- Brown, Tom, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, and Amanda Askell.** 2020. “Language Models Are Few-Shot Learners.” *Advances in Neural Information Processing Systems*, 33: 1877–1901.
- Buchheim, Lukas, Jonas Dovern, Carla Krolage, and Sebastian Link.** 2022. “Sentiment and Firm Behavior During the Covid-19 Pandemic.” *Journal of Economic Behavior & Organization*, 195: 186–198.
- Buchheim, Lukas, Sebastian Link, and Sascha Möhrle.** 2024. “Inflation and Wage Expectations of Firms and Employees.” IZA Discussion Paper 17269.
- Bunn, Philip, Lena S Anayi, Nicholas Bloom, Paul Mizen, Gregory Thwaites, and Ivan Yotzov.** 2022. “Firming up Price Inflation.” NBER Working Paper 30505.

- Busch, Christopher, David Domeij, Fatih Guvenen, and Rocio Madera.** 2022. “Skewed Idiosyncratic Income Risk over the Business Cycle: Sources and Insurance.” *American Economic Journal: Macroeconomics*, 14(2): 207–242.
- Bybee, Leland.** 2023. “The Ghost in the Machine: Generating Beliefs with Large Language Models.” Working Paper.
- Bybee, Leland, Bryan Kelly, Asaf Manela, and Dacheng Xiu.** 2024. “Business News and Business Cycles.” *Journal of Finance*, forthcoming.
- Caldara, Dario, and Matteo Iacoviello.** 2022. “Measuring Geopolitical Risk.” *American Economic Review*, 112(4): 1194–1225.
- Candia, Bernardo, Olivier Coibion, and Yuriy Gorodnichenko.** 2023. “The Macroeconomic Expectations of Firms.” In *Handbook of Economic Expectations*, ed. Rüdiger Bachmann, Giorgio Topa and Wilbert van der Klaauw, 321–353. Elsevier.
- Carroll, Christopher D.** 2003. “Macroeconomic Expectations of Households and Professional Forecasters.” *Quarterly Journal of Economics*, 118(1): 269–298.
- Carroll, Christopher D., Edmund Crawley, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White.** 2020. “Sticky Expectations and Consumption Dynamics.” *American Economic Journal: Macroeconomics*, 12(3): 40–76.
- Carstensen, Kai, Markus Heinrich, Magnus Reif, and Maik H. Wolters.** 2020. “Predicting Ordinary and Severe Recessions with a Three-State Markov-Switching Dynamic Factor Model: An Application to the German Business Cycle.” *International Journal of Forecasting*, 36(3): 829–850.
- Chahrour, Ryan, Kristoffer Nimark, and Stefan Pitschner.** 2021. “Sectoral Media Focus and Aggregate Fluctuations.” *American Economic Review*, 111(12): 3872–3922.
- Chen, Zhao, Xian Jiang, Zhikuo Liu, Juan Carlos Suárez Serrato, and Daniel Xu.** 2022. “Tax Policy and Lumpy Investment Behavior: Evidence from China’s Vat Reform.” *Review of Economic Studies*, 90(2): 634–674.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans.** 1999. “Monetary Policy Shocks: What Have We Learned and to What End?” In *Handbook of Macroeconomics*, ed. John B Taylor and Michael Woodford, 65–148. Elsevier.

- Cloyne, James, Ezgi Kurt, and Paolo Surico.** 2023. “Who Gains from Corporate Tax Cuts?” *NBER Working Paper*, 31278.
- Coibion, Olivier, and Yuriy Gorodnichenko.** 2015*a*. “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.** 2015*b*. “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.” *Quarterly Journal of Economics*, 135(1): 165–219.
- Colarieti, Roberto, Pierfrancesco Mei, and Stefanie Stantcheva.** 2024. “The How and Why of Household Reactions to Income Shocks.” NBER Working Paper 32191.
- Cooper, Russell W., and John C. Haltiwanger.** 2006. “On the Nature of Capital Adjustment Costs.” *Review of Economic Studies*, 73(3): 611–633.
- Council of Economic Advisers.** 2017. *Corporate Tax Reform and Wages: Theory and Evidence*. Washington, DC:Executive Office of the President of the United States.
- Curtis, E. Mark, Daniel G. Garrett, Eric C. Ohn, Kevin A. Roberts, and Juan Carlos Suárez Serrato.** 2021. “Capital Investment and Labor Demand.” *NBER Working Paper*, 29485.
- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam.** 1998. “Investor Psychology and Security Market Under- and Overreactions.” *Journal of Finance*, 53(6): 1839–85.
- de Bruin, Wändi, Keshav Dogra, Sebastian Heise, Edward S. Knotek II, Brent H. Meyer, Robert W. Rich, Raphael S. Schoenle, Topa, Giorgio, and Wilbert van der Klaauw.** 2023. “Estimates of Cost-Price Passthrough from Business Survey Data.” Working Paper.

- de Chaisemartin, Clément, and Xavier D'Haultfœuille.** 2022. "Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey." *NBER Working Paper*, 29691.
- Dedola, Luca, Mark Strøm Kristoffersen, and G Zullig.** 2021. "The Extensive and Intensive Margin of Price Adjustment to Cost Shocks: Evidence from Danish Multi-product Firms." Working Paper.
- Demirel, Ufuk Devrim.** 2021. "The Short-Term Effects of Tax Changes: The Role of State Dependence." *Journal of Monetary Economics*, 117: 918–934.
- Deutsche Bundesbank.** 2020. "Monthly Report June 2020." 72(6).
- Dibiasi, Andreas, Heiner Mikosch, and Samad Sarferaz.** 2021. "Uncertainty Shocks, Adjustment Costs and Firm Beliefs: Evidence From a Representative Survey." *KOF Working Papers*, 496.
- Dietrich, Alexander M, Gernot J. Müller, and Raphael Schoenle.** 2023. "Big News: Climate Change and the Business Cycle." Working Paper.
- Dimmock, Stephen G., Roy Kouwenberg, Olivia S. Mitchell, and Kim Peijnenburg.** 2016. "Ambiguity Aversion and Household Portfolio Choice Puzzles: Empirical Evidence." *Journal of Financial Economics*, 119(3): 559–577.
- Dobbins, Laura, and Martin Jacob.** 2016. "Do Corporate Tax Cuts Increase Investments?" *Accounting and Business Research*, 46(7): 731–759.
- Dolls, Mathias, Clemens Fuest, and Andreas Peichl.** 2012. "Automatic Stabilizers and Economic Crisis: Us Vs. Europe." *Journal of Public Economics*, 96(3): 279–294.
- Domar, Evsey D., and Richard A. Musgrave.** 1944. "Proportional Income Taxation and Risk-Taking." *Quarterly Journal of Economics*, 58(3): 388–422.
- Dotsey, Michael, and Robert G King.** 2005. "Implications of State-Dependent Pricing for Dynamic Macroeconomic Models." *Journal of Monetary Economics*, 52(1): 213–242.
- Dovern, Jonas, Lena Sophia Müller, and Klaus Wohlrabe.** 2023. "Local Information and Firm Expectations About Aggregates." *Journal of Monetary Economics*, 138: 1–13.

- Drechsel, Dirk, Heiner Mikosch, Samad Sarferaz, and Matthias Bannert.** 2022. “What Do Firm Managers Tell Us About the Transmission Channels of Oil Price Shocks?” KOF Working Papers.
- Dube, Arindrajit, Daniele Girardi, Oscar Jorda, and A Taylor.** 2023. “A Local Projections Approach to Difference-in-Differences Event Studies.” *NBER Working Paper*, 31184.
- Edgerton, Jesse.** 2010. “Investment Incentives and Corporate Tax Asymmetries.” *Journal of Public Economics*, 94(11): 936–952.
- Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo.** 2011. “Reference Prices, Costs, and Nominal Rigidities.” *American Economic Review*, 101(1): 234–262.
- Eichenbaum, Martin, Nir Jaimovich, Sergio Rebelo, and Josephine Smith.** 2014. “How Frequent Are Small Price Changes?” *American Economic Journal: Macroeconomics*, 6(2): 137–155.
- Elliott, Graham, Ivana Komunjer, and Allan Timmermann.** 2008. “Biases in Macroeconomic Forecasts: Irrationality or Asymmetric Loss?” *Journal of the European Economic Association*, 6(1): 122–157.
- Ellsberg, Daniel.** 1961. “Risk, Ambiguity, and the Savage Axioms.” *Quarterly Journal of Economics*, 75(4): 643–669.
- Enders, Zeno, Franziska Hünnekes, and Gernot J. Müller.** 2019. “Monetary Policy Announcements and Expectations: Evidence from German Firms.” *Journal of Monetary Economics*, 108: 45–63.
- Enders, Zeno, Franziska Hünnekes, and Gernot J. Müller.** 2022. “Firm Expectations and Economic Activity.” *Journal of the European Economic Association*, 20(6): 2396–2439.
- Epstein, Larry G., and Martin Schneider.** 2003a. “IID: Independently and Indistinguishably Distributed.” *Journal of Economic Theory*, 113(1): 32–50.
- Epstein, Larry G., and Martin Schneider.** 2003b. “Recursive Multiple-Priors.” *Journal of Economic Theory*, 113(1): 1–31.

- Epstein, Larry G., and Martin Schneider.** 2008. “Ambiguity, Information Quality, and Asset Pricing.” *Journal of Finance*, 63(1): 197–228.
- Fabiani, Silvia, Martine Druant, Ignacio Hernando, Claudia Kwapil, Bettina Landau, Claire Loupias, Fernando Martins, Thomas Y Mathä, Roberto Sabbatini, Harald Stahl, and Ad C J Stokman.** 2005. “The Pricing Behaviour of Firms in the Euro Area: New Survey Evidence.” ECB Working Paper Series 535.
- Farhi, Emmanuel, and Iván Werning.** 2019. “Monetary Policy, Bounded Rationality, and Incomplete Markets.” *American Economic Review*, 109(11): 887–392.
- Farmer, Leland, Emi Nakamura, and Jon Steinsson.** 2023. “Learning About the Long Run.” *Journal of Political Economy*, , (forthcoming).
- Fiore, Nicola Maria, Thorsten Martin, and Florian Nagler.** 2024. “Fiscal Constraints, Disaster Vulnerability, and Corporate Investment Decisions.” Working Paper.
- Fiori, Giuseppe, and Filippo Scoccianti.** 2023. “The Economic Effects of Firm-Level Uncertainty: Evidence Using Subjective Expectations.” *Journal of Monetary Economics*, 140: 92–105.
- Flynn, Joel P., and Karthik A. Sastry.** 2024. “Attention Cycles.” NBER Working Paper 32553.
- Focus Economics.** 2020. “Consensus Forecast Major Economies August 2020.” Consensus Forecast Major Economies.
- Foremny, Dirk, and Nadine Riedel.** 2014. “Business Taxes and the Electoral Cycle.” *Journal of Public Economics*, 115: 48–61.
- Friedman, Milton.** 1964. “Monetary Studies of the National Bureau.” *The National Bureau Enters Its 45th Year*, , (44th Annual Report): 7–25.
- Friedman, Milton.** 1993. “The “Plucking Model” of Business Fluctuations Revisited.” *Economic Inquiry*, 31(2): 171–177.
- Fuest, Clemens, and Florian Neumeier.** 2023. “Corporate Taxation.” *Annual Review of Economics*, 15: 425–450.

- Fuest, Clemens, Andreas Peichl, and Sebastian Siegloch.** 2018. “Do Higher Corporate Taxes Reduce Wages? Micro-Evidence from Germany.” *American Economic Review*, 108(2): 393–418.
- Furno, Francesco.** 2022. “The Macroeconomic Effects of Corporate Tax Reforms.” Working Paper.
- Fuster, Andreas, and Basit Zafar.** 2023. “Survey Experiments on Economic Expectations.” In *Handbook of Economic Expectations*, ed. Rüdiger Bachmann, Giorgio Topa and Wilbert van der Klaauw, 107–130. Elsevier.
- Fuster, Andreas, Greg Kaplan, and Basit Zafar.** 2021. “What Would You Do with \$500? Spending Responses to Gains, Losses, News, and Loans.” *Review of Economic Studies*, 88(4): 1760–1795.
- Gabaix, Xavier.** 2012. “Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance.” *Quarterly Journal of Economics*, 127(2): 645–700.
- Gabaix, Xavier.** 2020. “A Behavioral New Keynesian Model.” *American Economic Review*, (110): 2271–2327.
- Galí, Jordi.** 1999. “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?” *American Economic Review*, 89: 249–271.
- García-Schmidt, Mariana, and Michael Woodford.** 2019. “Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis.” *American Economic Review*, 109(1): 86–120.
- Garetto, Stefania.** 2016. “Firms’ Heterogeneity, Incomplete Information, and Pass-Through.” *Journal of International Economics*, 101: 168–179.
- Garrett, Daniel G., Eric Ohn, and Juan Carlos Suárez Serrato.** 2020. “Tax Policy and Local Labor Market Behavior.” *American Economic Review: Insights*, 2(1): 83–100.
- Gautier, Erwan, Cristina Conflitti, Riemer P. Faber, Brian Fabo, Ludmila Fadejeva, Valentin Jouvanceau, Jan-Oliver Menz, Teresa Messner, Pavlos Petroulas, Pau Roldan-Blanco, Fabio Rumler, Sergio Santoro, Elisabeth Wieland, and Hélène Zimmer.** 2023. “New Facts on Consumer Price Rigidity in the Euro Area.” *American Economic Journal: Macroeconomics*, forthcoming.

- Gemmi, Luca, and Rosen Valchev.** 2023. “Biased Surveys.” NBER Working Paper 31607.
- Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer.** 2016. “Expectations and Investment.” *NBER Macroeconomics Annual*, 30: 379–431.
- Gertler, Mark, and Peter Karadi.** 2015. “Monetary Policy Surprises, Credit Costs, and Economic Activity.” *American Economic Journal: Macroeconomics*, 7(1): 44–76.
- Ghassibe, Mishel, and Francesco Zanetti.** 2022. “State Dependence of Fiscal Multipliers: The Source of Fluctuations Matters.” *Journal of Monetary Economics*, 132: 1–23.
- Giglio, Stefano, Matteo Maggiori, Johannes Stroebe, and Stephen Utkus.** 2021. “Five Facts about Beliefs and Portfolios.” *American Economic Review*, 111(5): 1481–1522.
- Gilboa, Itzhak, and David Schmeidler.** 1989. “Maxmin Expected Utility with Non-Unique Prior.” *Journal of Mathematical Economics*, 18(2): 141–153.
- Giroud, Xavier, and Joshua Rauh.** 2019. “State Taxation and the Reallocation of Business Activity: Evidence from Establishment-Level Data.” *Journal of Political Economy*, 127(3): 1262–1316.
- Goetzmann, William N., Dasol Kim, and Robert J. Shiller.** 2024. “Emotions and Subjective Crash Beliefs.” National Bureau of Economic Research NBER Working Paper 32589, Cambridge, Mass.
- Golosov, Mikhail, and Robert E Lucas.** 2007. “Menu Costs and Phillips Curves.” *Journal of Political Economy*, 115(2): 171–199.
- Gopinath, Gita, and Oleg Itskhoki.** 2010. “Frequency of Price Adjustment and Pass-Through.” *Quarterly Journal of Economics*, 125(2): 675–727.
- Gopinath, Gita, Oleg Itskhoki, and Roberto Rigobon.** 2010. “Currency Choice and Exchange Rate Pass-Through.” *American Economic Review*, 100(1): 304–336.
- Gorodnichenko, Yuriy.** 2008. “Endogenous Information, Menu Costs and Inflation Persistence.” NBER Working Paper 14184.
- Gourio, François.** 2012. “Disaster Risk and Business Cycles.” *American Economic Review*, 102(6): 2734–2766.

- Graham, John R.** 2022. “Presidential Address: Corporate Finance and Reality.” *Journal of Finance*, 77(4): 1975–2049.
- Grasso, Adriana, and Tiziano Ropele.** 2018. “Firms’ Inflation Expectations and Investment Plans.” Banca d’Italia Working Paper 1203.
- Guceri, Irem, and Maciej Albinowski.** 2021. “Investment Responses to Tax Policy Under Uncertainty.” *Journal of Financial Economics*, 141(3): 1147–1170.
- Guerrieri, Veronica, Guido Lorenzoni, Iván Werning, and Ludwig Straub.** 2022. “Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?” *American Economic Review*, 112(5): 1437–1474.
- Guiso, Luigi, and Giuseppe Parigi.** 1999. “Investment and Demand Uncertainty.” *Quarterly Journal of Economics*, 114(1): 185–227.
- Guvenen, Fatih, Serdar Ozkan, and Jae Song.** 2014. “The Nature of Countercyclical Income Risk.” *Journal of Political Economy*, 122(3): 621–660.
- Gödl-Hanisch, Isabel, and Manuel Menkhoff.** 2022. “Inwieweit geben Unternehmen die gestiegenen Einkaufspreise an ihre Kunden weiter?” *ifo Schnelldienst*, 75(12): 32–34.
- Gödl-Hanisch, Isabel, and Manuel Menkhoff.** 2023. “Firms’ Pass-Through Dynamics: A Survey Approach.” CEPR Discussion Paper 18713.
- Haaland, Ingar, Christopher Roth, and Johannes Wohlfart.** 2023. “Designing Information Provision Experiments.” *Journal of Economic Literature*, 61(1): 3–40.
- Haaland, Ingar, Christopher Roth, Stefanie Stantcheva, and Johannes Wohlfart.** 2024. “Measuring What is Top of Mind.” NBER Working Paper 32421.
- Halevy, Yoram.** 2007. “Ellsberg Revisited: An Experimental Study.” *Econometrica*, 75(2): 503–536.
- Hall, Robert E., and Dale W. Jorgenson.** 1967. “Tax Policy and Investment Behavior.” *American Economic Review*, 57(3): 391–414.
- Handelsblatt.** 2001. “Putin lehnt Bushs Pläne ab.” *Handelsblatt Media Group*, 1.

- Harju, Jarkko, Aliisa Koivisto, and Tuomas Matikka.** 2022. “The Effects of Corporate Taxes on Small Firms.” *Journal of Public Economics*, 212: 104704.
- Hassett, Kevin A., and R. Glenn Hubbard.** 2002. “Tax Policy and Business Investment.” In *Handbook of Public Economics*. Vol. 3, , ed. Alan J. Auerbach and Martin Feldstein, 1293–1343. Elsevier.
- Hayo, Bernd, and Sascha Mierzwa.** 2021. “State-Dependent Effects of Tax Changes in Germany and the United Kingdom.” *MAGKS Papers on Economics*, 25.
- Hazell, Jonathon, Juan Herreno, Emi Nakamura, and Jón Steinsson.** 2022. “The Slope of the Phillips Curve: Evidence from US States.” *Quarterly Journal of Economics*, 137(3): 1299–1344.
- Hendren, Nathaniel, and Ben Sprung-Keyser.** 2020. “A Unified Welfare Analysis of Government Policies.” *Quarterly Journal of Economics*, 135(3): 1209–1318.
- Hennrich, Jonas, Stefan Sauer, and Klaus Wohlrabe.** 2023. “Who Reports the Mood in German Boardrooms? Evidence from the Ifo Business Survey.” CESifo Working Paper 10571.
- Hiersemenzel, Magdolna, Stefan Sauer, and Klaus Wohlrabe.** 2022. “On the Representativeness of the Ifo Business Survey.” CESifo Working Paper 9863.
- Horton, John J.** 2023. “Large Language Models as Simulated Economic Agents: What Can We Learn from Homo Silicus?” NBER Working Paper 31122.
- House, Christopher L., and Matthew D. Shapiro.** 2008. “Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation.” *American Economic Review*, 98(3): 737–768.
- IBS-IND.** 2023. *Ifo Business Survey Industry 1/1980 – 12/2023*. LMU-ifo Economics & Business Data Center, Munich, doi: 10.7805/ebdc-ibs-ind-2023b.
- IBS-SERV.** 2023. *Ifo Business Survey Service 10/2004-12/2023*. LMU-ifo Economics & Business Data Center, Munich, doi: 10.7805/ebdc-ibs-serv-2023b.
- IBS-TRA.** 2023. *Ifo Business Survey Trade 1/1990 – 12/2023*. LMU-ifo Economics & Business Data Center, Munich, doi: 10.7805/ebdc-ibs-tra-2023b.

- Ilut, Cosmin, and Hikaru Saijo.** 2021. “Learning, Confidence, and Business Cycles.” *Journal of Monetary Economics*, 117: 354–376.
- Ilut, Cosmin, and Martin Schneider.** 2014. “Ambiguous Business Cycles.” *American Economic Review*, 104(8): 2368–2399.
- Ilut, Cosmin, and Martin Schneider.** 2022. “Modeling Uncertainty as Ambiguity: A Review.” In *Handbook of Economic Expectations.*, ed. Rüdiger Bachmann, Giorgio Topa and Wilbert van der Klaauw, 749–777. Elsevier.
- IMS.** 2023. *ifo Management-Survey*. ifo Center for Macroeconomics and Surveys.
- Iseringhausen, Martin, Ivan Petrella, and Konstantinos Theodoridis.** 2023. “Aggregate Skewness and the Business Cycle.” *Review of Economics and Statistics*, 1–37.
- Isphording, Ingo, Andreas Lichter, Max Löffler, Thu-Van Nguyen, Felix Pöge, and Sebastian Siegloch.** 2021. “Profit Taxation, R&D Spending, and Innovation.” *CEPR Discussion Paper*, 16702.
- Ivanov, Ivan, Luke Pettit, and Toni M. Whited.** 2024. “Taxes Depress Corporate Borrowing: Evidence from Private Firms.” NBER Working Paper 32398.
- IVS.** 2019. *Ifo Investment Survey Industry 1964-2019*. LMU-ifo Economics & Business Data Center, Munich, doi: 10.7805/ebdc-ivs-ind-2019.
- Jappelli, Tullio, and Luigi Pistaferri.** 2014. “Fiscal Policy and MPC Heterogeneity.” *American Economic Journal: Macroeconomics*, 6(4): 107–136.
- Jappelli, Tullio, and Luigi Pistaferri.** 2020. “Reported MPC and Unobserved Heterogeneity.” *American Economic Journal: Economic Policy*, 12(4): 275–297.
- Jarociński, Marek, and Peter Karadi.** 2020. “Deconstructing Monetary Policy Surprises—The Role of Information Shocks.” *American Economic Journal: Macroeconomics*, 12(2): 1–43.
- Jones, Paul M., Eric Olson, and Mark E. Wohar.** 2015. “Asymmetric Tax Multipliers.” *Journal of Macroeconomics*, 43: 38–48.
- Jordà, Òscar.** 2005. “Estimation and Inference of Impulse Responses by Local Projections.” *American Economic Review*, 95(1): 161–182.

- Joussier, Raphael Lafrogne, Julien Martin, and Isabelle Mejean.** 2022. “Cost Pass-Through and the Rise of Inflation.” Working Paper.
- Juodis, Arturas, and Simas Kucinskas.** 2023. “Quantifying Noise in Survey Expectations.” *Quantitative Economics*, 14(2): 609–650.
- Jäger, Simon, Christopher Roth, Nina Roussille, and Benjamin Schoefer.** 2024. “Worker Beliefs About Outside Options.” *Quarterly Journal of Economics*, 139(3): 1505–1556.
- Karadi, Peter, and Adam Reiff.** 2019. “Menu Costs, Aggregate Fluctuations, and Large Shocks.” *American Economic Journal: Macroeconomics*, 11(3): 111–146.
- Kelly, Bryan, and Hao Jiang.** 2014. “Tail Risk and Asset Prices.” *Review of Financial Studies*, 27(10): 2841–2871.
- Kerssenfischer, Mark, and Maik Schmeling.** 2024. “What Moves Markets?” *Journal of Monetary Economics*, 145: 103560.
- Keynes, Johan Maynard.** 1921. “A Treatise on Probability.” In *The Collected Writings of John Maynard Keynes.*, ed. Elizabeth Johnson and Donald Moggridge. London:Macmillan.
- Kimball, Miles S.** 1995. “The Quantitative Analytics of the Basic Neomonetarist Model.” *Journal of Money, Credit, and Banking*, 27(4): 1241.
- Kim, Gwangmin, and Carola Binder.** 2023. “Learning-Through-Survey in Inflation Expectations.” *American Economic Journal: Macroeconomics*, 15(2): 254–278.
- Klenow, Peter J., and Jonathan L Willis.** 2016. “Real Rigidities and Nominal Price Changes.” *Economica*, 83(331): 443–472.
- Klenow, Peter J., and Oleksiy Kryvtsov.** 2008. “State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?” *Quarterly Journal of Economics*, 123(3): 863–904.
- Knight, H. Frank.** 1921. *Risk, Uncertainty and Profit*. Boston and New York:Houghton Mifflin Company.
- Kohlhas, Alexandre, and Ansgar Walther.** 2021. “Asymmetric Attention.” *American Economic Review*, 111(9): 2879–2925.

- Kohlhas, Alexandre, and Donald Robertson.** 2024. “Cautions Expectations.” Oxford University Mimeo.
- Kozlowski, Julian, Laura Veldkamp, and Venky Venkateswaran.** 2020. “The Tail That Wags the Economy: Beliefs and Persistent Stagnation.” *Journal of Political Economy*, 128(8): 2839–2879.
- Kucinkas, Simas, and Florian S. Peters.** 2022. “Measuring Under- and Overreaction in Expectation Formation.” *Review of Economics and Statistics*, 1–45.
- Kumar, Saten, Yuriy Gorodnichenko, and Olivier Coibion.** 2023. “The Effect of Macroeconomic Uncertainty on Firm Decisions.” *Econometrica*, 91(4): 1297–1332.
- Kydland, Finn E., and Edward C. Prescott.** 1982. “Time to Build and Aggregate Fluctuations.” *Econometrica*, 50(6): 1345–1370.
- Langenmayr, Dominika, and Martin Simmler.** 2021. “Firm Mobility and Jurisdictions’ Tax Rate Choices: Evidence from Immobile Firm Entry.” *Journal of Public Economics*, 204: 104530.
- Langenmayr, Dominika, and Rebecca Lester.** 2017. “Taxation and Corporate Risk-Taking.” *Accounting Review*, 93(3): 237–266.
- Lehmann, Robert.** 2023. “The Forecasting Power of the Ifo Business Survey.” *Journal of Business Cycle Research*, 19(1): 43–94.
- Lerche, Adrian.** 2022. “Investment Tax Credits and the Response of Firms.” *IZA Discussion Paper No. 15668*.
- Lettau, Martin, Matteo Maggiori, and Michael Weber.** 2014. “Conditional Risk Premia in Currency Markets and Other Asset Classes.” *Journal of Financial Economics*, 114(2): 197–225.
- Link, Sebastian, Andreas Peichl, Christopher Roth, and Johannes Wohlfart.** 2023. “Information Frictions Among Firms and Households.” *Journal of Monetary Economics*, 135: 99–115.
- Link, Sebastian, Andreas Peichl, Christopher Roth, and Johannes Wohlfart.** 2024a. “Attention to the Macroeconomy.” CESifo Working Paper 10858.

- Link, Sebastian, Manuel Menkhoff, Andreas Peichl, and Paul Schüle.** 2024b. “Downward Revision of Investment Decisions after Corporate Tax Hikes.” *American Economic Journal: Economic Policy*, 16(4): 194–222.
- Ljungqvist, Alexander, and Michael Smolyansky.** 2018. “To Cut or Not to Cut? On the Impact of Corporate Taxes on Employment and Income.” *NBER Working Paper*, 20753.
- Ljungqvist, Alexander, Liandong Zhang, and Luo Zuo.** 2017. “Sharing Risk with the Government: How Taxes Affect Corporate Risk Taking.” *Journal of Accounting Research*, 55(3): 669–707.
- Longstaff, Francis A., and Monika Piazzesi.** 2004. “Corporate Earnings and the Equity Premium.” *Journal of Financial Economics*, 74(3): 401–421.
- Lorenzoni, Guido.** 2009. “A Theory of Demand Shocks.” *American Economic Review*, 99(5): 2050–2084.
- Loupias, Claire, and Patrick Sevestre.** 2013. “Costs, Demand, and Producer Price Changes.” *Review of Economics and Statistics*, 95(1): 315–327.
- Lucas, Robert E.** 1972. “Expectations and the Neutrality of Money.” *Journal of Economic Theory*, 4(2): 103–124.
- L’Huillier, Jean-Paul, Sanjay R. Singh, and Donghoon Yoo.** 2024. “Incorporating Diagnostic Expectations into the New Keynesian Framework.” *Review of Economic Studies*, 91(5): 3013–3046.
- Mackowiak, Bartosz, and Mirko Wiederholt.** 2009. “Optimal Sticky Prices under Rational Inattention.” *American Economic Review*, 99(3): 769–803.
- Maffini, Giorgia, Jing Xing, and Michael P. Devereux.** 2019. “The Impact of Investment Incentives: Evidence from UK Corporation Tax Returns.” *American Economic Journal: Economic Policy*, 11(3): 361–389.
- Malmendier, Ulrike, and Geoffrey Tate.** 2015. “Behavioral CEOs: The Role of Managerial Overconfidence.” *Journal of Economic Perspectives*, 29(4): 37–60.

- Malmendier, Ulrike, and Stefan Nagel.** 2011. “Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?” *Quarterly Journal of Economics*, 126(1): 373–416.
- Malmendier, Ulrike, and Stefan Nagel.** 2016. “Learning from Inflation Experiences.” *Quarterly Journal of Economics*, 131(1): 53–87.
- Manski, Charles F.** 2018. “Survey Measurement of Probabilistic Macroeconomic Expectations: Progress and Promise.” *NBER Macroeconomics Annual*, 32: 411–471.
- Massenot, Baptiste, and Yuri Pettinicchi.** 2018. “Can Firms See into the Future? Survey Evidence from Germany.” *Journal of Economic Behavior & Organization*, 145: 66–79.
- Maćkowiak, Bartosz, and Mirko Wiederholt.** 2018. “Lack of Preparation for Rare Events.” *Journal of Monetary Economics*, 100: 35–47.
- Melosi, Leonardo.** 2014. “Estimating Models with Dispersed Information.” *American Economic Journal: Macroeconomics*, 6(1): 1–31.
- Mertens, Karel, and Morten O. Ravn.** 2013. “The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States.” *American Economic Review*, 103(4): 1212–1247.
- Midrigan, Virgiliu.** 2011. “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations.” *Econometrica*, 79(4): 1139–1180.
- Moll, Benjamin, Moritz Schularick, and Georg Zachmann.** 2023. “The Power of Substitution: The Great German Gas Debate in Retrospect.” *Brookings Papers on Economic Activity*, Fall 2023.
- Montiel Olea, José Luis, and Mikkel Plagborg-Møller.** 2021. “Local Projection Inference Is Simpler and More Robust Than You Think.” *Econometrica*, 89(4): 1789–1823.
- Mortensen, Dale, and Christopher Pissarides.** 2009. “Job Creation and Job Destruction in the Theory of Unemployment.” *Review of Economic Studies*, 63(2): 397–415.

- Nakamura, Emi, and Dawit Zerom.** 2010. “Accounting for Incomplete Pass-Through.” *Review of Economic Studies*, 77(3): 1192–1230.
- Nakamura, Emi, and Jón Steinsson.** 2008. “Five Facts about Prices: A Reevaluation of Menu Cost Models.” *Quarterly Journal of Economics*, 123(4): 1415–1464.
- Nakamura, Emi, and Jón Steinsson.** 2010. “Monetary Non-Neutrality in a Multisector Model.” *Quarterly Journal of Economics*, 125(3): 961–1013.
- Nakamura, Emi, and Jón Steinsson.** 2013. “Price Rigidity: Microeconomic Evidence and Macroeconomic Implications.” *Annual Review of Economics*, 5(1): 133–163.
- Nakamura, Emi, and Jón Steinsson.** 2014. “Fiscal Stimulus in a Monetary Union: Evidence from US Regions.” *American Economic Review*, 104(3): 753–792.
- Nakamura, Emi, Jón Steinsson, Robert Barro, and José Ursúa.** 2013. “Crises and Recoveries in an Empirical Model of Consumption Disasters.” *American Economic Journal: Macroeconomics*, 5(3): 35–74.
- Neiman, Brent.** 2010. “Stickiness, Synchronization, and Passthrough in Intrafirm Trade Prices.” *Journal of Monetary Economics*, 57: 295–308.
- Nicodème, Gaetan.** 2001. “Computing Effective Corporate Tax Rates: Comparisons and Results.” *MPRA Working Paper*, 3808.
- Ohrn, Eric.** 2018. “The Effect of Corporate Taxation on Investment and Financial Policy: Evidence from the DPAD.” *American Economic Journal: Economic Policy*, 10(2): 272–301.
- Orlik, Anna, and Laura Veldkamp.** 2022. “Understanding Uncertainty Shocks and the Role of Black Swans.” Working Paper.
- Phelps, Edmund S.** 1970. “Introduction: The New Microeconomics in Employment and Inflation Theory.” In *Microeconomic Foundations of Employment and Inflation Theory*. New York: Norton.
- Phillips, A. W.** 1958. “The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–19571.” *Economica*, 25(100): 283–299.

- Ramey, Valerie A, and Sarah Zubairy.** 2018. “Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data.” *Journal of Political Economy*, 126(2): 850–901.
- Riedel, Nadine, and Martin Simmler.** 2021. “Large and Influential: Firm Size and Governments’ Corporate Tax Rate Choice.” *Canadian Journal of Economics/Revue Canadienne d’Économie*, 54(2): 812–839.
- Rietz, Thomas A.** 1988. “The Equity Risk Premium a Solution.” *Journal of Monetary Economics*, 22(1): 117–131.
- Robinson, Sarah, and Alisa Tazhitdinova.** 2022. “What Drives Tax Policy? Political, Institutional and Economic Determinants of State Tax Policy.” *NBER Working Paper*, 31268.
- Romer, Christina D., and David H. Romer.** 2004. “A New Measure of Monetary Shocks: Derivation and Implications.” *American Economic Review*, 94(4): 1055–1084.
- Romer, Christina D., and David H. Romer.** 2010. “The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks.” *American Economic Review*, 100(3): 763–801.
- Ropele, Tiziano, Yuriy Gorodnichenko, and Olivier Coibion.** 2022. “Inflation Expectations and Corporate Borrowing Decisions: New Causal Evidence.” NBER Working Paper 30537.
- Ropele, Tiziano, Yuriy Gorodnichenko, and Olivier Coibion.** 2024. “Inflation Expectations and Misallocation of Resources: Evidence from Italy.” *American Economic Review: Insights*, 6(2): 246–261.
- Rosolia, Alfonso.** 2021. “Does Information About Current Inflation Affect Expectations and Decisions? Another Look at Italian Firms.” Banca d’Italia Working Paper 1353.
- Roth, Christopher, and Johannes Wohlfart.** 2020. “How Do Expectations about the Macroeconomy Affect Personal Expectations and Behavior?” *Review of Economics and Statistics*, 102(4): 731–748.
- Roth, Christopher, Mirko Wiederholt, and Johannes Wohlfart.** 2023. “The Effects of Monetary Policy: Theory with Measured Expectations.” CESifo Working Paper 10216.

- Salgado, Sergio, Fatih Guvenen, and Nicholas Bloom.** 2023. “Skewed Business Cycles.” NBER Working Paper 26565.
- Sauer, Stefan, and Klaus Wohlrabe.** 2019. “CEO or Intern – Who Actually Answers the Questionnaires in the Ifo Business Survey?” *CESifo Forum*, 20(2): 29–31.
- Sauer, Stefan, Moritz Schasching, and Klaus Wohlrabe.** 2023. *Handbook of ifo Surveys. ifo Beiträge zur Wirtschaftsforschung*.
- Savignac, Frédérique, Erwan Gautier, Yuriy Gorodnichenko, and Olivier Coibion.** 2024. “Firms’ Inflation Expectations: New Evidence from France.” *Journal of the European Economic Association*, 22(6): 2748–2781.
- Shiller, Robert J.** 2000. “Measuring Bubble Expectations and Investor Confidence.” *Journal of Psychology and Financial Markets*, 1(1): 49–60.
- Shiller, Robert J.** 2017. “Narrative Economics.” *American Economic Review*, 107(4): 967–1004.
- SIGE.** 2023. *Survey on Inflation and Growth Expectations*. Banca d’Italia.
- Smets, Frank, and Rafael Wouters.** 2003. “An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area.” *Journal of the European Economic Association*, 1(5): 1123–1175.
- Smets, Frank, and Rafael Wouters.** 2007. “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach.” *American Economic Review*, 97(3): 586–606.
- Smith, V. Kerry, Jr. Taylor, Donald H., Frank A. Sloan, F. Reed Johnson, and William H. Desvousges.** 2001. “Do Smokers Respond to Health Shocks?” *Review of Economics and Statistics*, 83(4): 675–687.
- Stantcheva, Stefanie.** 2023. “How to Run Surveys: A Guide to Creating Your Own Identifying Variation and Revealing the Invisible.” *Annual Review of Economics*, 15: 205–234.
- Sun, Liyang, and Sarah Abraham.** 2021. “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects.” *Journal of Econometrics*, 225(2): 175–199.

- Suárez Serrato, Juan Carlos, and Owen Zidar.** 2016. “Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms.” *American Economic Review*, 106(9): 2582–2624.
- Tauchen, George.** 1986. “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions.” *Economics Letters*, 20(2): 177–181.
- Taylor, Shelley E., and Suzanne C. Thompson.** 1982. “Stalking the Elusive “Vividness” Effect.” *Psychological Review*, , (2): 155–181.
- Theil, Henri.** 1955. “Recent Experiences with the Munich Business Test: An Expository Article.” *Econometrica*, 23(2): 184–192.
- Trautmann, Stefan T., and Gijs Van De Kuilen.** 2015. “Ambiguity Attitudes.” In *The Wiley Blackwell Handbook of Judgment and Decision Making*, ed. Gideon Keren and George Wu, 89–116. Wiley.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin.** 2017. “Attention is All you Need.” Vol. 30. Curran Associates, Inc.
- Vavra, Joseph.** 2014. “Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation.” *Quarterly Journal of Economics*, 129(1): 215–258.
- Viscusi, Kip, and Richard Zeckhauser.** 2015. “The Relative Weights of Direct and Indirect Experiences in the Formation of Environmental Risk Beliefs.” *Risk Analysis*, 35(2): 318–331.
- Wachter, Jessica A.** 2013. “Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?” *Journal of Finance*, 68(3): 987–1035.
- Wang, Olivier, and Iván Werning.** 2022. “Dynamic Oligopoly and Price Stickiness.” *American Economic Review*, 112(8): 2815–2849.
- Winberry, Thomas.** 2021. “Lumpy Investment, Business Cycles, and Stimulus Policy.” *American Economic Review*, 111(1): 364–396.
- Wiswall, Matthew, and Basit Zafar.** 2021. “Human Capital Investments and Expectations about Career and Family.” *Journal of Political Economy*, 129(5): 1361–1424.
- Woodford, Michael.** 2003. *Interest and Prices*. Princeton University Press.

- Woodford, Michael.** 2010. “Robustly Optimal Monetary Policy with Near-Rational Expectations.” *American Economic Review*, 100(1): 274–303.
- Woodford, Michael, and Julio J Rotemberg.** 1997. “An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy.” *NBER Macroeconomics Annual*, 12: 297–346.
- Xu, Qiping, and Eric Zwick.** 2022. “Tax Policy and Abnormal Investment Behavior.” NBER Working Paper 27363.
- Yagi, Tomoyuki, Yoshiyuki Kurachi, Masato Takahashi, Kotone Yamada, and Hiroshi Kawata.** 2022. “Pass-Through of Cost-Push Pressures to Consumer Prices.” Bank of Japan Working Paper.
- Zwick, Eric, and James Mahon.** 2017. “Tax Policy and Heterogeneous Investment Behavior.” *American Economic Review*, 107(1): 217–248.